



Junos[®] OS

Routing Protocols and Policies Command Reference

Release
12.3



Published: 2013-01-30

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

This product includes the Envoy SNMP Engine, developed by Epilogue Technology, an Integrated Systems Company. Copyright © 1986-1997, Epilogue Technology Corporation. All rights reserved. This program and its documentation were developed at private expense, and no part of them is in the public domain.

This product includes memory allocation software developed by Mark Moraes, copyright © 1988, 1989, 1993, University of Toronto.

This product includes FreeBSD software developed by the University of California, Berkeley, and its contributors. All of the documentation and software included in the 4.4BSD and 4.4BSD-Lite Releases is copyrighted by the Regents of the University of California. Copyright © 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994. The Regents of the University of California. All rights reserved.

GateD software copyright © 1995, the Regents of the University. All rights reserved. Gate Daemon was originated and developed through release 3.0 by Cornell University and its collaborators. Gated is based on Kirton's EGP, UC Berkeley's routing daemon (routed), and DCN's HELLO routing protocol. Development of Gated has been supported in part by the National Science Foundation. Portions of the GateD software copyright © 1988, Regents of the University of California. All rights reserved. Portions of the GateD software copyright © 1991, D. L. S. Associates.

This product includes software developed by Maker Communications, Inc., copyright © 1996, 1997, Maker Communications, Inc.

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

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

Products made or sold by Juniper Networks or components thereof might be covered by one or more of the following patents that are owned by or licensed to Juniper Networks: U.S. Patent Nos. 5,473,599, 5,905,725, 5,909,440, 6,192,051, 6,333,650, 6,359,479, 6,406,312, 6,429,706, 6,459,579, 6,493,347, 6,538,518, 6,538,899, 6,552,918, 6,567,902, 6,578,186, and 6,590,785.

Junos® OS Routing Protocols and Policies Command Reference

Release 12.3R1

Copyright © 2013, Juniper Networks, Inc.

All rights reserved.

Revision History

January 2013—R1 Junos OS 12.3

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Abbreviated Table of Contents

| | | |
|------------|--|------|
| | About This Guide | xiii |
| Part 1 | Protocols | |
| Chapter 1 | ANCP Operational Mode Commands | 3 |
| Chapter 2 | BFD Operational Mode Commands | 25 |
| Chapter 3 | BGP Operational Mode Commands | 37 |
| Chapter 4 | ES-IS Operational Mode Commands | 77 |
| Chapter 5 | IP Multicast Operational Mode Commands | 87 |
| Chapter 6 | IPv6 Operational Mode Commands | 301 |
| Chapter 7 | IS-IS Operational Mode Commands | 309 |
| Chapter 8 | LLDP Operational Mode Commands | 363 |
| Chapter 9 | MVRP Operational Mode Commands | 379 |
| Chapter 10 | OSPF Operational Mode Commands | 391 |
| Chapter 11 | Protocol-Independent Routing Operational Mode Commands | 475 |
| Chapter 12 | PTP Operational Mode Commands | 677 |
| Chapter 13 | RIP Operational Mode Commands | 691 |
| Chapter 14 | RIPng Operational Mode Commands | 701 |
| Part 2 | Policy Framework | |
| Chapter 15 | Firewall Filter Operational Mode Commands | 711 |
| Chapter 16 | Forwarding Operational Mode Commands | 731 |
| Chapter 17 | Routing Policy Operational Mode Commands | 765 |
| Part 3 | MPLS | |
| Chapter 18 | LDP Operational Mode Commands | 775 |
| Chapter 19 | MPLS Operational Mode Commands | 807 |
| Chapter 20 | RSVP Operational Mode Commands | 865 |
| Part 4 | Layer 2 Bridging and Switching Operational Mode Commands | |
| Chapter 21 | Layer 2 Bridging and Switching Operational Mode Commands | 899 |
| Chapter 22 | Spanning Tree Operational Mode Commands | 927 |
| Part 5 | VPNs | |
| Chapter 23 | VPN Operational Mode Commands | 947 |

Part 6

Index

| | |
|--|------|
| Index | 1007 |
| Index of Statements and Commands | 1019 |

Table of Contents

| | | |
|------------------|--|-------------|
| | About This Guide | xiii |
| | Junos OS Documentation and Release Notes | xiii |
| | Objectives | xiv |
| | Audience | xiv |
| | Supported Platforms | xv |
| | Using the Indexes | xv |
| | Documentation Conventions | xv |
| | Documentation Feedback | xvii |
| | Requesting Technical Support | xvii |
| | Self-Help Online Tools and Resources | xviii |
| | Opening a Case with JTAC | xviii |
| Part 1 | Protocols | |
| Chapter 1 | ANCP Operational Mode Commands | 3 |
| | clear ancp neighbor | 4 |
| | clear ancp subscriber | 6 |
| | request ancp oam interface | 8 |
| | request ancp oam neighbor | 9 |
| | show ancp cos | 10 |
| | show ancp neighbor | 13 |
| | show ancp subscriber | 19 |
| Chapter 2 | BFD Operational Mode Commands | 25 |
| | clear bfd adaptation | 26 |
| | clear bfd session | 27 |
| | show bfd session | 28 |
| Chapter 3 | BGP Operational Mode Commands | 37 |
| | clear bgp damping | 38 |
| | clear bgp neighbor | 39 |
| | clear bgp table | 41 |
| | show bgp bmp | 43 |
| | show bgp group | 44 |
| | show bgp group traffic-statistics | 51 |
| | show bgp neighbor | 53 |
| | show bgp replication | 67 |
| | show bgp summary | 69 |
| | show policy damping | 74 |

| | | |
|------------------|---|-----------|
| Chapter 4 | ES-IS Operational Mode Commands | 77 |
| | clear esis adjacency | 78 |
| | clear esis statistics | 79 |
| | show esis adjacency | 80 |
| | show esis interface | 82 |
| | show esis statistics | 84 |
| Chapter 5 | IP Multicast Operational Mode Commands | 87 |
| | clear igmp membership | 92 |
| | clear igmp snooping membership | 95 |
| | clear igmp snooping statistics | 96 |
| | clear igmp statistics | 97 |
| | clear mld membership | 99 |
| | clear mld statistics | 100 |
| | clear msdp cache | 101 |
| | clear msdp statistics | 102 |
| | clear multicast bandwidth-admission | 103 |
| | clear multicast forwarding-cache | 105 |
| | clear multicast scope | 106 |
| | clear multicast sessions | 107 |
| | clear multicast snooping statistics | 108 |
| | clear multicast statistics | 109 |
| | clear pgm negative-acknowledgments | 110 |
| | clear pgm source-path-messages | 111 |
| | clear pgm statistics | 112 |
| | clear pim join | 113 |
| | clear pim join-distribution | 114 |
| | clear pim register | 116 |
| | clear pim snooping join | 118 |
| | clear pim snooping statistics | 120 |
| | clear pim statistics | 122 |
| | request pim multicast-tunnel rebalance | 125 |
| | show dvmrp interfaces | 126 |
| | show dvmrp neighbors | 128 |
| | show dvmrp prefix | 130 |
| | show dvmrp prunes | 132 |
| | show igmp group | 134 |
| | show igmp interface | 138 |
| | show igmp snooping interface | 142 |
| | show igmp snooping membership | 145 |
| | show igmp snooping statistics | 149 |
| | show igmp statistics | 152 |
| | show mld group | 155 |
| | show mld interface | 160 |
| | show mld statistics | 164 |
| | show msdp | 167 |
| | show msdp source | 169 |
| | show msdp source-active | 171 |
| | show msdp statistics | 174 |

| | | |
|------------------|--|------------|
| | show multicast backup-pe-groups | 178 |
| | show multicast flow-map | 180 |
| | show multicast forwarding-cache statistics | 182 |
| | show multicast interface | 184 |
| | show multicast mrinfo | 186 |
| | show multicast next-hops | 188 |
| | show multicast pim-to-igmp-proxy | 191 |
| | show multicast pim-to-mld-proxy | 193 |
| | show multicast route | 195 |
| | show multicast rpf | 201 |
| | show multicast scope | 205 |
| | show multicast sessions | 207 |
| | show multicast snooping next-hops | 210 |
| | show multicast snooping route | 212 |
| | show multicast snooping statistics | 215 |
| | show multicast statistics | 218 |
| | show multicast usage | 221 |
| | show pgm negative-acknowledgments | 224 |
| | show pgm source-path-messages | 226 |
| | show pgm statistics | 227 |
| | show pim bootstrap | 230 |
| | show pim interfaces | 232 |
| | show pim join | 235 |
| | show pim mdt | 244 |
| | show pim mdt data-mdt-joins | 248 |
| | show pim mdt data-mdt-limit | 250 |
| | show pim mvpn | 252 |
| | show pim neighbors | 253 |
| | show pim rps | 258 |
| | show pim snooping interfaces | 265 |
| | show pim snooping join | 269 |
| | show pim snooping neighbors | 273 |
| | show pim snooping statistics | 278 |
| | show pim source | 283 |
| | show pim statistics | 285 |
| | show sap listen | 298 |
| | test msdp | 299 |
| Chapter 6 | IPv6 Operational Mode Commands | 301 |
| | clear ipv6 neighbors | 302 |
| | clear ipv6 router-advertisement | 303 |
| | show ipv6 neighbors | 304 |
| | show ipv6 router-advertisement | 306 |
| Chapter 7 | IS-IS Operational Mode Commands | 309 |
| | clear isis adjacency | 311 |
| | clear isis database | 313 |
| | clear isis overload | 315 |
| | clear isis statistics | 317 |
| | show isis adjacency | 319 |

| | | |
|-------------------|--|------------|
| | show isis authentication | 323 |
| | show isis backup coverage | 325 |
| | show isis backup label-switched-path | 327 |
| | show isis backup spf results | 329 |
| | show isis context-identifier | 333 |
| | show isis database | 335 |
| | show isis hostname | 342 |
| | show isis interface | 343 |
| | show isis overview | 348 |
| | show isis route | 351 |
| | show isis spf | 355 |
| | show isis statistics | 360 |
| Chapter 8 | LLDP Operational Mode Commands | 363 |
| | clear lldp neighbor | 364 |
| | clear lldp statistics | 365 |
| | show lldp | 366 |
| | show lldp local-information | 369 |
| | show lldp neighbors | 371 |
| | show lldp remote-global-statistics | 375 |
| | show lldp statistics | 377 |
| Chapter 9 | MVRP Operational Mode Commands | 379 |
| | show mvrp | 380 |
| | show mvrp applicant-state | 382 |
| | show mvrp dynamic-vlan-memberships | 384 |
| | show mvrp interface | 385 |
| | show mvrp registration-state | 386 |
| | show mvrp statistics | 388 |
| Chapter 10 | OSPF Operational Mode Commands | 391 |
| | clear (ospf ospf3) database | 393 |
| | clear (ospf ospf3) database-protection | 396 |
| | clear (ospf ospf3) io-statistics | 397 |
| | clear (ospf ospf3) neighbor | 398 |
| | clear (ospf ospf3) overload | 400 |
| | clear (ospf ospf3) statistics | 401 |
| | show (ospf ospf3) backup coverage | 403 |
| | show (ospf ospf3) backup lsp | 406 |
| | show (ospf ospf3) backup neighbor | 408 |
| | show (ospf ospf3) backup spf | 410 |
| | show ospf context-identifier | 418 |
| | show ospf database | 420 |
| | show ospf3 database | 429 |
| | show (ospf ospf3) interface | 440 |
| | show (ospf ospf3) io-statistics | 446 |
| | show (ospf ospf3) log | 448 |
| | show (ospf ospf3) neighbor | 452 |
| | show (ospf ospf3) overview | 459 |
| | show (ospf ospf3) route | 464 |

| | | |
|-------------------|---|------------|
| Chapter 11 | show (ospf ospf3) statistics | 470 |
| | Protocol-Independent Routing Operational Mode Commands | 475 |
| | show as-path | 478 |
| | show as-path domain | 482 |
| | show as-path summary | 484 |
| | show route | 486 |
| | show route active-path | 492 |
| | show route advertising-protocol | 497 |
| | show route all | 502 |
| | show route aspath-regex | 504 |
| | show route best | 506 |
| | show route brief | 509 |
| | show route ccc | 511 |
| | show route community | 512 |
| | show route community-name | 514 |
| | show route damping | 516 |
| | show route detail | 522 |
| | show route exact | 537 |
| | show route export | 539 |
| | show route export vrf-target | 541 |
| | show route extensive | 543 |
| | show route flow validation | 558 |
| | show route forwarding-table | 560 |
| | show route forwarding-table interface-name | 574 |
| | show route hidden | 578 |
| | show route inactive-path | 581 |
| | show route inactive-prefix | 585 |
| | show route instance | 587 |
| | show route label | 594 |
| | show route label-switched-path | 596 |
| | show route localization | 598 |
| | show route martians | 600 |
| | show route next-hop | 602 |
| | show route no-community | 608 |
| | show route output | 611 |
| | show route protocol | 615 |
| | show route range | 627 |
| | show route receive-protocol | 632 |
| | show route resolution | 641 |
| | show route snooping | 644 |
| | show route source-gateway | 652 |
| | show route summary | 658 |
| | show route table | 662 |
| | show route terse | 673 |
| Chapter 12 | PTP Operational Mode Commands | 677 |
| | show ptp clock | 678 |
| | show ptp hybrid | 681 |
| | show ptp lock-status | 683 |

| | | |
|-------------------|--|------------|
| | show ptp master | 685 |
| | show ptp port | 687 |
| | show ptp slave | 689 |
| Chapter 13 | RIP Operational Mode Commands | 691 |
| | clear rip general-statistics | 692 |
| | clear rip statistics | 693 |
| | show rip general-statistics | 694 |
| | show rip neighbor | 696 |
| | show rip statistics | 698 |
| Chapter 14 | RIPng Operational Mode Commands | 701 |
| | clear ripng general-statistics | 702 |
| | clear ripng statistics | 703 |
| | show ripng general-statistics | 704 |
| | show ripng neighbor | 705 |
| | show ripng statistics | 707 |
| Part 2 | Policy Framework | |
| Chapter 15 | Firewall Filter Operational Mode Commands | 711 |
| | clear firewall | 712 |
| | show firewall | 714 |
| | show firewall filter version | 721 |
| | show firewall log | 722 |
| | show firewall prefix-action-stats | 725 |
| | show firewall templates-in-use | 726 |
| | show policer | 728 |
| Chapter 16 | Forwarding Operational Mode Commands | 731 |
| | clear dhcp relay binding | 732 |
| | clear dhcp relay statistics | 734 |
| | clear dhcpv6 relay binding | 737 |
| | clear dhcpv6 relay statistics | 741 |
| | clear helper statistics | 743 |
| | show dhcp relay binding | 745 |
| | show dhcp relay statistics | 750 |
| | show dhcpv6 relay binding | 753 |
| | show dhcpv6 relay statistics | 759 |
| | show helper statistics | 762 |
| Chapter 17 | Routing Policy Operational Mode Commands | 765 |
| | show policy | 766 |
| | show policy conditions | 768 |
| | test policy | 770 |
| Part 3 | MPLS | |
| Chapter 18 | LDP Operational Mode Commands | 775 |
| | clear ldp neighbor | 777 |
| | clear ldp session | 778 |

| | | |
|-------------------|---|------------|
| | clear ldp statistics | 779 |
| | show ldp database | 780 |
| | show ldp fec-filters | 784 |
| | show ldp interface | 785 |
| | show ldp neighbor | 787 |
| | show ldp path | 789 |
| | show ldp route | 791 |
| | show ldp session | 795 |
| | show ldp statistics | 801 |
| | show ldp traffic-statistics | 805 |
| Chapter 19 | MPLS Operational Mode Commands | 807 |
| | clear mpls lsp | 809 |
| | request mpls lsp adjust-autobandwidth | 811 |
| | show connections | 813 |
| | show link-management | 816 |
| | show link-management peer | 820 |
| | show link-management routing | 822 |
| | show link-management statistics | 825 |
| | show link-management te-link | 827 |
| | show mpls admin-groups | 829 |
| | show mpls call-admission-control | 830 |
| | show mpls cspf | 832 |
| | show mpls diffserv-te | 834 |
| | show mpls interface | 836 |
| | show mpls lsp | 838 |
| | show mpls path | 851 |
| | show mpls srlg | 852 |
| | show mpls static-lsp | 853 |
| | show ted database | 856 |
| | show ted link | 861 |
| | show ted protocol | 863 |
| Chapter 20 | RSVP Operational Mode Commands | 865 |
| | clear rsvp session | 866 |
| | clear rsvp statistics | 868 |
| | show rsvp interface | 869 |
| | show rsvp neighbor | 875 |
| | show rsvp session | 880 |
| | show rsvp statistics | 890 |
| | show rsvp version | 894 |
| Part 4 | Layer 2 Bridging and Switching Operational Mode Commands | |
| Chapter 21 | Layer 2 Bridging and Switching Operational Mode Commands | 899 |
| | clear bridge mac-table | 900 |
| | clear error bpdu | 901 |
| | clear error mac-rewrite | 902 |
| | show bridge domain | 903 |
| | show bridge flood | 905 |

| | | |
|-------------------|--|------------|
| | show bridge mac-table | 912 |
| | show bridge statistics | 916 |
| | show l2-learning global-information | 920 |
| | show l2-learning global-mac-count | 921 |
| | show l2-learning instance | 922 |
| | show l2-learning interface | 924 |
| | show mac-rewrite interface | 926 |
| Chapter 22 | Spanning Tree Operational Mode Commands | 927 |
| | clear spanning-tree protocol-migration | 928 |
| | clear spanning-tree statistics | 929 |
| | show spanning-tree bridge | 930 |
| | show spanning-tree interface | 935 |
| | show spanning-tree mstp configuration | 941 |
| | show spanning-tree statistics | 943 |
| Part 5 | VPNs | |
| Chapter 23 | VPN Operational Mode Commands | 947 |
| | clear vpls mac-address | 949 |
| | clear vpls mac-table | 950 |
| | request l2circuit-switchover | 951 |
| | show dynamic-tunnels database | 952 |
| | show hfr profiles | 954 |
| | show ingress-replication mvpn | 956 |
| | show l2circuit connections | 957 |
| | show l2vpn connections | 963 |
| | show mvpn c-multicast | 969 |
| | show mvpn instance | 971 |
| | show mvpn neighbor | 975 |
| | show vpls connections | 979 |
| | show vpls flood event-queue | 990 |
| | show vpls flood instance | 992 |
| | show vpls flood route | 994 |
| | show vpls mac-table | 996 |
| | show vpls statistics | 1001 |
| Part 6 | Index | |
| | Index | 1007 |
| | Index of Statements and Commands | 1019 |

About This Guide

This preface provides the following guidelines for using the *Junos[®] OS Routing Protocols and Policies Command Reference*:

- Junos OS Documentation and Release Notes on page xiii
- Objectives on page xiv
- Audience on page xiv
- Supported Platforms on page xv
- Using the Indexes on page xv
- Documentation Conventions on page xv
- Documentation Feedback on page xvii
- Requesting Technical Support on page xvii

Junos OS Documentation and Release Notes

For a list of related Junos OS documentation, see <http://www.juniper.net/techpubs/software/junos/>.

If the information in the latest release notes differs from the information in the documentation, follow the *Junos OS Release Notes*.

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

Juniper Networks supports a technical book program to publish books by Juniper Networks engineers and subject matter experts with book publishers around the world. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration using the Junos operating system (Junos OS) and Juniper Networks devices. In addition, the Juniper Networks Technical Library, published in conjunction with O'Reilly Media, explores improving network security, reliability, and availability using Junos OS configuration techniques. All the books are for sale at technical bookstores and book outlets around the world. The current list can be viewed at <http://www.juniper.net/books>.

Objectives

This guide provides descriptions of the Juniper Networks Junos OS commands that you use to monitor and troubleshoot routing protocols, protocol-independent features, and policies, including firewall filters, forwarding options, and routing policies.

- Junos OS Operational Mode Commands
- Junos OS Operational Mode Commands



NOTE: For additional information about the Junos OS—either corrections to or information that might have been omitted from this guide—see the software release notes at <http://www.juniper.net/>.

For information about configuration statements and guidelines related to the commands described in this reference, see the following configuration guides:

- Junos OS Routing Protocols Configuration Guide—Includes configuration statements and guidelines for routing protocols and protocol-independent features.
- Routing Policy Configuration Guide—Includes configuration statements and guidelines for policies, including firewall filters, forwarding options, and routing policies.
- Junos OS MPLS Applications Configuration Guide—Includes configuration statements and guidelines for Multiprotocol Label Switching (MPLS) traffic engineering.
- Junos OS VPNs Configuration Guide—Includes configuration statements and guidelines for Layer 2 and Layer 3 virtual private networks (VPNs), virtual private LAN service (VPLS), and Layer 2 circuits.

For information about MPLS-related tasks performed by network operations center (NOC) personnel, see the *Junos MPLS Network Operations Guide*.

Audience

This guide is designed for network administrators who are configuring and monitoring a Juniper Networks M Series, MX Series, T Series, EX Series, or J Series router or switch.

To use this guide, you need a broad understanding of networks in general, the Internet in particular, networking principles, and network configuration. You must also be familiar with one or more of the following Internet routing protocols:

- Border Gateway Protocol (BGP)
- Distance Vector Multicast Routing Protocol (DVMRP)
- Intermediate System-to-Intermediate System (IS-IS)
- Internet Control Message Protocol (ICMP) router discovery
- Internet Group Management Protocol (IGMP)

- Multiprotocol Label Switching (MPLS)
- Open Shortest Path First (OSPF)
- Protocol-Independent Multicast (PIM)
- Resource Reservation Protocol (RSVP)
- Routing Information Protocol (RIP)
- Simple Network Management Protocol (SNMP)

Personnel operating the equipment must be trained and competent; must not conduct themselves in a careless, willfully negligent, or hostile manner; and must abide by the instructions provided by the documentation.

Supported Platforms

For the features described in this manual, the Junos OS currently supports the following platforms:

- J Series
- M Series
- MX Series
- SRX Series
- T Series
- EX Series

Using the Indexes

This reference contains two indexes: a standard index with topic entries, and an index of commands.

Documentation Conventions

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

Table 1: Notice Icons

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

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

Table 2: Text and Syntax Conventions

| Convention | Description | Examples |
|------------------------------|--|---|
| Bold text like this | Represents text that you type. | To enter configuration mode, type the configure command: user@host> configure |
| Fixed-width text like this | Represents output that appears on the terminal screen. | user@host> show chassis alarms No alarms currently active |
| <i>Italic text like this</i> | <ul style="list-style-type: none"> Introduces or emphasizes important new terms. Identifies book names. Identifies RFC and Internet draft titles. | <ul style="list-style-type: none"> A policy <i>term</i> is a named structure that defines match conditions and actions. <i>Junos OS System Basics Configuration Guide</i> RFC 1997, <i>BGP Communities Attribute</i> |
| <i>Italic text like this</i> | Represents variables (options for which you substitute a value) in commands or configuration statements. | Configure the machine's domain name: [edit] root@# set system domain-name <i>domain-name</i> |
| Text like this | Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components. | <ul style="list-style-type: none"> To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level. The console port is labeled CONSOLE. |
| < > (angle brackets) | Enclose optional keywords or variables. | stub <default-metric <i>metric</i> >; |

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

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

Documentation Feedback

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

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

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or JNASC support contract,

or are covered under warranty, and need postsales technical support, you can access our tools and resources online or open a case with JTAC.

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

Self-Help Online Tools and Resources

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

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

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

Opening a Case with JTAC

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

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

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

PART 1

Protocols

- [ANCP Operational Mode Commands on page 3](#)
- [BFD Operational Mode Commands on page 25](#)
- [BGP Operational Mode Commands on page 37](#)
- [ES-IS Operational Mode Commands on page 77](#)
- [IP Multicast Operational Mode Commands on page 87](#)
- [IPv6 Operational Mode Commands on page 301](#)
- [IS-IS Operational Mode Commands on page 309](#)
- [LLDP Operational Mode Commands on page 363](#)
- [MVRP Operational Mode Commands on page 379](#)
- [OSPF Operational Mode Commands on page 391](#)
- [Protocol-Independent Routing Operational Mode Commands on page 475](#)
- [PTP Operational Mode Commands on page 677](#)
- [RIP Operational Mode Commands on page 691](#)
- [RIPng Operational Mode Commands on page 701](#)

CHAPTER 1

ANCP Operational Mode Commands

Table 3 on page 3 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot Access Node Control Protocol (ANCP) operations. Commands are listed in alphabetical order.

Table 3: ANCP Operational Mode Commands

| Task | Command |
|---|---|
| Clear ANCP neighbors. | <code>clear ancp neighbor</code> |
| Clear ANCP subscriber connections. | <code>clear ancp subscriber</code> |
| Trigger the access node to run a loopback test on the local loop specified by an ANCP interface or interface set. | <code>request ancp oam interface</code> |
| Trigger the access node to run a loopback test on the local loop specified by an ANCP neighbor. | <code>request ancp oam neighbor</code> |
| Display ANCP class-of-service information. | <code>show ancp cos</code> |
| Display ANCP neighbor information. | <code>show ancp neighbor</code> |
| Display ANCP subscriber information. | <code>show ancp subscriber</code> |



NOTE: For information about how to configure ANCP, see the *Junos Subscriber Access Configuration Guide*.

clear ancp neighbor

| | |
|--------------------------|---|
| Syntax | <code>clear ancp neighbor</code> <code><ip-address <i>ip-address</i>></code> <code><system-name <i>mac-address</i>></code> |
| Release Information | Command introduced in Junos OS Release 9.4. |
| Description | Clear the connection with all ANCP neighbors or with the specified ANCP neighbor. |
| Options | none —Clear all ANCP neighbors. ip-address <i>ip-address</i> —(Optional) Clear the ANCP neighbor specified by the IP address. system-name <i>mac-address</i> —(Optional) Clear the ANCP neighbor specified by the MAC address. |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• show ancp neighbor on page 13 |
| List of Sample Output | clear ancp neighbor on page 5 show ancp neighbor on page 5 |
| Output Fields | When you enter this command, you are provided no feedback on the status of your request. You can enter the show ancp neighbor command before and after clearing the ANCP neighbors to verify the clear operation. |

Sample Output

clear ancp neighbor user@host> clear ancp neighbor

show ancp neighbor The following sample output displays the connections with ANCP neighbors before and after the **clear ancp neighbor** command was issued.

user@host> show ancp neighbor

| IP Address | MAC Address | State | Subscriber Count | Capabilities |
|------------|-------------------|-------------|------------------|--------------|
| 10.10.10.2 | ba:ad:be:ef:10:10 | Established | 5 | Topo |
| 12.12.12.2 | ba:ad:be:ef:10:12 | Established | 5 | Topo |
| 13.13.13.2 | ba:ad:be:ef:10:13 | Established | 5 | Topo |
| 14.14.14.2 | ba:ad:be:ef:10:14 | Established | 5 | Topo |

user@host> clear ancp neighbor ip-address 10.10.10.2

user@host> show ancp neighbor

| IP Address | MAC Address | State | Subscriber Count | Capabilities |
|------------|-------------------|-------------|------------------|--------------|
| 12.12.12.2 | ba:ad:be:ef:10:12 | Established | 5 | Topo |
| 13.13.13.2 | ba:ad:be:ef:10:13 | Established | 5 | Topo |
| 14.14.14.2 | ba:ad:be:ef:10:14 | Established | 5 | Topo |

clear ancp subscriber

| | |
|--------------------------|---|
| Syntax | clear ancp subscriber <identifier <i>identifier-string</i> <ip-address <i>ip-address</i> system-name <i>mac-address</i>>> |
| Release Information | Command introduced in Junos OS Release 11.4. |
| Description | Clear the connection with all ANCP subscribers or with the specified ANCP subscriber. |
| Options | <p>none—Clear all ANCP subscribers.</p> <p>identifier <i>identifier-string</i>—(Optional) Clear the ANCP subscriber identified by the access loop ID on all neighbors.</p> <p>identifier <i>identifier-string</i> ip-address <i>ip-address</i>—(Optional) Clear the ANCP subscriber specified by the access loop ID and the neighbor IP address.</p> <p>identifier <i>identifier-string</i> system-name <i>mac-address</i>—(Optional) Clear the ANCP subscriber specified by the access loop ID and the neighbor MAC address.</p> |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• show ancp subscriber on page 19 |
| List of Sample Output | show ancp subscriber brief on page 7 show ancp subscriber brief on page 7 clear ancp subscriber on page 7 |
| Output Fields | When you enter this command, you are provided no feedback on the status of your request. You can enter the show ancp subscriber command before and after clearing the ANCP neighbors to verify the clear operation. |

Sample Output

**show ancp subscriber
brief**

```
user@host> show ancp subscriber brief
Loop Identifier      Type      Interface      Rate      Neighbor
                    Kbps
port-1-10            VDSL2     set-ge-10410   64        10.10.10.2
port-1-11            VDSL2     set-ge-10411   64        11.11.11.2
port-2-10            VDSL2     ge-1/0/4.12    64        10.12.12.2
port-2-10            VDSL2     ge-1/0/4.12    64        10.12.12.3
port-2-11            VDSL2     ge-1/0/4.13    64        10.13.13.2
```

```
user@host> clear ancp subscriber identifier port-2-10
```

```
user@host> show ancp subscriber brief
Loop Identifier      Type      Interface      Rate      Neighbor
                    Kbps
port-1-10            VDSL2     set-ge-10410   64        10.10.10.2
port-1-11            VDSL2     set-ge-10411   64        11.11.11.2
port-2-11            VDSL2     ge-1/0/4.13    64        10.13.13.2
```

**show ancp subscriber
brief**

```
user@host> show ancp subscriber brief
Loop Identifier      Type      Interface      Rate      Neighbor
                    Kbps
port-1-10            VDSL2     set-ge-10410   64        10.10.10.2
port-1-11            VDSL2     set-ge-10411   64        11.11.11.2
port-2-10            VDSL2     ge-1/0/4.12    64        10.12.12.2
port-2-10            VDSL2     ge-1/0/4.12    64        10.12.12.3
port-2-11            VDSL2     ge-1/0/4.13    64        10.13.13.2
```

```
user@host> clear ancp subscriber identifier port-2-10 ip-address 10.12.12.3
```

```
user@host> show ancp subscriber brief
Loop Identifier      Type      Interface      Rate      Neighbor
                    Kbps
port-1-10            VDSL2     set-ge-10410   64        10.10.10.2
port-1-11            VDSL2     set-ge-10411   64        11.11.11.2
port-2-10            VDSL2     ge-1/0/4.12    64        10.12.12.2
port-2-11            VDSL2     ge-1/0/4.13    64        10.13.13.2
```

clear ancp subscriber

```
user@host> clear ancp subscriber
```

request ancp oam interface

| | |
|---------------------------------|--|
| Syntax | request ancp oam interface (<i>interface-name</i> interface-set <i>set-name</i>) <count <i>count</i> > <timeout <i>duration</i> > |
| Release Information | Command introduced in Junos OS Release 11.4. |
| Description | Trigger the access node to run a loopback test on the local loop between the access node and the customer premises equipment. You must specify either an ANCP interface or an ANCP interface set. The access node responds to the NAS with the results of the test. |
| Options | <p><i>interface-name</i>—Name of the ANCP interface on whose local loop the loopback test is run.</p> <p>interface-set <i>set-name</i>—Name of the ANCP interface set on whose local loop the loopback test is run.</p> <p>count <i>count</i>—(Optional) Number of times a loopback message is sent on the local loop. Range: 1 through 32. Default: 1.</p> <p>timeout <i>duration</i>—(Optional) Period of time in seconds that the NAS waits for a response to the OAM request. Range: 0 through 255. Default: 5 .</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">Triggering ANCP OAM |
| List of Sample Output | request ancp oam interface on page 8 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request, including the result of the test, the response code, and the response string returned with the OAM response in the event of failure, an error code is displayed. |

Sample Output

| | |
|-----------------------------------|---|
| request ancp oam interface | user@host> request ancp oam interface ge-1/0/4.12 count 5 timeout 40 request succeeded 0x503 : DSL line status showtime DEFAULT RESPONSE |
|-----------------------------------|---|

request ancp oam neighbor

| | |
|---------------------------------|--|
| Syntax | request ancp oam neighbor (ip-address <i>ip-address</i> system-name <i>neighbor-name</i>) subscriber <i>identifier-string</i> <count <i>count</i> > <timeout <i>duration</i> > |
| Release Information | Command introduced in Junos OS Release 11.4. |
| Description | Trigger the access node to run a loopback test on the local loop between the access node and the customer premises equipment. You must specify both the access node and the subscriber. The access node responds to the NAS with the results of the test. |
| Options | <p>ip-address <i>ip-address</i>—IP address that specifies the access node on whose local loop the loopback test is run.</p> <p>system-name <i>neighbor-name</i>—System name that specifies the access node on whose local loop the loopback test is run.</p> <p>subscriber <i>identifier-string</i>—Access identifier that specifies the subscriber on whose local loop the loopback test is run.</p> <p>count <i>count</i>—(Optional) Number of times a loopback message is sent on the local loop. Range: 1 through 32. Default: 1.</p> <p>timeout <i>duration</i>—(Optional) Period of time in seconds that the NAS waits for a response to the OAM request. Range: 0 through 255. Default: 5.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> Triggering ANCP OAM |
| List of Sample Output | request ancp oam subscriber on page 9 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request, including the result of the test, the response code, and the response string returned with the OAM response in the event of failure, an error code is displayed. |

Sample Output

```
request ancp oam subscriber
user@host> request ancp oam neighbor 10.10.10.1 subscriber "dslam port-1-11"
request succeeded
0x503 : DSL line status showtime
DEFAULT RESPONSE
```

show ancp cos

| | |
|---------------------------------|---|
| Syntax | show ancp cos <identifier <i>identifier</i>> <last-update> <pending-update> |
| Release Information | Command introduced in Junos OS Release 9.4. |
| Description | Display information about the CoS state for subscriber traffic. |
| Options | <p>identifier <i>identifier</i>—(Optional) Display information about the local loops for the specified access identifier.</p> <p>last-update—(Optional) Display the most recently updated CoS information.</p> <p>pending-update—(Optional) Display the pending update of CoS information.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • show ancp neighbor on page 13 • show ancp subscriber on page 19 |
| List of Sample Output | show ancp cos on page 12 show ancp cos last-update on page 12 show ancp cos pending-update on page 12 |
| Output Fields | Table 4 on page 10 lists the output fields for the show ancp cos command. Output fields are listed in the approximate order in which they appear. |

Table 4: show ancp cos Output Fields

| Field Name | Field Description |
|------------------------|---|
| QoS Adjust Flag | State of QoS adjust: TRUE (configured) or FALSE (not configured). |
| Keepalive Timer | Interval between the keepalive messages that ANCP sends to CoS. |
| Cos State | State of the ANCP-CoS interaction: <ul style="list-style-type: none"> • ANCPD_COS_CONNECT_NEEDED • ANCPD_COS_CONNECT_PENDING • ANCPD_COS_CONNECT_DONE • ANCPD_COS_SESSION_SENT • ANCPD_COS_WRITE_READY |
| Connect Time | Time at which ANCP connected to CoS; useful for debugging. |
| Session Time | Time at which ANCP sent a session connect message to CoS; useful for debugging. |

Table 4: show ancp cos Output Fields (*continued*)

| Field Name | Field Description |
|------------------------------|--|
| Routing Instance Time | Time at which ANCP sent the routing instance to CoS; useful for debugging. |
| Keepalive Time | Time at which the last keepalive message was sent. |
| Rate Update Time | Time at which the shaping rate was last updated. |
| Type | Subscriber access type: ifl indicates that a single VLAN carries subscriber traffic and iflset indicates that a set of VLANs carries subscriber traffic. |
| Name | System-wide name of the particular subscriber access. |
| Index | Access identifier. |
| Pending Update | Actual downstream data rate to be applied next to this local loop, in Kbps. |
| Last Update | Actual downstream data rate last applied to this local loop, in Kbps. |

Sample Output

show ancp cos

```
user@host> show ancp cos
```

```
Qos Adjust Flag:      TRUE
Keepalive Timer:      45 secs
Cos State:            WRITE_READY
Connect Time:         Mon Nov 17 15:03:01 2008
Session Time:         Mon Nov 17 15:03:13 2008
Routing Instance Time: Mon Nov 17 15:03:14 2008
Keepalive Time:       Not Set
Rate Update Time:     Mon Nov 17 15:03:15 2008
```

| Type | Name | Index | Pending Update | Last Update |
|--------|--------------|-------|----------------|-------------|
| iflset | set-ge-10410 | 1 | None | 64 Kbps |
| iflset | set-ge-10411 | 2 | None | 64 Kbps |
| ifl | ge-1/0/4.2 | 71 | None | 64 Kbps |
| ifl | ge-1/0/4.3 | 72 | None | 64 Kbps |

show ancp cos last-update

```
user@host> show ancp cos last-update
```

```
Qos Adjust Flag:      TRUE
Keepalive Timer:      45 secs
Cos State:            WRITE_READY
Connect Time:         Mon Nov 17 15:03:01 2008
Session Time:         Mon Nov 17 15:03:13 2008
Routing Instance Time: Mon Nov 17 15:03:14 2008
Keepalive Time:       Wed Nov 19 15:32:14 2008
Rate Update Time:     Mon Nov 17 15:03:15 2008
```

| Type | Name | Index | Pending Update | Last Update |
|--------|---------|-------|----------------|-------------|
| iflset | iflset0 | 1 | None | 64 Kbps |
| iflset | iflset1 | 2 | None | 64 Kbps |

show ancp cos pending-update

```
user@host> show ancp cos pending-update
```

```
Qos Adjust Flag:      TRUE
Keepalive Timer:      45 secs
Cos State:            WRITE_READY
Connect Time:         Mon Nov 17 15:03:01 2008
Session Time:         Mon Nov 17 15:03:13 2008
Routing Instance Time: Mon Nov 17 15:03:14 2008
Keepalive Time:       Wed Nov 19 15:32:29 2008
Rate Update Time:     Mon Nov 17 15:03:15 2008
```

show ancp neighbor

| | |
|---------------------------------|---|
| Syntax | show ancp neighbor <brief detail extensive terse> <ip-address <i>ip-address</i> > <system-name <i>mac-address</i> > |
| Release Information | Command introduced in Junos OS Release 9.4. |
| Description | Display information about all ANCP neighbors or the specified ANCP neighbor. |
| Options | brief detail extensive terse —(Optional) Display the specified level of detail. ip-address <i>ip-address</i> —(Optional) IP address of the ANCP neighbor (access node). system-name <i>mac-address</i> —(Optional) MAC address of the ANCP neighbor (access node). |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • show ancp cos on page 10 • show ancp subscriber on page 19 |
| List of Sample Output | show ancp neighbor on page 16 show ancp neighbor detail on page 16 show ancp neighbor ip-address on page 17 show ancp neighbor system-name on page 17 |
| Output Fields | Table 5 on page 13 lists the output fields for the show ancp neighbor command. Output fields are listed in the approximate order in which they appear. |

Table 5: show ancp neighbor Output Fields

| Field Name | Field Description |
|------------------|---|
| IP Address | IP address of the ANCP neighbor. |
| System Name | MAC address of the ANCP neighbor. |
| State | State of the ANCP adjacency: <ul style="list-style-type: none"> • Established—ANCP session has been established. • Init—ANCP session has been initiated. • SynSent—ANCP has sent a SYN message. • SynReceived—ANCP has sent a SYNACK message. |
| Subscriber Count | Number of subscribers associated with the ANCP neighbor (access local loop). |

Table 5: show ancp neighbor Output Fields (*continued*)

| Field Name | Field Description |
|-----------------------------|---|
| Capabilities | Negotiated ANCP capability: <ul style="list-style-type: none"> • Topo—Topology discovery. • OAM—Performance of local Operations Administration Maintenance (OAM) procedures on an access loop controlled by the router. |
| TCP Port | TCP port on which ANCP messages are exchanged. |
| System Instance | Number identifying the ANCP link instance from the edge device's perspective. |
| Peer Instance | Number identifying the ANCP instance from the access node's perspective. This number is unique and changes when the node or link comes back up after going down. |
| Timer | Adjacency timer value advertised by the ANCP peer in 100 ms increments; the interval between ANCP ACK messages. This value remains constant for the duration of an ANCP session. |
| Partition Type | Number that identifies whether partitions are used and how the ID is negotiated: <ul style="list-style-type: none"> • 0—No partition. • 1—Fixed partition requested. • 2—Fixed partition assigned. |
| Partition Flag | Number that specifies the type of partition requested: 1 (new adjacency) or 2 (recovered adjacency). |
| Partition Identifier | Number that associates the ANCP message with a specific partition. |
| Dead Timer | Remaining period that the edge device waits for adjacency packets from a neighbor before declaring the neighbor to be down. The maximum dead time value is three times the configured adjacency timer value. This field displays the current value based on the time that the last adjacency packet was received. |

Table 5: show ancpl neighbor Output Fields (*continued*)

| Field Name | Field Description |
|---|--|
| Received | <p>Count of the following ANCP message packets received by the node from the neighbor:</p> <ul style="list-style-type: none"> • Syn Count—Synchronization message used to maintain an adjacency. • Synack Count—Neighbor response to the node's synchronization messages. • Rstack Count—Message indicating that the link to the neighbor needs to be reset. • Ack Count—Acknowledgment message periodically received after an adjacency has been established. • Port Up Count—Status message indicating that a port has transitioned to the up state. • Port Down Count—Status message indicating that a port has transitioned to the down state. • OAM Response Count—Number of OAM responses received in reply to request commands. • Other Count—Count of all other ANCP received message packets that do not fit into one of the other categories. |
| Sent | <p>Count of the following ANCP message packets sent by the node:</p> <ul style="list-style-type: none"> • Syn Count—Synchronization message used to maintain an adjacency. • Synack Count—Node response to the neighbor's synchronization messages. • Rstack Count—Message indicating that the link to the node needs to be reset. • Ack Count—Acknowledgment message periodically sent after an adjacency has been established. • OAM Request Count—Number of OAM request commands sent. |
| Max Discovery Limit Exceed Count | Number of times that the maximum number of discovery table entries accepted from the neighbor has been exceeded. |

Sample Output

show ancp neighbor

user@host> show ancp neighbor

| IP Address | State | Up Time | Subscriber Count | Capabilities |
|------------|-------------|---------|------------------|--------------|
| 10.10.10.2 | Established | 3 | 2 | Topo, OAM |
| 11.11.11.2 | Established | 3 | 2 | Topo, OAM |

show ancp neighbor detail

user@host> show ancp neighbor detail

Neighbor Information

```

IP Address           : 192.168.10.1
System Name          : 00:00:64:1b:01:02
  Up Time              : 38
  TCP Port              : 64959
  State                 : Established
  Subscriber Count      : 7
  Capabilities          : Topology Discovery
  System Instance       : 11
  Peer Instance         : 1
  Adjacency Timer (in 100ms) : 50
  Peer Adjacency Timer (in 100ms) : 100
  Partition Type        : 0
  Partition Flag        : 1
  Partition Identifier   : 0
  Dead Timer            : 22
  Received Syn Count     : 47
  Received Synack Count  : 48
  Received Rstack Count  : 2
  Received Ack Count     : 12
  Received Port Up Count : 8
  Received Port Down Count : 2
  Received Other Count   : 0
  Sent Syn Count         : 48
  Sent Synack Count      : 47
  Sent Rstack Count      : 1
  Sent Ack Count         : 12
  Max Discovery Limit Exceed Count : 0
IP Address           : 192.168.9.1
System Name          : 00:00:64:1c:01:02
  Up Time              : 36
  TCP Port              : 61408
  State                 : Established
  Subscriber Count      : 1
  Capabilities          : Topology Discovery
  System Instance       : 12
  Peer Instance         : 1
  Adjacency Timer (in 100ms) : 50
  Peer Adjacency Timer (in 100ms) : 100
  Partition Type        : 0
  Partition Flag        : 1
  Partition Identifier   : 0
  Dead Timer            : 23
  Received Syn Count     : 24
  Received Synack Count  : 20
  Received Rstack Count  : 2
  Received Ack Count     : 9

```

```

Received Port Up Count      : 5
Received Port Down Count    : 0
Received OAM Responses Count : 2
Received Other Count        : 0
Sent Syn Count              : 20
Sent Synack Count           : 24
Sent Rstack Count           : 1
Sent Ack Count              : 9
Sent OAM Requests Count     : 4
Max Discovery Limit Exceed Count : 0

```

**show ancp neighbor
ip-address**

user@host> show ancp neighbor 10.10.10.2

```

Neighbor Information
  IP Address      : 10.10.10.2
  System Name     : ba:ad:be:ef:10:10
  TCP Port        : 3332
  State           : Established
  Subscriber Count : 5
  Capabilities    : Topology Discovery
  System Instance : 6
  Peer Instance   : 1695
  Timer           : 250
  Partition Type  : 0
  Partition Flag   : 1
  Partition Identifier : 0
  Dead Timer      : 63

```

**show ancp neighbor
system-name**

user@host> show ancp neighbor ba:ad:be:ef:10:10 detail

```

Neighbor Information
  IP Address      : 10.100.0.1
  System Name     : 00:00:64:1b:01:02
  Up Time         : 19
  TCP Port        : 1028
  State           : Established
  Subscriber Count : 2
  Capabilities    : Topology Discovery, OAM
  System Instance : 1
  Peer Instance   : 10
  Adjacency Timer (in 100ms) : 100
  Peer Adjacency Timer (in 100ms) : 250
  Partition Type  : 0
  Partition Flag   : 1
  Partition Identifier : 0
  Dead Timer      : 55
  Received Syn Count : 1

  Received Synack Count : 1
  Received Rstack Count : 0
  Received Ack Count    : 1
  Received Port Up Count : 34
  Received Port Down Count : 0
  Received OAM Responses Count : 2
  Received Other Count   : 0
  Sent Syn Count         : 1
  Sent Synack Count      : 1
  Sent Rstack Count      : 0
  Sent Ack Count         : 3
  Sent OAM Requests Count : 4

```

Max Discovery Limit Exceed Count : 3

show ancp subscriber

| | |
|---------------------------------|--|
| Syntax | <code>show ancp subscriber</code> <code><brief detail></code> <code><identifier <i>identifier</i>></code> <code><neighbor <i>ip-address</i>></code> |
| Release Information | Command introduced in Junos OS Release 9.4. |
| Description | Display information about all subscribers (local access loops), the identified subscriber, or the subscriber associated with the specified ANCP neighbor (access node). |
| Options | <p>brief detail—(Optional) Display the specified level of detail.</p> <p>identifier <i>identifier</i>—(Optional) Display information about the subscriber specified by the access identifier.</p> <p>neighbor <i>ip-address</i> —(Optional) Display information about the local loops connected to the access node specified by the IP address.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • clear ancp subscriber on page 6 • show ancp cos on page 10 • show ancp neighbor on page 13 |
| List of Sample Output | show ancp subscriber brief on page 22 show ancp subscriber detail on page 22 show ancp subscriber identifier identifier-string detail on page 23 |
| Output Fields | Table 6 on page 19 lists the output fields for the show ancp subscriber command. Output fields are listed in the approximate order in which they appear. |

Table 6: show ancp subscriber Output Fields

| Field Name | Field Description |
|------------------------|--|
| Loop Identifier | Access loop identifier as sent by the access node and configured to map the subscriber to an interface. |
| Type | Type of digital subscriber line employed by the access node: ADSL1 , ADSL2 , ADSL2+ , VDSL1 , VDSL2 , SDSL , or UNKNOWN . |
| State | State of the DSL line: Idle , Show Time , or Silent . |
| Rate Kbps | Actual downstream data rate for this local loop. |
| Neighbor | IP address of ANCP neighbor (access node). |

Table 6: show ancp subscriber Output Fields (*continued*)

| Field Name | Field Description |
|--|--|
| Access Node Identifier | Access node identifier as sent by the access node and configured to map the subscriber to an interface. |
| Neighbor IP Address | IP address of the ANCP neighbor (access node). |
| Aggregate Circuit Identifier Binary | Binary identifier for the VLAN circuit ID. |
| DSL Type | Type of digital subscriber line employed by the access node: ADSL1 , ADSL2 , ADSL2+ , VDSL1 , VDSL2 , SDSL , or UNKNOWN . |
| Interface Type | Type of interface employed for subscriber traffic: ifl for a single VLAN or interface-set for a configured group of VLANs. |
| Interface | Name of the interface set or logical interface. |
| DSL Line State | State of the DSL line: Idle , Show Time , or Silent . |
| Actual Net Data Upstream | Actual upstream data rate for this local loop, in Kbps. |
| Actual Net Data Downstream | Actual downstream data rate for this local loop, in Kbps. |
| DSL Line Data Link | Data link protocol employed on the access loop: AAL5 or Ethernet . |
| DSL Line Encapsulation | Encapsulation type on the access loop, for Ethernet only: <ul style="list-style-type: none"> • 0—NA, type not conveyed • 1—Untagged Ethernet • 2—Single-tagged Ethernet |
| DSL Line Encapsulation Payload | Payload carried across the access loop: <ul style="list-style-type: none"> • 0—NA, type not conveyed • 1—PPPoA LLC • 2—PPPoA null • 3—IPoA LLC • 4—IPoA null • 5—Ethernet over AAL5 LLC with FCS • 6—Ethernet over AAL5 LLC without FCS • 7—Ethernet over AAL5 null with FCS • 8—Ethernet over AAL5 null without FCS |
| Minimum Net Data Upstream | Minimum upstream data rate desired by the operator for this local loop, in Kbps. |

Table 6: show ancpc subscriber Output Fields (*continued*)

| Field Name | Field Description |
|--|---|
| Minimum Net Data Downstream | Minimum downstream data rate desired by the operator for this local loop, in Kbps. |
| Maximum Net Data Upstream | Maximum upstream data rate desired by the operator for this local loop, in Kbps. |
| Maximum Net Data Downstream | Maximum downstream data rate desired by the operator for this local loop, in Kbps. |
| Attainable Net Data Upstream | Maximum attainable upstream data rate for this local loop, in Kbps. |
| Attainable Net Data Downstream | Maximum attainable downstream data rate for this local loop, in Kbps. |
| Minimum Low Power Data Downstream | Minimum downstream data rate desired by the operator for this local loop in low power state, in Kbps. |
| Minimum Low Power Data Upstream | Minimum upstream data rate desired by the operator for this local loop in low power state, in Kbps. |
| Maximum Interleave Delay Downstream | Maximum interleaving delay for downstream data, in milliseconds. |
| Maximum Interleave Delay Upstream | Maximum interleaving delay for upstream data, in milliseconds. |
| Actual Interleave Delay Downstream | Actual interleaving delay for downstream data, in milliseconds. |
| Actual Interleave Delay Upstream | Actual interleaving delay for upstream data, in milliseconds. |

Sample Output

**show ancp subscriber
brief**

user@host> show ancp subscriber brief

| Loop Identifier | Type | Interface | Rate Kbps | Neighbor |
|-----------------|-------|--------------|--------------|------------|
| port-1-10 | VDSL2 | set-ge-10410 | 64 | 10.10.10.2 |
| port-1-11 | VDSL2 | set-ge-10411 | 64 | 11.11.11.2 |
| port-2-10 | VDSL2 | ge-1/0/4.12 | 64 | 10.12.12.2 |
| port-2-11 | VDSL2 | ge-1/0/4.13 | 64 | 10.13.13.2 |

**show ancp subscriber
detail**

user@host> show ancp subscriber detail

```
Subscriber Information
  Access Loop Identifier : port-2-11
    Neighbor IP Address      : 10.11.11.2
    Aggregate Circuit Identifier Binary : 0/0
    DSL Type                 : VDSL2
    Interface Type           : ifl
    Interface                : ge-1/0/4.10
    DSL Line State           : Show Time
    Actual Net Data Upstream : 64
    Actual Net Data Downstream : 64
    DSL Line Data Link        : AAL5
    DSL Line Encapsulation    : N/A
    DSL Line Encapsulation Payload : N/A
    Minimum Net Data Upstream : 64
    Minimum Net Data Downstream : 64
    Maximum Net Data Upstream : 64
    Maximum Net Data Downstream : 64
    Attainable Net Data Upstream : 64
    Attainable Net Data Downstream : 64
    Minimum Low Power Data Downstream : 64
    Minimum Low Power Data Upstream : 64
    Maximum Interleave Delay Downstream : 50
    Maximum Interleave Delay Upstream : 50
    Actual Interleave Delay Downstream : 50
    Actual Interleave Delay Upstream : 50
  Access Loop Identifier : port-1-11
    Neighbor IP Address      : 10.11.11.2
    Aggregate Circuit Identifier Binary : 0/0
    DSL Type                 : DSL 0
    Interface Type           : interface-set
    Interface                : set-ge-10411
    DSL Line State           : Show Time
    Actual Net Data Upstream : 64
    Actual Net Data Downstream : 64
    DSL Line Data Link        : AAL5
    DSL Line Encapsulation    : N/A
    DSL Line Encapsulation Payload : N/A
    Minimum Net Data Upstream : 64
    Minimum Net Data Downstream : 64
    Maximum Net Data Upstream : 64
    Maximum Net Data Downstream : 64
    Attainable Net Data Upstream : 64
    Attainable Net Data Downstream : 64
    Minimum Low Power Data Downstream : 64
```



```

Minimum Low Power Data Upstream      : 64
Maximum Interleave Delay Downstream  : 50
Maximum Interleave Delay Upstream    : 50
Actual Interleave Delay Downstream    : 50
Actual Interleave Delay Upstream      : 50

```

show ancp subscriber
identifier
identifier-string detail

user@host> show ancp subscriber identifier port-1-11 detail

```

Access Loop Identifier : port-1-11
Neighbor IP Address    : 10.11.11.2
Aggregate Circuit Identifier Binary : 0/0
DSL Type               : DSL 0
Interface Type         : interface-set
Interface              : set-ge-10411
DSL Line State         : Show Time
Actual Net Data Upstream : 64
Actual Net Data Downstream : 64
DSL Line Data Link     : AAL5
DSL Line Encapsulation : N/A
DSL Line Encapsulation Payload : N/A
Minimum Net Data Upstream : 64
Minimum Net Data Downstream : 64
Maximum Net Data Upstream : 64
Maximum Net Data Downstream : 64
Attainable Net Data Upstream : 64
Attainable Net Data Downstream : 64
Minimum Low Power Data Downstream : 64
Minimum Low Power Data Upstream : 64
Maximum Interleave Delay Downstream : 50
Maximum Interleave Delay Upstream : 50
Actual Interleave Delay Downstream : 50
Actual Interleave Delay Upstream : 50

```


CHAPTER 2

BFD Operational Mode Commands

Table 7 on page 25 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot Bidirectional Forwarding Detection (BFD) sessions. Commands are listed in alphabetical order.

Table 7: BFD Operational Mode Commands

| Task | Command |
|---------------------------------|-----------------------------------|
| Clear BFD parameters. | <code>clear bfd adaptation</code> |
| Clear BFD sessions. | <code>clear bfd session</code> |
| Display BFD session statistics. | <code>show bfd session</code> |



NOTE: The protocol client for which the BFD session is active can be either IS-IS or OSPF.



NOTE: For information about how to configure BFD, see the *Junos Routing Protocols Configuration Guide*.

clear bfd adaptation

| | |
|---------------------------------|--|
| Syntax | <code>clear bfd adaptation</code> <code><address <i>session-address</i>></code> <code><discriminator <i>discr-number</i>></code> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | <p>Clear adaptation for Bidirectional Forwarding Detection (BFD) sessions. BFD is a simple hello mechanism that detects failures in a network. Configured BFD interval timers can change, adapting to network situations. Use this command to return BFD interval timers to their configured values.</p> <p>The clear bfd adaptation command is hitless, meaning that the command does not affect traffic flow on the routing device.</p> |
| Options | <p>none—Clear adaptation for all BFD sessions.</p> <p>address <i>session-address</i>—(Optional) Clear adaptation for all BFD sessions matching the specified address.</p> <p>discriminator <i>discr-number</i>—(Optional) Clear adaptation for the local BFD session matching the specified discriminator.</p> |
| Additional Information | For more information, see the description of the bfd-liveness-detection configuration statement in the <i>Junos Routing Protocols Configuration Guide</i> . |
| Required Privilege Level | clear |
| List of Sample Output | clear bfd adaptation on page 26 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear bfd adaptation user@host> clear bfd adaptation

clear bfd session

| | |
|---|---|
| Syntax | clear bfd session <address <i>session-address</i> > <discriminator <i>discr-number</i> > <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switch and QFX Series) | clear bfd session <address <i>session-address</i> > <discriminator <i>discr-number</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 12.1 for the QFX Series. |
| Description | Drop one or more Bidirectional Forwarding Detection (BFD) sessions. |
| Options | <p>none—Drop all BFD sessions.</p> <p>address <i>session-address</i>—(Optional) Drop all BFD sessions matching the specified address.</p> <p>discriminator <i>discr-number</i>—(Optional) Drop the local BFD session matching the specified discriminator.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none"> • show bfd session on page 28 |
| List of Sample Output | clear bfd session on page 27 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

```
clear bfd session      user@host> clear bfd session
```

show bfd session

| | |
|---|---|
| Syntax | <code>show bfd session</code> <code><brief detail extensive summary></code> <code><address <i>address</i>></code> <code><discriminator <i>discriminator</i>></code> <code><logical-system (all <i>logical-system-name</i>)></code> <code><prefix <i>address</i>></code> |
| Syntax (EX Series Switch and QFX Series) | <code>show bfd session</code> <code><brief detail extensive summary></code> <code><address <i>address</i>></code> <code><discriminator <i>discriminator</i>></code> <code><prefix <i>address</i>></code> |
| Release Information | Command introduced before Junos OS Release 7.4. Options discriminator and address introduced in Junos OS Release 8.2. Option prefix introduced in Junos OS Release 9.0. Command introduced in Junos OS Release 12.1 for the QFX Series. |
| Description | Display information about active Bidirectional Forwarding Detection (BFD) sessions. |
| Options | none —(Same as brief) Display information about active BFD sessions. brief detail extensive summary —(Optional) Display the specified level of output. address <i>address</i> —(Optional) Display information about the BFD session for the specified neighbor address. discriminator <i>discriminator</i> —(Optional) Display information about the BFD session using the specified local discriminator. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. prefix <i>address</i> —(Optional) Display information about all of the BFD sessions for the specified LDP forwarding equivalence class (FEC). |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">• clear bfd session on page 27• Examples: Configuring BFD for Static Routes• Example: Configuring BFD for OSPF• Example: Configuring BFD for BGP• Configuring PIM and the Bidirectional Forwarding Detection (BFD) Protocol• Example: Configuring BFD for IS-IS |

List of Sample Output [show bfd session on page 33](#)
[show bfd session brief on page 33](#)
[show bfd session detail on page 33](#)
[show bfd session detail \(with Authentication\) on page 33](#)
[show bfd session address extensive on page 33](#)
[show bfd session extensive on page 34](#)
[show bfd session extensive \(with Authentication\) on page 34](#)
[show bfd session summary on page 35](#)

Output Fields [Table 8 on page 29](#) describes the output fields for the **show bfd session** command. Output fields are listed in the approximate order in which they appear.

Table 8: show bfd session Output Fields

| Field Name | Field Description | Level of Output |
|--------------------------|---|---------------------------------------|
| Address | Address on which the BFD session is active. | brief detail extensive none |
| State | State of the BFD session: Up , Down , Init (initializing), or Failing . | brief detail extensive none |
| Interface | Interface on which the BFD session is active. | brief detail extensive none |
| Detect Time | Negotiated time interval, in seconds, used to detect BFD control packets. | brief detail extensive none |
| Transmit Interval | Time interval, in seconds, used by the transmitting system to send BFD control packets. | brief detail extensive none |
| Multiplier | Negotiated multiplier by which the time interval is multiplied to determine the detection time for the transmitting system. | detail extensive |
| Session up time | How long a BFD session has been established. | detail extensive |
| Client | Protocol for which the BFD session is active: ISIS , OSPF , or Static . | detail extensive |
| TX interval | Time interval, in seconds, used by the host system to transmit BFD control packets. | brief detail extensive none |
| RX interval | Time interval, in seconds, used by the host system to receive BFD control packets. | brief detail extensive none |
| Authenticate | Indicates that BFD authentication is configured. | detail extensive |
| keychain | Name of the security authentication keychain being used by a specific client. BFD authentication information for a client is provided in a single line and includes the keychain , algo , and mode parameters. Multiple clients can be configured on a BFD session. | extensive |

Table 8: show bfd session Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-----------------------------------|--|-------------------------|
| algo | <p>BFD authentication algorithm being used for a specific client: keyed-md5, keyed-sha-1, meticulous-keyed-md5, meticulous-keyed-sha-1, or simple-password.</p> <p>BFD authentication information for a client is provided in a single line and includes the keychain, algo, and mode parameters. Multiple clients can be configured on a BFD session.</p> | extensive |
| mode | <p>Level of BFD authentication enforcement being used by a specific client: strict or loose. Strict enforcement indicates that authentication is configured at both ends of the session (the default). Loose enforcement indicates that one end of the session might not be authenticated.</p> <p>BFD authentication information for a client is provided in a single line and includes the keychain, algo, and mode parameters. Multiple clients can be configured on a BFD session.</p> | extensive |
| Local diagnostic | Local diagnostic information about failing BFD sessions. | detail extensive |
| Remote diagnostic | Remote diagnostic information about failing BFD sessions. | detail extensive |
| Remote state | Reports whether the remote system's BFD packets have been received and whether the remote system is receiving transmitted control packets. | detail extensive |
| Version | BFD version: 0 or 1 . | extensive |
| Replicated | The replicated flag appears when nonstop routing or graceful Routing Engine switchover is configured and the BFD session has been replicated to the backup Routing Engine. | detail extensive |
| Min async interval | Minimum amount of time, in seconds, between asynchronous control packet transmissions across the BFD session. | extensive |
| Min slow interval | Minimum amount of time, in seconds, between synchronous control packet transmissions across the BFD session. | extensive |
| Adaptive async TX interval | Transmission interval being used because of adaptation. | extensive |
| RX interval | Minimum required receive interval. | extensive |
| Local min TX interval | Minimum amount of time, in seconds, between control packet transmissions on the local system. | extensive |
| Local min RX interval | Minimum amount of time, in seconds, between control packet detections on the local system. | extensive |
| Remote min TX interval | Minimum amount of time, in seconds, between control packet transmissions on the remote system. | extensive |
| Remote min RX interval | Minimum amount of time, in seconds, between control packet detections on the remote system. | extensive |

Table 8: show bfd session Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-------------------------------------|---|------------------|
| Threshold transmission interval | Threshold for notification if the transmission interval increases. | extensive |
| Threshold for detection time | Threshold for notification if the detection time increases. | extensive |
| Local discriminator | Authentication code used by the local system to identify that BFD session. | extensive |
| Remote discriminator | Authentication code used by the remote system to identify that BFD session. | extensive |
| Echo mode | Information about the state of echo transmissions on the BFD session. | extensive |
| Prefix | LDP FEC address associated with the BFD session. | All levels |
| Egress, Destination | Displays the LDP FEC destination address. This field is displayed only on a router at the egress of an LDP FEC, where the BFD session has an LDP Operation, Administration, and Maintenance (OAM) client. | All levels |
| Remote is control-plane independent | <p>The BFD session on the remote peer is running on its Packet Forwarding Engine. In this case, when the remote node undergoes a graceful restart, the local peer can help the remote peer with the graceful restart.</p> <p>The following BFD sessions are not distributed to the Packet Forwarding Engine: multihop sessions, tunnel-encapsulated sessions, and sessions over aggregated Ethernet and integrated routing and bridging (IRB) interfaces.</p> | extensive |
| Authentication | <p>Summary status of BFD authentication:</p> <ul style="list-style-type: none"> • status—enabled/active indicates authentication is configured and active. enabled/inactive indicates authentication is configured but not active. This only occurs when the remote end of the session does not support authentication and loose checking is configured. • keychain—Name of the security authentication keychain associated with the specified BFD session. • algo—BFD authentication algorithm being used: keyed-md5, keyed-sha-1, meticulous-keyed-md5, meticulous-keyed-sha-1, or simple-password. • mode—Level of BFD authentication enforcement: strict or loose. Strict enforcement indicates authentication is configured at both ends of the session (the default). Loose enforcement indicates that one end of the session might not be authenticated. <p>This information is only shown if BFD authentication is configured.</p> | extensive |
| Session ID | The BFD session ID number that represents the protection using MPLS fast reroute (FRR) and loop-free alternate (LFA). | detail extensive |
| sessions | Total number of active BFD sessions. | All levels |
| clients | Total number of clients that are hosting active BFD sessions. | All levels |

Table 8: show bfd session Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------------------|--|------------------|
| Cumulative transmit rate | Total number of BFD control packets transmitted per second on all active sessions. | All levels |
| Cumulative receive rate | Total number of BFD control packets received per second on all active sessions. | All levels |
| Multi-hop, min-recv-TTL | Minimum time to live (TTL) accepted if the session is configured for multihop. | extensive |
| route table | Route table used if the session is configured for multihop. | extensive |
| local address | Local address of the source used if the session is configured for multihop. | extensive |

Sample Output

show bfd session

```
user@host> show bfd session
```

| Address | State | Interface | Detect Time | Transmit Interval | Multiplier |
|-----------|-------|------------|-------------|-------------------|------------|
| 10.9.1.33 | Up | so-7/1/0.0 | 0.600 | 0.200 | 3 |
| 10.9.1.29 | Up | ge-4/0/0.0 | 0.600 | 0.200 | 3 |

2 sessions, 2 clients
Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps

show bfd session brief

The output for the **show bfd session brief** command is identical to that for the **show bfd session** command. For sample output, see [show bfd session on page 33](#).

show bfd session detail

```
user@host> show bfd session detail
```

| Address | State | Interface | Detect Time | Transmit Interval | Multiplier |
|--|-------|------------|-------------|-------------------|------------|
| 10.9.1.33 | Up | so-7/1/0.0 | 0.600 | 0.200 | 3 |
| Client OSPF, TX interval 0.200, RX interval 0.200, multiplier 3 | | | | | |
| Session up time 3d 00:34 | | | | | |
| Local diagnostic None, remote diagnostic None | | | | | |
| Remote state Up, version 1 | | | | | |
| Replicated | | | | | |
| 10.9.1.29 | Up | ge-4/0/0.0 | 0.600 | 0.200 | 3 |
| Client ISIS L2, TX interval 0.200, RX interval 0.200, multiplier 3 | | | | | |
| Session up time 3d 00:29, previous down time 00:00:01 | | | | | |
| Local diagnostic NbrSignal, remote diagnostic AdminDown | | | | | |
| Remote state Up, version 1 | | | | | |

2 sessions, 2 clients
Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps

show bfd session detail (with Authentication)

```
user@host> show bfd session detail
```

| Address | State | Interface | Detect Time | Transmit Interval | Multiplier |
|--|-------|------------|-------------|-------------------|------------|
| 10.9.1.33 | Up | so-7/1/0.0 | 0.600 | 0.200 | 3 |
| Client OSPF, TX interval 0.200, RX interval 0.200, multiplier 3, Authenticate | | | | | |
| Session up time 3d 00:34 | | | | | |
| Local diagnostic None, remote diagnostic None | | | | | |
| Remote state Up, version 1 | | | | | |
| Replicated | | | | | |
| 10.9.1.29 | Up | ge-4/0/0.0 | 0.600 | 0.200 | 3 |
| Client ISIS L2, TX interval 0.200, RX interval 0.200, multiplier 3 | | | | | |
| Session up time 3d 00:29, previous down time 00:00:01 | | | | | |
| Local diagnostic NbrSignal, remote diagnostic AdminDown | | | | | |
| Remote state Up, version 1 | | | | | |

2 sessions, 2 clients
Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps

show bfd session address extensive

```
user@host> show bfd session 10.255.245.212 extensive
```

| Address | State | Interface | Detect Time | Transmit Interval | Multiplier |
|---|-------|-----------|-------------|-------------------|------------|
| 10.255.245.212 | Up | | 1.200 | 0.400 | 3 |
| Client Static, TX interval 0.400, RX interval 0.400, multiplier 3 | | | | | |

```

Session up time 00:17:03, previous down time 00:00:14
Local diagnostic CtlExpire, remote diagnostic NbrSignal
Remote state Up, version 1
Replicated
Min async interval 0.400, min slow interval 1.000
Adaptive async tx interval 0.400, rx interval 0.400
Local min tx interval 0.400, min rx interval 0.400, multiplier 3
Remote min tx interval 0.400, min rx interval 0.400, multiplier 3
Threshold transmission interval 0.000, Threshold for detection time 0.000
Local discriminator 6, remote discriminator 16
Echo mode disabled/inactive
Multi-hop, min-recv-TTL 255, route-table 0, local-address 10.255.245.205

```

```

1 sessions, 1 clients
Cumulative transmit rate 2.5 pps, cumulative receive rate 2.5 pps

```

show bfd session extensive

```

user@host> show bfd session extensive
10.31.1.2 Up ge-2/1/8.0 0.030 0.010 3
Client OSPF realm ospf-v2 Area 0.0.0.0, TX interval 0.010, RX interval 0.010
Session up time 00:10:13
Local diagnostic None, remote diagnostic None
Remote state Up, version 1
Replicated
Min async interval 0.010, min slow interval 1.000
Adaptive async TX interval 0.010, RX interval 0.010
Local min TX interval 0.010, minimum RX interval 0.010, multiplier 3
Remote min TX interval 0.010, min RX interval 0.010, multiplier 3
Local discriminator 12, remote discriminator 4
Echo mode disabled/inactive
Remote is control-plane independent
Session ID: 0x201

```

| Address | State | Interface | Detect Time | Transmit Interval | Multiplier |
|-----------|-------|------------|-------------|-------------------|------------|
| 10.31.2.2 | Up | ge-2/1/4.0 | 0.030 | 0.010 | 3 |

```

Client OSPF realm ospf-v2 Area 0.0.0.0, TX interval 0.010, RX interval 0.010
Session up time 00:10:14
Local diagnostic None, remote diagnostic NbrSignal
Remote state Up, version 1
Replicated
Min async interval 0.010, min slow interval 1.000
Adaptive async TX interval 0.010, RX interval 0.010
Local min TX interval 0.010, minimum RX interval 0.010, multiplier 3
Remote min TX interval 0.010, min RX interval 0.010, multiplier 3
Local discriminator 13, remote discriminator 5
Echo mode disabled/inactive
Remote is control-plane independent
Session ID: 0x202

```

```

2 sessions, 2 clients
Cumulative transmit rate 200.0 pps, cumulative receive rate 200.0 pps

```

show bfd session extensive (with Authentication)

```

user@host> show bfd session extensive

```

| Address | State | Interface | Detect Time | Transmit Interval | Multiplier |
|----------------|-------|------------|-------------|-------------------|------------|
| 192.168.208.26 | Up | so-1/0/0.0 | 2.400 | 0.800 | 10 |

```

Client Static, TX interval 0.600, RX interval 0.600, Authenticate
keychain bfd, algo keyed-md5, mode loose
Session up time 00:18:07

```

```
Local diagnostic None, remote diagnostic NbrSignal
Remote state Up, version 1
Replicated
Min async interval 0.600, min slow interval 1.000
Adaptive async TX interval 0.600, RX interval 0.600
Local min TX interval 0.600, minimum RX interval 0.600, multiplier 10
Remote min TX interval 0.800, min RX interval 0.800, multiplier 3
Local discriminator 2, remote discriminator 3
Echo mode disabled/inactive
Authentication enabled/active, keychain bfd, algo keyed-md5, mode loose
```

```
1 sessions, 1 clients
Cumulative transmit rate 1.2 pps, cumulative receive rate 1.2 pps
```

**show bfd session
summary**

```
user@host> show bfd session summary
2 sessions, 2 clients
Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps
```


CHAPTER 3

BGP Operational Mode Commands

Table 9 on page 37 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot the Border Gateway Protocol (BGP). Commands are listed in alphabetical order.

Table 9: BGP Operational Mode Commands

| Task | Command |
|--|--|
| Remove damping information. | <code>clear bgp damping</code> |
| Remove entries from the neighbor database. | <code>clear bgp neighbor</code> |
| Request BGP to refresh routes. | <code>clear bgp table</code> |
| Display information about the BGP Monitoring Protocol. | <code>show bgp bmp</code> |
| Display entries in the BGP group database. | <code>show bgp group</code> |
| Display traffic statistics for BGP groups. | <code>show bgp group traffic-statistics</code> |
| Display entries in the BGP neighbor database. | <code>show bgp neighbor</code> |
| Display the BGP state replication status for nonstop active routing-enabled devices. | <code>show bgp replication</code> |
| Display BGP summary information. | <code>show bgp summary</code> |
| Display BGP damping parameters. | <code>show policy damping</code> |



NOTE: For more BGP-related commands, such as `show route protocol`, `show route instance`, and `show route table`, see Protocol-Independent Routing Operational Mode Commands.



NOTE: For information about how to configure BGP, see the *Junos Routing Protocols Configuration Guide*.

clear bgp damping

| | |
|---|---|
| Syntax | clear bgp damping <logical-system (all <i>logical-system-name</i>)> < <i>prefix</i> > |
| Syntax (EX Series Switch and QFX Series) | clear bgp damping < <i>prefix</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | Clear BGP route flap damping information. |
| Options | none —Clear all BGP route flap damping information. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. <i>prefix</i> —(Optional) Clear route flap damping information for only the specified destination prefix. |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• show policy damping on page 74• show route damping on page 516 |
| List of Sample Output | clear bgp damping on page 38 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear bgp damping user@host> clear bgp damping

clear bgp neighbor

| | |
|---|--|
| Syntax | <pre>clear bgp neighbor <as <i>as-number</i>> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <neighbor> <soft soft-inbound> <soft-minimum-igp></pre> |
| Syntax (EX Series Switch and QFX Series) | <pre>clear bgp neighbor <as <i>as-number</i>> <instance <i>instance-name</i>> <neighbor> <soft soft-inbound> <soft-minimum-igp></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> |
| Description | <p>Perform one of the following tasks:</p> <ul style="list-style-type: none"> • Change the state of one or more BGP neighbors to IDLE. For neighbors in the ESTABLISHED state, this command drops the TCP connection to the neighbors and then reestablishes the connection. • (soft or soft-inbound keyword only) Reapply export policies or import policies, respectively, and send refresh updates to one or more BGP neighbors without changing their state. |
| Options | <p>none—Change the state of all BGP neighbors to IDLE.</p> <p>as <i>as-number</i>—(Optional) Apply this command only to neighbors in the specified autonomous system (AS).</p> <p>instance <i>instance-name</i>—(Optional) Apply this command only to neighbors for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>neighbor—(Optional) IP address of a BGP peer. Apply this command only to the specified neighbor.</p> <p>soft—(Optional) Reapply any export policies and send refresh updates to neighbors without clearing the state.</p> <p>soft-inbound—(Optional) Reapply any import policies and send refresh updates to neighbors without clearing the state.</p> |

soft-minimum-igp—(Optional) Provides soft refresh of the outbound state when the interior gateway protocol (IGP) metric is reset.

Required Privilege Level clear

Related Documentation • [show bgp neighbor on page 53](#)

List of Sample Output [clear bgp neighbor on page 40](#)

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

Sample Output

clear bgp neighbor user@host> clear bgp neighbor

clear bgp table

| | |
|---|--|
| Syntax | <code>clear bgp table <i>table-name</i></code> <code><logical-system (all <i>logical-system-name</i>)></code> |
| Syntax (EX Series Switch and QFX Series) | <code>clear bgp table <i>table-name</i></code> |
| Release Information | Command introduced in Junos OS Release 9.0. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | Request that BGP refresh routes in a specified routing table. |
| Options | <code>logical-system (all <i>logical-system-name</i>)</code> —(Optional) Perform this operation on all logical systems or on a particular logical system. <code><i>table-name</i></code> —Request that BGP refresh routes in the specified table. |
| Additional Information | In some cases, a prefix limit is associated with a routing table for a VPN instance. When this limit is exceeded (for example, because of a network misconfiguration), some routes might not be inserted in the table. Such routes need to be added to the table after the network issue is resolved. Use the clear bgp table command to request that BGP refresh routes in a VPN instance table. |
| Required Privilege Level | clear |
| List of Sample Output | clear bgp table private.inet.0 on page 42 clear bgp table inet.6 logical-system all on page 42 clear bgp table private.inet.6 logical-system ls1 on page 42 clear bgp table logical-system all inet.0 on page 42 clear bgp table logical-system ls2 private.inet.0 on page 42 |
| Output Fields | This command produces no output. |

Sample Output

clear bgp table
private.inet.0

user@host> clear bgp table private.inet.0

clear bgp table inet.6
logical-system all

user@host> clear bgp table inet.6 logical-system all

clear bgp table
private.inet.6
logical-system ls1

user@host> clear bgp table private.inet.6 logical-system ls1

clear bgp table
logical-system all
inet.0

user@host> clear bgp table logical-system all inet.0

clear bgp table
logical-system ls2
private.inet.0

user@host> clear bgp table logical-system ls2 private.inet.0

show bgp bmp

| | |
|---------------------------------|--|
| Syntax | show bgp bmp |
| Release Information | Command introduced in Junos OS Release 9.5. Command introduced in Junos OS Release 9.5 for EX Series switches. |
| Description | Display information about the BGP Monitoring Protocol (BMP). |
| Options | This command has no options. |
| Required Privilege Level | view |
| List of Sample Output | show bgp bmp on page 43 |
| Output Fields | Table 10 on page 43 lists the output fields for the show bgp bmp command. Output fields are listed in the approximate order in which they appear. |

Table 10: show bgp bmp Output Fields

| Field Name | Field Description |
|-------------------------------------|--|
| BMP station address/port | IP address and port number of the monitoring station to which BGP Monitoring Protocol (BMP) statistics are sent. |
| BMP session state | Status of the BMP session: UP or DOWN . |
| Memory consumed by BMP | Memory used by the active BMP session. |
| Statistics timeout | Amount of time, in seconds, between transmissions of BMP data to the monitoring station. |
| Memory limit | Threshold, in bytes, at which the routing device stops collecting BMP data. |
| Memory-connect retry timeout | Amount of time, in seconds, after which the routing device attempts to resume a BMP session that was ended after the configured memory threshold was exceeded. |

Sample Output

```

show bgp bmp
user@host> show bgp bmp
  BMP station address/port: 172.24.24.157+5454
  BMP session state: DOWN
  Memory consumed by BMP: 0
  Statistics timeout: 15
  Memory limit: 10485760
  Memory connect retry timeout: 600

```

show bgp group

| | |
|---|--|
| Syntax | <pre>show bgp group <brief detail summary> <group-name> <exact-instance instance-name> <instance instance-name> <logical-system (all logical-system-name)> <rtf></pre> |
| Syntax (EX Series Switch and QFX Series) | <pre>show bgp group <brief detail summary> <group-name> <exact-instance instance-name> <instance instance-name></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>exact-instance option introduced in Junos OS Release 11.4.</p> |
| Description | Display information about the configured BGP groups. |
| Options | <p>none—Display group information about all BGP groups.</p> <p>brief detail summary—(Optional) Display the specified level of output.</p> <p>group-name—(Optional) Display group information for the specified group.</p> <p>exact-instance instance-name—(Optional) Display information for the specified instance only.</p> <p>instance instance-name—(Optional) Display information about BGP groups for all routing instances whose name begins with this string (for example, cust1, cust11, and cust111 are all displayed when you run the show bgp group instance cust1 command). The instance name can be master for the main instance, or any valid configured instance name or its prefix.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>rtf—(Optional) Display BGP group route targeting information.</p> |
| Required Privilege Level | view |
| List of Sample Output | <p>show bgp group on page 49</p> <p>show bgp group brief on page 49</p> <p>show bgp group detail on page 49</p> <p>show bgp group rtf detail on page 50</p> <p>show bgp group summary on page 50</p> |

Output Fields Table 11 on page 45 describes the output fields for the **show bgp group** command. Output fields are listed in the approximate order in which they appear.

Table 11: show bgp group Output Fields

| Field Name | Field Description | Level of Output |
|------------------------------------|--|-----------------------------|
| Group Type or Group | Type of BGP group: Internal or External . | All levels |
| group-index | Index number for the BGP peer group. The index number differentiates between groups when a single BGP group is split because of different configuration options at the group and peer levels. | rtf detail |
| AS | AS number of the peer. For internal BGP (IBGP), this number is the same as Local AS . | brief detail none |
| Local AS | AS number of the local routing device. | brief detail none |
| Name | Name of a specific BGP group. | brief detail none |
| Index | Unique index number of a BGP group. | brief detail none |
| Flags | Flags associated with the BGP group. This field is used by Juniper Networks customer support. | brief detail none |
| Remove-private options | Options associated with the remove-private statement. | brief detail none |
| Holdtime | Maximum number of seconds allowed to elapse between successive keepalive or update messages that BGP receives from a peer in the BGP group, after which the connection to the peer is closed and routing devices through that peer become unavailable. | brief detail none |
| Export | Export policies configured for the BGP group with the export statement. | brief detail none |
| MED tracks IGP metric update delay | Time, in seconds, that updates to multiple exit discriminator (MED) are delayed. Also displays the time remaining before the interval is set to expire | All levels |
| Traffic Statistics Interval | Time between sample periods for labeled-unicast traffic statistics, in seconds. | brief detail none |
| Total peers | Total number of peers in the group. | brief detail none |
| Established | Number of peers in the group that are in the established state. | All levels |

Table 11: show bgp group Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--|--|-----------------|
| Active/Received/Accepted/Damped | <p>Multipurpose field that displays information about BGP peer sessions. The field's contents depend upon whether a session is established and whether it was established in the main routing device or in a routing instance.</p> <ul style="list-style-type: none"> If a peer is not established, the field shows the state of the peer session: Active, Connect, or Idle. If a BGP session is established in the main routing device, the field shows the number of active, received, accepted, and damped routes that are received from a neighbor and appear in the inet.0 (main) and inet.2 (multicast) routing tables. For example, 8/10/10/2 and 2/4/4/0 indicate the following: <ul style="list-style-type: none"> 8 active routes, 10 received routes, 10 accepted routes, and 2 damped routes from a BGP peer appear in the inet.0 routing table. 2 active routes, 4 received routes, 4 accepted routes, and no damped routes from a BGP peer appear in the inet.2 routing table. | summary |
| ip-addresses | List of peers who are members of the group. The address is followed by the peer's port number. | All levels |
| Route Queue Timer | Number of seconds until queued routes are sent. If this time has already elapsed, this field displays the number of seconds by which the updates are delayed. | detail |
| Route Queue | Number of prefixes that are queued up for sending to the peers in the group. | detail |
| inet.number | <p>Number of active, received, accepted, and damped routes in the routing table. For example, inet.0: 7/10/9/0 indicates the following:</p> <ul style="list-style-type: none"> 7 active routes, 10 received routes, 9 accepted routes, and no damped routes from a BGP peer appear in the inet.0 routing table. | none |

Table 11: show bgp group Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------------|---|---------------------|
| Table inet.number | Information about the routing table. <ul style="list-style-type: none"> • Received prefixes—Total number of prefixes from the peer, both active and inactive, that are in the routing table. • Active prefixes—Number of prefixes received from the peer that are active in the routing table. • Suppressed due to damping—Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols. • Advertised prefixes—Number of prefixes advertised to a peer. • Received external prefixes—Total number of prefixes from the external BGP (EBGP) peers, both active and inactive, that are in the routing table. • Active external prefixes—Number of prefixes received from the EBGP peers that are active in the routing table. • Externals suppressed—Number of routes received from EBGP peers currently inactive because of damping or other reasons. • Received internal prefixes—Total number of prefixes from the IBGP peers, both active and inactive, that are in the routing table. • Active internal prefixes—Number of prefixes received from the IBGP peers that are active in the routing table. • Internals suppressed—Number of routes received from IBGP peers currently inactive because of damping or other reasons. • RIB State—Status of the graceful restart process for this routing table: BGP restart is complete, BGP restart in progress, VPN restart in progress, or VPN restart is complete. | detail |
| Groups | Total number of groups. | All levels |
| Peers | Total number of peers. | All levels |
| External | Total number of external peers. | All levels |
| Internal | Total number of internal peers. | All levels |
| Down peers | Total number of unavailable peers. | All levels |
| Flaps | Total number of flaps that occurred. | All levels |
| Table | Name of a routing table. | brief , none |
| Tot Paths | Total number of routes. | brief , none |
| Act Paths | Number of active routes. | brief , none |
| Suppressed | Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols. | brief , none |

Table 11: show bgp group Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------|---|--------------------|
| History | Number of withdrawn routes stored locally to keep track of damping history. | brief, none |
| Damp State | Number of active routes with a figure of merit greater than zero, but lower than the threshold at which suppression occurs. | brief, none |
| Pending | Routes being processed by the BGP import policy. | brief, none |
| Group | Group the peer belongs to in the BGP configuration. | detail |
| Receive mask | Mask of the received target included in the advertised route. | detail |
| Entries | Number of route entries received. | detail |
| Target | Route target that is to be passed by route-target filtering. If a route advertised from the provider edge (PE) routing device matches an entry in the route-target filter, the route is passed to the peer. | detail |
| Mask | Mask which specifies that the peer receive routes with the given route target. | detail |

Sample Output

show bgp group

```
user@host> show bgp group
Groups: 2 Peers: 2 External: 0 Internal: 2 Down peers: 1 Flaps: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending

inet.0
0 0 0 0 0 0

bgp.l3vpn.0
0 0 0 0 0 0

bgp.rtarget.0
2 0 0 0 0 0
```

show bgp group brief

```
user@host> show bgp group brief
Groups: 2 Peers: 2 External: 0 Internal: 2 Down peers: 1 Flaps: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending

inet.0
0 0 0 0 0 0

bgp.l3vpn.0
0 0 0 0 0 0

bgp.rtarget.0
2 0 0 0 0 0
```

show bgp group detail

```
user@host> show bgp group detail
Group Type: Internal AS: 1 Local AS: 1
Name: ibgp Index: 0 Flags: <Export Eval>
Holdtime: 0
Total peers: 3 Established: 0
22.0.0.2
22.0.0.8
22.0.0.5

Groups: 1 Peers: 3 External: 0 Internal: 3 Down peers: 3 Flaps: 3
Table bgp.l3vpn.0
Received prefixes: 0
Accepted prefixes: 0
Active prefixes: 0
Suppressed due to damping: 0
Received external prefixes: 0
Active external prefixes: 0
Externals suppressed: 0
Received internal prefixes: 0
Active internal prefixes: 0
Internals suppressed: 0
RIB State: BGP restart is complete
RIB State: VPN restart is complete
Table bgp.mdt.0
Received prefixes: 0
Accepted prefixes: 0
Active prefixes: 0
Suppressed due to damping: 0
Received external prefixes: 0
Active external prefixes: 0
```

```

Externals suppressed:      0
Received internal prefixes: 0
Active internal prefixes:  0
Internals suppressed:      0
RIB State: BGP restart is complete
RIB State: VPN restart is complete
Table VPN-A.inet.0
Received prefixes:         0
Accepted prefixes:         0
Active prefixes:           0
Suppressed due to damping: 0
Received external prefixes: 0
Active external prefixes:  0
Externals suppressed:      0
Received internal prefixes: 0
Active internal prefixes:  0
Internals suppressed:      0
RIB State: BGP restart is complete
RIB State: VPN restart is complete
Table VPN-A.mdt.0
Received prefixes:         0
Accepted prefixes:         0
Active prefixes:           0
Suppressed due to damping: 0
Received external prefixes: 0
Active external prefixes:  0
Externals suppressed:      0
Received internal prefixes: 0
Active internal prefixes:  0
Internals suppressed:      0
RIB State: BGP restart is complete
RIB State: VPN restart is complete

```

show bgp group rtf detail

```

user@host> show bgp group rtf detail
Group: internal (group-index: 0)
  Receive mask: 00000002
  Table: bgp.rtarget.0
    Target      Mask      Entries: 2
    100:100/64  00000002
    200:201/64  (Group)
Group: internal (group-index: 1)
  Table: bgp.rtarget.0
    Target      Mask      Entries: 1
    200:201/64  (Group)

```

show bgp group summary

```

user@host> show bgp group summary
Group      Type      Peers      Established      Active/Received/Accepted/Damped
ibgp       Internal  3          0
Groups: 1  Peers: 3   External: 0   Internal: 3   Down peers: 3   Flaps: 3
bgp.l3vpn.0 : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0
bgp.mdt.0   : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0
VPN-A.inet.0 : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0
VPN-A.mdt.0 : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0

```

show bgp group traffic-statistics

| | |
|---------------------------------|---|
| Syntax | show bgp group traffic-statistics <brief detail> <group-name> <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display the traffic statistics for configured Border Gateway Protocol (BGP) groups. |
| Options | <p>none—Display traffic statistics for all BGP groups.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>group-name—(Optional) Display BGP traffic statistics for only the specified group.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show bgp group traffic-statistics (Per-Group-Label Not Configured) on page 52 show bgp group traffic-statistics (Per-Group-Label Configured) on page 52 |
| Output Fields | Table 12 on page 51 describes the output fields for the show bgp group traffic-statistics command. Output fields are listed in the approximate order in which they appear. |

Table 12: show bgp group traffic-statistics Output Fields

| Field Name | Field Description |
|--------------------|--|
| Group name | Name of a specific BGP group. |
| Group Index | Index number for the BGP group. |
| NLRI | Network layer reachability information (NLRI) indicating the source of the traffic statistics for the BGP group. |
| FEC | Forwarding equivalence classes (FECs) associated with the BGP group. |
| Packets | Number of packets sent through each FEC. |
| Bytes | Number of bytes transmitted through each FEC. |
| EgressAS | Autonomous system (AS) number of the egress router. |
| AdvLabel | Label associated with each FEC. |

Sample Output

show bgp group
traffic-statistics
(Per-Group-Label
Not Configured)

```
user@host> show bgp group traffic-statistics
Group Name: ext1      Group Index: 0      NLRI: inet-labeled-unicast
FEC                   Packets             Bytes      EgressAS  AdvLabel
10.255.245.55         0                   0          I         100224
10.255.245.57         0                   0          I         100240
100.101.0.0           550                 48400      25        100256
100.102.0.0           550                 48400      25        100256
100.103.0.0           550                 48400      25        100272
100.104.0.0           550                 48400      25        100272
192.168.25.0          0                   0          I         100288

Group Name: ext2      Group Index: 1      NLRI: inet-labeled-unicast
FEC                   Packets             Bytes      EgressAS  AdvLabel
10.255.245.55         0                   0          I         100224
10.255.245.57         0                   0          I         100240
100.101.0.0           550                 48400      25        100256
100.102.0.0           550                 48400      25        100256
100.103.0.0           550                 48400      25        100272
100.104.0.0           550                 48400      25        100272
192.168.25.0          0                   0          I         100288
```

show bgp group
traffic-statistics
(Per-Group-Label
Configured)

```
user@host> show bgp group traffic-statistics
Group Name: ext1      Group Index: 0      NLRI: inet-labeled-unicast
FEC                   Packets             Bytes      EgressAS  AdvLabel
10.255.245.55         0                   0          I         100384
10.255.245.57         0                   0          I         100400
100.101.0.0           101                 8888       25        100416
100.102.0.0           101                 8888       25        100416
100.103.0.0           0                   0          25        100432
100.104.0.0           0                   0          25        100432
192.168.25.0          0                   0          I         100448

Group Name: ext2      Group Index: 1      NLRI: inet-labeled-unicast
FEC                   Packets             Bytes      EgressAS  AdvLabel
10.255.245.55         0                   0          I         100304
10.255.245.57         0                   0          I         100320
100.101.0.0           0                   0          25        100336
100.102.0.0           0                   0          25        100336
100.103.0.0           101                 8888       25        100352
100.104.0.0           101                 8888       25        100352
192.168.25.0          0                   0          I         100368
```

show bgp neighbor

| | |
|---|---|
| Syntax | <pre>show bgp neighbor <exact-instance <i>instance-name</i>> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <neighbor-address> <orf (detail <i>neighbor-address</i>)</pre> |
| Syntax (EX Series Switch and QFX Series) | <pre>show bgp neighbor <instance <i>instance-name</i>> <exact-instance <i>instance-name</i>> <neighbor-address> <orf (<i>neighbor-address</i> detail)</pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>orf option introduced in Junos OS Release 9.2.</p> <p>exact-instance option introduced in Junos OS Release 11.4.</p> |
| Description | Display information about BGP peers. |
| Options | <p>none—Display information about all BGP peers.</p> <p>exact-instance <i>instance-name</i>—(Optional) Display information for the specified instance only.</p> <p>instance <i>instance-name</i>—(Optional) Display information about BGP peers for all routing instances whose name begins with this string (for example, cust1, cust11, and cust111 are all displayed when you run the show bgp neighbor instance cust1 command).</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>neighbor-address—(Optional) Display information for only the BGP peer at the specified IP address.</p> <p>orf (detail <i>neighbor-address</i>)—(Optional) Display outbound route-filtering information for all BGP peers or only for the BGP peer at the specified IP address. The default is to display brief output. Use the detail option to display detailed output.</p> |
| Additional Information | For information about the local-address , nlri , hold-time , and preference statements, see the Junos OS Routing Protocols Configuration Guide. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • clear bgp neighbor on page 39 |

List of Sample Output [show bgp neighbor on page 61](#)
[show bgp neighbor \(CLNS\) on page 61](#)
[show bgp neighbor \(Layer 2 VPN\) on page 62](#)
[show bgp neighbor \(Layer 3 VPN\) on page 64](#)
[show bgp neighbor neighbor-address on page 65](#)
[show bgp neighbor neighbor-address on page 65](#)
[show bgp neighbor orf neighbor-address detail on page 66](#)

Output Fields [Table 13 on page 54](#) describes the output fields for the **show bgp neighbor** command. Output fields are listed in the approximate order in which they appear.

Table 13: show bgp neighbor Output Fields

| Field Name | Field Description |
|--------------|---|
| Peer | Address of the BGP neighbor. The address is followed by the neighbor port number. |
| AS | AS number of the peer. |
| Local | Address of the local routing device. The address is followed by the peer port number. |
| Type | Type of peer: Internal or External . |
| State | <p>Current state of the BGP session:</p> <ul style="list-style-type: none"> • Active—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an Open message. • Connect—BGP is waiting for the transport protocol connection to be completed. • Established—The BGP session has been established, and the peers are exchanging update messages. • Idle—This is the first stage of a connection. BGP is waiting for a Start event. • OpenConfirm—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message. • OpenSent—BGP has sent an open message and is waiting to receive an open message from the peer. |
| Flags | <p>Internal BGP flags:</p> <ul style="list-style-type: none"> • Aggregate Label—BGP has aggregated a set of incoming labels (labels received from the peer) into a single forwarding label. • CleanUp—The peer session is being shut down. • Delete—This peer has been deleted. • Idled—This peer has been permanently idled. • ImportEval—At the last commit operation, this peer was identified as needing to reevaluate all received routes. • Initializing—The peer session is initializing. • SendRtn—Messages are being sent to the peer. • Sync—This peer is synchronized with the rest of the peer group. • TryConnect—Another attempt is being made to connect to the peer. • Unconfigured—This peer is not configured. • WriteFailed—An attempt to write to this peer failed. |

Table 13: show bgp neighbor Output Fields (*continued*)

| Field Name | Field Description |
|-------------------|---|
| Last state | <p>Previous state of the BGP session:</p> <ul style="list-style-type: none"> • Active—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an Open message. • Connect—BGP is waiting for the transport protocol connection to be completed. • Established—The BGP session has been established, and the peers are exchanging update messages. • Idle—This is the first stage of a connection. BGP is waiting for a Start event. • OpenConfirm—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message. • OpenSent—BGP has sent an open message and is waiting to receive an open message from the peer. |
| Last event | <p>Last activity that occurred in the BGP session:</p> <ul style="list-style-type: none"> • Closed—The BGP session closed. • ConnectRetry—The transport protocol connection failed, and BGP is trying again to connect. • HoldTime—The session ended because the hold timer expired. • KeepAlive—The local routing device sent a BGP keepalive message to the peer. • Open—The local routing device sent a BGP open message to the peer. • OpenFail—The local routing device did not receive an acknowledgment of a BGP open message from the peer. • RecvKeepAlive—The local routing device received a BGP keepalive message from the peer. • RecvNotify—The local routing device received a BGP notification message from the peer. • RecvOpen—The local routing device received a BGP open message from the peer. • RecvUpdate—The local routing device received a BGP update message from the peer. • Start—The peering session started. • Stop—The peering session stopped. • TransportError—A TCP error occurred. |
| Last error | <p>Last error that occurred in the BGP session:</p> <ul style="list-style-type: none"> • Cease—An error occurred, such as a version mismatch, that caused the session to close. • Finite State Machine Error—In setting up the session, BGP received a message that it did not understand. • Hold Time Expired—The session's hold time expired. • Message Header Error—The header of a BGP message was malformed. • Open Message Error—A BGP open message contained an error. • None—No errors occurred in the BGP session. • Update Message Error—A BGP update message contained an error. |
| Export | Name of the export policy that is configured on the peer. |
| Import | Name of the import policy that is configured on the peer. |

Table 13: show bgp neighbor Output Fields (*continued*)

| Field Name | Field Description |
|--|--|
| Options | Configured BGP options: <ul style="list-style-type: none"> • AddressFamily—Configured address family: inet or inet-vpn. • AuthKeyChain—Authentication key change is enabled. • DropPathAttributes—Certain path attributes are configured to be dropped from neighbor updates during inbound processing. • GracefulRestart—Graceful restart is configured. • HoldTime—Hold time configured with the hold-time statement. The hold time is three times the interval at which keepalive messages are sent. • IgnorePathAttributes—Certain path attributes are configured to be ignored in neighbor updates during inbound processing. • Local Address—Address configured with the local-address statement. • Multihop—Allow BGP connections to external peers that are not on a directly connected network. • NLRI—Configured MBGP state for the BGP group: multicast, unicast, or both if you have configured nlri any. • Peer AS—Configured peer autonomous system (AS). • Preference—Preference value configured with the preference statement. • Refresh—Configured to refresh automatically when the policy changes. • Rib-group—Configured routing table group. |
| Path-attributes dropped | Path attribute codes that are dropped from neighbor updates. |
| Path-attributes ignored | Path attribute codes that are ignored during neighbor updates. |
| Authentication key change | (appears only if the authentication-keychain statement has been configured) Name of the authentication keychain enabled. |
| Authentication algorithm | (appears only if the authentication-algorithm statement has been configured) Type of authentication algorithm enabled: hmac or md5 . |
| Address families configured | Names of configured address families for the VPN. |
| Local Address | Address of the local routing device. |
| Remove-private options | Options associated with the remove-private statement. |
| Holdtime | Hold time configured with the hold-time statement. The hold time is three times the interval at which keepalive messages are sent. |
| Flags for NLRI inet-label-unicast | Flags related to labeled-unicast: <ul style="list-style-type: none"> • TrafficStatistics—Collection of statistics for labeled-unicast traffic is enabled. |

Table 13: show bgp neighbor Output Fields (*continued*)

| Field Name | Field Description |
|--|---|
| Traffic statistics | Information about labeled-unicast traffic statistics: <ul style="list-style-type: none"> • Options—Options configured for collecting statistics about labeled-unicast traffic. • File—Name and location of statistics log files. • size—Size of all the log files, in bytes. • files—Number of log files. |
| Traffic Statistics Interval | Time between sample periods for labeled-unicast traffic statistics, in seconds. |
| Preference | Preference value configured with the preference statement. |
| Outbound Timer | Time for which the route is available in Junos OS routing table before it is exported to BGP. This field is displayed in the output only if the out-delay parameter is configured to a non-zero value. |
| Number of flaps | Number of times the BGP session has gone down and then come back up. |
| Peer ID | Router identifier of the peer. |
| Group index | Index number for the BGP peer group. The index number differentiates between groups when a single BGP group is split because of different configuration options at the group and peer levels. |
| Peer index | Index that is unique within the BGP group to which the peer belongs. |
| Local ID | Router identifier of the local routing device. |
| Local Interface | Name of the interface on the local routing device. |
| Active holdtime | Hold time that the local routing device negotiated with the peer. |
| Keepalive Interval | Keepalive interval, in seconds. |
| BFD | Status of BFD failure detection. |
| Local Address | Name of directly connected interface over which direct EBGP peering is established. |
| NLRI for restart configured on peer | Names of address families configured for restart. |
| NLRI advertised by peer | Address families supported by the peer: unicast or multicast . |
| NLRI for this session | Address families being used for this session. |
| Peer supports Refresh capability | Remote peer's ability to send and request full route table readvertisement (route refresh capability). For more information, see RFC 2918, <i>Route Refresh Capability for BGP-4</i> . |
| Restart time configured on peer | Configured time allowed for restart on the neighbor. |

Table 13: show bgp neighbor Output Fields (*continued*)

| Field Name | Field Description |
|--|---|
| Stale routes from peer are kept for | When graceful restart is negotiated, the maximum time allowed to hold routes from neighbors after the BGP session has gone down. |
| Peer does not support Restarter functionality | Graceful restart restarter-mode is disabled on the peer. |
| Peer does not support Receiver functionality | Graceful restart helper-mode is disabled on the peer. |
| Restart time requested by this peer | Restart time requested by this neighbor during capability negotiation. |
| Restart flag received from the peer | When this field appears, the BGP speaker has restarted (Restarting), and this peer should not wait for the end-of-rib marker from the speaker before advertising routing information to the speaker. |
| NLRI that peer supports restart for | Neighbor supports graceful restart for this address family. |
| NLRI peer can save forwarding state | Neighbor supporting this address family saves all forwarding states. |
| NLRI that peer saved forwarding for | Neighbor saves all forwarding states for this address family. |
| NLRI that restart is negotiated for | Router supports graceful restart for this address family. |
| NLRI of received end-of-rib markers | Address families for which end-of-routing-table markers are received from the neighbor. |
| NLRI of all end-of-rib markers sent | Address families for which end-of-routing-table markers are sent to the neighbor. |
| Peer supports 4 byte AS extension (peer-as 1) | Peer understands 4-byte AS numbers in BGP messages. The peer is running Junos OS Release 9.1 or later. |
| NLRIs for which peer can receive multiple paths | Appears in the command output of the local router if the downstream peer is configured to receive multiple BGP routes to a single destination, instead of only receiving the active route. Possible value is inet-unicast . |
| NLRIs for which peer can send multiple paths: inet-unicast | Appears in the command output of the local router if the upstream peer is configured to send multiple BGP routes to a single destination, instead of only sending the active route. Possible value is inet-unicast . |

Table 13: show bgp neighbor Output Fields (*continued*)

| Field Name | Field Description |
|-------------------------------|--|
| Table inet.number | <p>Information about the routing table:</p> <ul style="list-style-type: none"> • RIB State—BGP is in the graceful restart process for this routing table: restart is complete or restart in progress. • Bit—Number that represents the entry in the routing table for this peer. • Send state—State of the BGP group: in sync, not in sync, or not advertising. • Active prefixes—Number of prefixes received from the peer that are active in the routing table. • Received prefixes—Total number of prefixes from the peer, both active and inactive, that are in the routing table. • Accepted prefixes—Total number of prefixes from the peer that have been accepted by a routing policy. • Suppressed due to damping—Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols. |
| Last traffic (seconds) | Last time any traffic was received from the peer or sent to the peer, and the last time the local routing device checked. |
| Input messages | Messages that BGP has received from the receive socket buffer, showing the total number of messages, number of update messages, number of times a policy is changed and refreshed, and the buffer size in octets. The buffer size is 16 KB. |
| Output messages | Messages that BGP has written to the transmit socket buffer, showing the total number of messages, number of update messages, number of times a policy is changed and refreshed, and the buffer size in octets. The buffer size is 16 KB. |
| Input dropped path attributes | <p>Information about dropped path attributes:</p> <ul style="list-style-type: none"> • Code—Path attribute code. • Count—Path attribute count. |
| Input ignored path attributes | <p>Information about ignored path attributes:</p> <ul style="list-style-type: none"> • Code—Path attribute code. • Count—Path attribute count. |
| Output queue | Number of BGP packets that are queued to be transmitted to a particular neighbor for a particular routing table. Output queue 0 is for unicast NLRIs, and queue 1 is for multicast NLRIs. |
| Trace options | Configured tracing of BGP protocol packets and operations. |
| Trace file | Name of the file to receive the output of the tracing operation. |
| Filter Updates rcv | <p>(orf option only) Number of outbound-route filters received for each configured address family.</p> <p>NOTE: The counter is cumulative. For example, the counter is increased after the remote peer either resends or clears the outbound route filtering prefix list.</p> |

Table 13: show bgp neighbor Output Fields (*continued*)

| Field Name | Field Description |
|--------------------------------|---|
| Immediate | (orf option only) Number of route updates received with the immediate flag set. The immediate flag indicates that the BGP peer should readvertise the updated routes. NOTE: The counter is cumulative. For example, the counter is increased after the remote peer either resends or clears the outbound route filtering prefix list. |
| Filter | (orf option only) Type of prefix filter received: prefix-based or extended-community . |
| Received filter entries | (orf option only) List of received filters displayed. |
| seq | (orf option only) Numerical order assigned to this prefix entry among all the received outbound route filter prefix entries. |
| prefix | (orf option only) Address for the prefix entry that matches the filter. |
| minlength | (orf option only) Minimum prefix length, in bits, required to match this prefix. |
| maxlength | (orf option only) Maximum prefix length, in bits, required to match this prefix. |
| match | (orf option only) For this prefix match, whether to permit or deny route updates. |

Sample Output

show bgp neighbor

```

user@host > show bgp neighbor
Peer: 10.255.7.250+179 AS 10   Local: 10.255.7.248+63740 AS 10
  Type: Internal   State: Established   Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Export: [ redist_static ]
  Options: <Preference LocalAddress PeerAS Refresh>
  Local Address: 10.255.7.248 Holdtime: 90 Preference: 170 Outbound Timer: 50
  Number of flaps: 0
  Peer ID: 10.255.7.250   Local ID: 10.255.7.248   Active Holdtime: 90
  Keepalive Interval: 30   Group index: 0   Peer index: 0
  BFD: disabled, down
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Peer supports 4 byte AS extension (peer-as 10)
  Peer does not support Addpath
  Table inet.0 Bit: 10000
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:          1
    Received prefixes:        1
    Accepted prefixes:        1
    Suppressed due to damping: 0
    Advertised prefixes:      1
  Last traffic (seconds): Received 9   Sent 5   Checked 5
  Input messages: Total 36   Updates 2   Refreshes 0   Octets 718
  Output messages: Total 37   Updates 1   Refreshes 0   Octets 796
  Output Queue[0]: 0

Peer: 10.255.162.214+52193 AS 100 Local: 10.255.167.205+179 AS 100
  Type: Internal   State: Established (route reflector client)Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress Cluster AddressFamily Rib-group Refresh>
  Address families configured: inet-unicast inet-vpn-unicast route-target
  Local Address: 10.255.167.205 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.255.162.214   Local ID: 10.255.167.205   Active Holdtime: 90
  Keepalive Interval: 30   Group index: 0   Peer index: 1

```

show bgp neighbor (CLNS)

```

user@host> show bgp neighbor
Peer: 10.245.245.1+179 AS 200 Local: 10.245.245.3+3770 AS 100
  Type: External   State: Established   Flags: <ImportEval Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Options: <Multihop Preference LocalAddress HoldTime AddressFamily PeerAS
  Rib-group Refresh>
  Address families configured: iso-vpn-unicast
  Local Address: 10.245.245.3 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.245.245.1   Local ID: 10.245.245.3   Active Holdtime: 90

```

```

Keepalive Interval: 30          Peer index: 0
NLRI advertised by peer: iso-vpn-unicast
NLRI for this session: iso-vpn-unicast
Peer supports Refresh capability (2)
Table bgp.isovpn.0 Bit: 10000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:              3
  Received prefixes:            3
  Suppressed due to damping:    0
  Advertised prefixes:          3
Table aaaa.iso.0
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: not advertising
  Active prefixes:              3
  Received prefixes:            3
  Suppressed due to damping:    0
Last traffic (seconds): Received 6   Sent 5   Checked 5
Input messages: Total 1736   Updates 4   Refreshes 0   Octets 33385
Output messages: Total 1738   Updates 3   Refreshes 0   Octets 33305
Output Queue[0]: 0
Output Queue[1]: 0

```

show bgp neighbor (Layer 2 VPN)

```

user@host> show bgp neighbor
Peer: 10.69.103.2      AS 65100 Local: 10.69.103.1      AS 65103
  Type: External      State: Active      Flags: <ImportEval>
  Last State: Idle    Last Event: Start
  Last Error: None
  Export: [ BGP-INET-import ]
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily PeerAS
Refresh>
  Address families configured: inet-unicast
  Local Address: 10.69.103.1 Holdtime: 90 Preference: 170
  Number of flaps: 0
Peer: 10.69.104.2      AS 65100 Local: 10.69.104.1      AS 65104
  Type: External      State: Active      Flags: <ImportEval>
  Last State: Idle    Last Event: Start
  Last Error: None
  Export: [ BGP-L-import ]
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily PeerAS
Refresh>
  Address families configured: inet-labeled-unicast
  Local Address: 10.69.104.1 Holdtime: 90 Preference: 170
  Number of flaps: 0
Peer: 10.255.14.182+179 AS 69   Local: 10.255.14.176+2131 AS 69
  Type: Internal      State: Established  Flags: <ImportEval>
  Last State: OpenConfirm Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
Rib-group Refresh>
  Address families configured: inet-vpn-unicast l2vpn
  Local Address: 10.255.14.176 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.255.14.182   Local ID: 10.255.14.176   Active Holdtime: 90
  Keepalive Interval: 30
  NLRI for restart configured on peer: inet-vpn-unicast l2vpn
  NLRI advertised by peer: inet-vpn-unicast l2vpn
  NLRI for this session: inet-vpn-unicast l2vpn
  Peer supports Refresh capability (2)

```



```

Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-vpn-unicast 12vpn
NLRI peer can save forwarding state: inet-vpn-unicast 12vpn
NLRI that peer saved forwarding for: inet-vpn-unicast 12vpn
NLRI that restart is negotiated for: inet-vpn-unicast 12vpn
NLRI of received end-of-rib markers: inet-vpn-unicast 12vpn
Table bgp.13vpn.0 Bit: 10000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          10
  Received prefixes:        10
  Suppressed due to damping: 0
Table bgp.12vpn.0 Bit: 20000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        1
  Suppressed due to damping: 0
Table BGP-INET.inet.0 Bit: 30000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2
  Suppressed due to damping: 0
Table BGP-L.inet.0 Bit: 40000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2
  Suppressed due to damping: 0
Table LDP.inet.0 Bit: 50000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        1
  Suppressed due to damping: 0
Table OSPF.inet.0 Bit: 60000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2
  Suppressed due to damping: 0
Table RIP.inet.0 Bit: 70000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2
  Suppressed due to damping: 0
Table STATIC.inet.0 Bit: 80000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync

```

```

Active prefixes:          1
Received prefixes:       1
Suppressed due to damping: 0
Table L2VPN.l2vpn.0 Bit: 90000
RIB State: BGP restart is complete
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:          1
Received prefixes:       1
Suppressed due to damping: 0
Last traffic (seconds): Received 0      Sent 0      Checked 0
Input messages: Total 14      Updates 13     Refreshes 0      Octets 1053
Output messages: Total 3      Updates 0      Refreshes 0      Octets 105
Output Queue[0]: 0
Output Queue[1]: 0
Output Queue[2]: 0
Output Queue[3]: 0
Output Queue[4]: 0
Output Queue[5]: 0
Output Queue[6]: 0
Output Queue[7]: 0
Output Queue[8]: 0

```

show bgp neighbor (Layer 3 VPN)

```

user@host> show bgp neighbor
Peer: 4.4.4.4+179      AS 10045 Local: 5.5.5.5+1214      AS 10045
Type: Internal      State: Established      Flags: <ImportEval>
Last State: OpenConfirm      Last Event: RecvKeepAlive
Last Error: None
Export: [ match-all ] Import: [ match-all ]
Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
Rib-group Refresh>
Address families configured: inet-vpn-unicast
Local Address: 5.5.5.5 Holdtime: 90 Preference: 170
Flags for NLRI inet-labeled-unicast: TrafficStatistics
Traffic Statistics: Options: all File: /var/log/bstat.log
                        size 131072 files 10

Traffic Statistics Interval: 60
Number of flaps: 0
Peer ID: 192.168.1.110      Local ID: 192.168.1.111      Active Holdtime: 90
Keepalive Interval: 30
NLRI for restart configured on peer: inet-vpn-unicast
NLRI advertised by peer: inet-vpn-unicast
NLRI for this session: inet-vpn-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-vpn-unicast
NLRI peer can save forwarding state: inet-vpn-unicast
NLRI that peer saved forwarding for: inet-vpn-unicast
NLRI that restart is negotiated for: inet-vpn-unicast
NLRI of received end-of-rib markers: inet-vpn-unicast
NLRI of all end-of-rib markers sent: inet-vpn-unicast
Table bgp.l3vpn.0 Bit: 10000
RIB State: BGP restart is complete
RIB State: VPN restart is complete
Send state: in sync
Active prefixes:          2
Received prefixes:       2
Suppressed due to damping: 0
Table vpn-green.inet.0 Bit: 20001

```

```

RIB State: BGP restart is complete
RIB State: VPN restart is complete
Send state: in sync
Active prefixes:          2
Received prefixes:        2
Suppressed due to damping: 0
Last traffic (seconds): Received 15   Sent 20   Checked 20
Input messages:  Total 40   Updates 2   Refreshes 0   Octets 856
Output messages: Total 44   Updates 2   Refreshes 0   Octets 1066
Output Queue[0]: 0
Output Queue[1]: 0
Trace options: detail packets
Trace file: /var/log/bgpr.log size 131072 files 10

```

show bgp neighbor neighbor-address

```

user@host> show bgp neighbor 192.168.1.111
Peer: 10.255.245.12+179 AS 35 Local: 10.255.245.13+2884 AS 35
Type: Internal State: Established (route reflector client)Flags: <Sync>
Last State: OpenConfirm Last Event: RecvKeepAlive
Last Error: None
Options: <Preference LocalAddress HoldTime Cluster AddressFamily Rib-group
Refresh>
Address families configured: inet-vpn-unicast inet-labeled-unicast
Local Address: 10.255.245.13 Holdtime: 90 Preference: 170
Flags for NLRI inet-vpn-unicast: AggregateLabel
Flags for NLRI inet-labeled-unicast: AggregateLabel
Number of flaps: 0
Peer ID: 10.255.245.12 Local ID: 10.255.245.13 Active Holdtime: 90
Keepalive Interval: 30
BFD: disabled
NLRI advertised by peer: inet-vpn-unicast inet-labeled-unicast
NLRI for this session: inet-vpn-unicast inet-labeled-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 300
Stale routes from peer are kept for: 60
Restart time requested by this peer: 300
NLRI that peer supports restart for: inet-unicast inet6-unicast
NLRI that restart is negotiated for: inet-unicast inet6-unicast
NLRI of received end-of-rib markers: inet-unicast inet6-unicast
NLRI of all end-of-rib markers sent: inet-unicast inet6-unicast
Table inet.0 Bit: 10000
RIB State: restart is complete
Send state: in sync
Active prefixes: 4
Received prefixes: 6
Suppressed due to damping: 0
Table inet6.0 Bit: 20000
RIB State: restart is complete
Send state: in sync
Active prefixes: 0
Received prefixes: 2
Suppressed due to damping: 0
Last traffic (seconds): Received 3 Sent 3 Checked 3
Input messages:  Total 9 Updates 6 Refreshes 0 Octets 403
Output messages: Total 7 Updates 3 Refreshes 0 Octets 365
Output Queue[0]: 0
Output Queue[1]: 0
Trace options: detail packets
Trace file: /var/log/bgpr size 131072 files 10

user@host> show bgp neighbor 192.168.4.222

```

**show bgp neighbor
neighbor-address**

```

Peer: 192.168.4.222+4902 AS 65501 Local: 192.168.4.221+179 AS 65500
Type: External      State: Established      Flags: <Sync>
Last State: OpenConfirm  Last Event: RecvKeepAlive
Last Error: Cease
Export: [ export-policy ] Import: [ import-policy ]
Options: <Preference HoldTime AddressFamily PeerAS PrefixLimit Refresh>
Address families configured: inet-unicast inet-multicast
Holdtime: 60000 Preference: 170
Number of flaps: 4
Last flap event: RecvUpdate
Error: 'Cease' Sent: 5 Recv: 0
Peer ID: 10.255.245.6      Local ID: 10.255.245.5      Active Holdtime: 60000
Keepalive Interval: 20000      Peer index: 0
BFD: disabled, down
Local Interface: fxp0.0
NLRI advertised by peer: inet-unicast inet-multicast
NLRI for this session: inet-unicast inet-multicast
Peer supports Refresh capability (2)
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          8
  Received prefixes:        10
  Accepted prefixes:        10
  Suppressed due to damping: 0
  Advertised prefixes:      3
Table inet.2 Bit: 20000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        0
  Accepted prefixes:        0
  Suppressed due to damping: 0
  Advertised prefixes:      0
Last traffic (seconds): Received 357 Sent 357 Checked 357
Input messages: Total 4 Updates 2 Refreshes 0 Octets 211
Output messages: Total 4 Updates 1 Refreshes 0 Octets 147
Output Queue[0]: 0
Output Queue[1]: 0
Trace options: all
Trace file: /var/log/bgp size 10485760 files 10

```

**show bgp neighbor orf
neighbor-address
detail**

```

user@host > show bgp neighbor orf 192.168.165.56 detail
Peer: 192.168.165.56+179 Type: External
Group: ext1

inet-unicast
  Filter updates rcv:          1 Immediate:          1
  Filter: prefix-based receive
  Received filter entries:
    seq 1: prefix 2.2.2.2/32: minlen 32: maxlen 32: match deny:

inet6-unicast
  Filter updates rcv:          0 Immediate:          1
  Filter: prefix-based receive
  Received filter entries:
    *.*

```

show bgp replication

| | |
|---------------------------------|--|
| Syntax | show bgp replication |
| Release Information | Command introduced in JUNOS Release 8.5. Command introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | Displays the status of BGP state replication between the master and backup Routing Engines on devices that have nonstop active routing configured on them. |
| Options | This command has no options. |
| Required Privilege Level | view |
| List of Sample Output | show bgp replication (for Master) on page 68 show bgp replication (for Backup) on page 68 |
| Output Fields | Table 14 on page 67 lists the output fields for the show bgp replication command. Output fields are listed in the approximate order in which they appear. |

Table 14: show bgp replication Output Fields

| Field Name | Field Description |
|--------------------------------|---|
| session state | State of the current internal BGP state replication session, Up or Down, and the duration for which the session has been in the indicated state. |
| flaps | Total number of flaps that occurred. |
| protocol state | Current state of the protocol operation, Active, Connect, Idle, and the duration for which the protocol has been in the indicated state. |
| synchronization state | Synchronization state at the time of executing the command. The states can be: <ul style="list-style-type: none"> • Idle • Neighbor—Indicates that the neighbor state synchronization is in progress. • AckWait—Indicates that the request processing is over. • ORF—Indicates that the outbound routing filter synchronization is in progress. • RIB—Indicates that the routing table synchronization is in progress. • Complete |
| number of peers waiting | Total number of peers waiting for various messages: <ul style="list-style-type: none"> • AckWait—Number of peers waiting for a connection establishment or completed acknowledgment messages. • SoWait—Number of peers waiting for TCP socket-related operations. • Scheduled—Number of peers being synchronized. |

Table 14: show bgp replication Output Fields (*continued*)

| Field Name | Field Description |
|--------------------------|--|
| messages sent | <p>Number of various types of messages that have been sent since internal replication session became active:</p> <ul style="list-style-type: none"> • Open—Number of Open messages sent. • Establish—Number of connection establishment acknowledgment messages sent. • Update—Number of update messages sent. • Error—Number of error messages sent. • Complete—Number of connection complete acknowledgment messages sent. |
| messages received | <p>Total number of messages received:</p> <ul style="list-style-type: none"> • Open—Number of Open messages received. • Request—Number of request messages received: <ul style="list-style-type: none"> • Wildcard—Number of requests received that used wildcards in the target address. • Targeted—Number of requests received that used a specific address. • EstablishAck—Number of connection establishment acknowledgement messages received. • CompleteAck—Number of connection completed acknowledgement messages received. |

Sample Output

show bgp replication (for Master)

```
user@host> show bgp replication
Synchronization master:
  Session state: Up, Since: 44:07
  Flaps: 0
  Protocol state: Idle, Since: 14
  Synchronization state: Complete
  Number of peers waiting: AckWait: 0, SoWait: 0, Scheduled: 0
  Messages sent: Open 1, Establish 924, Update 381, Error 60, Complete 114
  Messages received: Open 1, Request 1 wildcard 113 targeted, EstablishAck 924,
  CompleteAck 114
```

show bgp replication (for Backup)

```
user@host> show bgp replication
Synchronization backup:
  State: Established 13 ago
  , Unsync timer: 2

  Unsync entry queue:
    Instance: 0 Neighbor: 30.30.30.1 elapsed: 7
    Instance: 0 Neighbor: 40.40.40.3 elapsed: 7
    Instance: 0 Neighbor: 40.40.40.4 elapsed: 7
    Instance: 0 Neighbor: 40.40.40.5 elapsed: 7
    Instance: 0 Neighbor: 40.40.40.6 elapsed: 7
    Instance: 0 Neighbor: 40.40.40.1 elapsed: 7
    Instance: 0 Neighbor: 40.40.40.2 elapsed: 7
```

show bgp summary

| | |
|---|---|
| Syntax | <pre>show bgp summary <exact-instance <i>instance-name</i>> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)></pre> |
| Syntax (EX Series Switch and QFX Series) | <pre>show bgp summary <exact-instance <i>instance-name</i>> <instance <i>instance-name</i>></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>exact-instance option introduced in Junos OS Release 11.4.</p> |
| Description | Display BGP summary information. |
| Options | <p>none—Display BGP summary information for all routing instances.</p> <p>exact-instance <i>instance-name</i>—(Optional) Display information for the specified instance only.</p> <p>instance <i>instance-name</i>—(Optional) Display information for all routing instances whose name begins with this string (for example, cust1, cust11, and cust111 are all displayed when you run the show bgp summary instance cust1 command). The instance name can be master for the main instance, or any valid configured instance name or its prefix.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | <p>show bgp summary (When a Peer Is Not Established) on page 72</p> <p>show bgp summary (When a Peer Is Established) on page 72</p> <p>show bgp summary (CLNS) on page 72</p> <p>show bgp summary (Layer 2 VPN) on page 72</p> <p>show bgp summary (Layer 3 VPN) on page 73</p> |
| Output Fields | <p>Table 15 on page 69 describes the output fields for the show bgp summary command. Output fields are listed in the approximate order in which they appear.</p> |

Table 15: show bgp summary Output Fields

| Field Name | Field Description |
|------------|-----------------------|
| Groups | Number of BGP groups. |
| Peers | Number of BGP peers. |

Table 15: show bgp summary Output Fields (*continued*)

| Field Name | Field Description |
|---------------------|--|
| Down peers | Number of down BGP peers. |
| Table | Name of routing table. |
| Tot Paths | Total number of paths. |
| Act Paths | Number of active routes. |
| Suppressed | Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols. |
| History | Number of withdrawn routes stored locally to keep track of damping history. |
| Damp State | Number of routes with a figure of merit greater than zero, but still active because the value has not reached the threshold at which suppression occurs. |
| Pending | Routes in process by BGP import policy. |
| Peer | Address of each BGP peer. Each peer has one line of output. |
| AS | Peer's AS number. |
| InPkt | Number of packets received from the peer. |
| OutPkt | Number of packets sent to the peer. |
| OutQ | Number of BGP packets that are queued to be transmitted to a particular neighbor. It normally is 0 because the queue usually is emptied quickly. |
| Flaps | Number of times the BGP session has gone down and then come back up. |
| Last Up/Down | Last time since the neighbor transitioned to or from the established state. |

Table 15: show bgp summary Output Fields (*continued*)

| Field Name | Field Description |
|---|--|
| State #Active /Received/Accepted /Damped | <p>Multipurpose field that displays information about BGP peer sessions. The field's contents depend upon whether a session is established and whether it was established on the main routing device or in a routing instance.</p> <ul style="list-style-type: none"> If a peer is not established, the field shows the state of the peer session: Active, Connect, or Idle. In general, the Idle state is the first stage of a connection. BGP is waiting for a Start event. A session can be idle for other reasons as well. The reason that a session is idle is sometimes displayed. For example: Idle (Removal in progress) or Idle (LicenseFailure). If a BGP session is established on the main routing device, the field shows the number of active, received, accepted, and damped routes that are received from a neighbor and appear in the inet.0 (main) and inet.2 (multicast) routing tables. For example, 8/10/10/2 and 2/4/4/0 indicate the following: <ul style="list-style-type: none"> 8 active routes, 10 received routes, 10 accepted routes, and 2 damped routes from a BGP peer appear in the inet.0 routing table. 2 active routes, 4 received routes, 4 accepted routes, and no damped routes from a BGP peer appear in the inet.2 routing table. If a BGP session is established in a routing instance, the field indicates the established (Establ) state, identifies the specific routing table that receives BGP updates, and shows the number of active, received, and damped routes that are received from a neighbor. For example, Establ VPN-AB.inet.0: 2/4/0 indicates the following: <ul style="list-style-type: none"> The BGP session is established. Routes are received in the VPN-AB.inet.0 routing table. The local routing device has two active routes, four received routes, and no damped routes from a BGP peer. <p>When a BGP session is established, the peers are exchanging update messages.</p> |

Sample Output

show bgp summary
(When a Peer Is Not Established)

```

user@host> show bgp summary
Groups: 2 Peers: 4 Down peers: 1
Table          Tot Paths  Act Paths Suppressed  History  Damp State   Pending
inet.0          6          4          0          0          0          0
Peer          AS      InPkt    OutPkt    OutQ    Flaps Last Up/Dwn
State|#Active/Received/Damped...
10.0.0.3      65002      86       90       0        2      42:54 0/0/0

0/0/0
10.0.0.4      65002      90       91       0        1      42:54 0/2/0

0/0/0
10.0.0.6      65002      87       90       0        3          3 Active
10.1.12.1     65001      89       89       0        1      42:54 4/4/0

0/0/0

```

show bgp summary
(When a Peer Is Established)

```

user@host> show bgp summary
Groups: 1 Peers: 3 Down peers: 0
Table          Tot Paths  Act Paths Suppressed  History  Damp State   Pending
inet.0          6          4          0          0          0          0
Peer          AS      InPkt    OutPkt    OutQ    Flaps Last Up/Dwn
State|#Active/Received/Damped...
10.0.0.2      65002    88675    88652     0        2      42:38 2/4/0

0/0/0
10.0.0.3      65002    54528    54532     0        1     2w4d22h 0/0/0

0/0/0
10.0.0.4      65002    51597    51584     0        0     2w3d22h 2/2/0

0/0/0

```

show bgp summary
(CLNS)

```

user@host> show bgp summary
Groups: 1 Peers: 1 Down peers: 0
Peer          AS      InPkt    OutPkt    OutQ    Flaps Last Up/Dwn
State|#Active/Received/Damped...
10.245.245.1   200     1735     1737     0        0     14:26:12 Establ
  bgp.isovpn.0: 3/3/0
  aaaa.iso.0: 3/3/0

```

show bgp summary
(Layer 2 VPN)

```

user@host> show bgp summary
Groups: 1 Peers: 5 Down peers: 0
Table          Tot Paths  Act Paths Suppressed  History  Damp State   Pending
bgp.l2vpn.0      1          1          0          0          0          0
inet.0           0          0          0          0          0          0
Peer          AS      InPkt    OutPkt    OutQ    Flaps Last Up/Dwn
State|#Active/Received/Damped...
10.255.245.35   65299      72       74       0        1      19:00 Establ
  bgp.l2vpn.0: 1/1/0
  frame-vpn.l2vpn.0: 1/1/0
10.255.245.36   65299    2164     2423     0        4      19:50 Establ
  bgp.l2vpn.0: 0/0/0
  frame-vpn.l2vpn.0: 0/0/0
10.255.245.37   65299      36       37       0        4      17:07 Establ

```

```

inet.0: 0/0/0
10.255.245.39 65299 138 168 0 6 53:48 Estab1
bgp.12vpn.0: 0/0/0
frame-vpn.12vpn.0: 0/0/0
10.255.245.69 65299 134 140 0 6 53:42 Estab1
inet.0: 0/0/0

```

show bgp summary (Layer 3 VPN)

```

user@host> show bgp summary
Groups: 2 Peers: 2 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
bgp.13vpn.0 2 2 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Damped...
10.39.1.5 2 21 22 0 0 6:26 Estab1
VPN-AB.inet.0: 1/1/0
10.255.71.15 1 19 21 0 0 6:17 Estab1
bgp.13vpn.0: 2/2/0
VPN-A.inet.0: 1/1/0
VPN-AB.inet.0: 2/2/0
VPN-B.inet.0: 1/1/0

```

show policy damping

| | |
|---|--|
| Syntax | show policy damping <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switch and QFX Series) | show policy damping |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | Display information about BGP route flap damping parameters. |
| Options | <p>none—Display information about BGP route flap damping parameters.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Additional Information | In the output from this command, figure-of-merit values correlate with the probability of future instability of a routing device. Routes with higher figure-of-merit values are suppressed for longer periods of time. The figure-of-merit value decays exponentially over time. A figure-of-merit value of zero is assigned to each new route. The value is increased each time the route is withdrawn or readvertised, or when one of its path attributes changes. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • “Configuring BGP Flap Damping Parameters” in the Routing Policy Configuration Guide • clear bgp damping on page 38 • show route damping on page 516 |
| List of Sample Output | show policy damping on page 75 |
| Output Fields | Table 16 on page 74 describes the output fields for the show policy damping command. Output fields are listed in the approximate order in which they appear. |

Table 16: show policy damping Output Fields

| Field Name | Field Description |
|-----------------|--|
| Halflife | Decay half-life, in minutes. The value represents the period during which the accumulated figure-of-merit value is reduced by half if the route remains stable. If a route has flapped, but then becomes stable, the figure-of-merit value for the route decays exponentially. For example, for a route with a figure-of-merit value of 1500, if no incidents occur, its figure-of-merit value is reduced to 750 after 15 minutes and to 375 after another 15 minutes. |

Table 16: show policy damping Output Fields (*continued*)

| Field Name | Field Description |
|------------------------------|--|
| Reuse merit | Figure-of-merit value below which a suppressed route can be used again. A suppressed route becomes reusable when its figure-of-merit value decays to a value below a reuse threshold, and the route once again is considered usable and can be installed in the forwarding table and exported from the routing table. |
| Suppress/cutoff merit | Figure-of-merit value above which a route is suppressed for use or inclusion in advertisements. When a route's figure-of-merit value reaches a particular level, called the cutoff or suppression threshold, the route is suppressed. When a route is suppressed, the routing table no longer installs the route into the forwarding table and no longer exports this route to any of the routing protocols. |
| Maximum suppress time | Maximum hold-down time, in minutes. The value represents the maximum time that a route can be suppressed no matter how unstable it has been before this period of stability. |
| Computed values | <ul style="list-style-type: none"> • Merit ceiling—Maximum merit that a flapping route can collect. • Maximum decay—Maximum decay half-life, in minutes. |

Sample Output

```

show policy damping      user@host> show policy damping
                          Default damping information:
                          Halflife: 15 minutes
                          Reuse merit: 750 Suppress/cutoff merit: 3000
                          Maximum suppress time: 60 minutes
                          Computed values:
                          Merit ceiling: 12110
                          Maximum decay: 6193
                          Damping information for "standard-damping":
                          Halflife: 10 minutes
                          Reuse merit: 4000 Suppress/cutoff merit: 8000
                          Maximum suppress time: 30 minutes
                          Computed values:
                          Merit ceiling: 32120
                          Maximum decay: 12453

```


ES-IS Operational Mode Commands

Table 17 on page 77 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot the End System-to-Intermediate System (ES-IS) protocol. Commands are listed in alphabetical order.

Table 17: ES-IS Operational Mode Commands

| Task | Command |
|--|------------------------------------|
| Clear ES-IS adjacencies. | <code>clear esis adjacency</code> |
| Clear ES-IS statistics for packets sent or received. | <code>clear esis statistics</code> |
| Display ES-IS adjacencies. | <code>show esis adjacency</code> |
| Display ES-IS interfaces. | <code>show esis interface</code> |
| Display ES-IS statistics for packets sent or received. | <code>show esis statistics</code> |



NOTE: ES-IS is supported only on J Series routers. For information about how to configure ES-IS, see the *J Series Services Router Basic LAN and WAN Access Configuration Guide* or the *Junos OS Routing Protocols Configuration Guide*.

clear esis adjacency

| | |
|---------------------------------|--|
| Syntax | clear esis adjacency <instance <i>instance-name</i> > <interface <i>interface-name</i> > <neighbor> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Clear End System-to-Intermediate System (ES-IS) adjacencies. |
| Options | none —Clear all ES-IS adjacencies. instance <i>instance-name</i> —(Optional) Clear adjacencies for the specified routing instance only. interface <i>interface-name</i> —(Optional) Clear adjacencies for the specified interface only. neighbor —(Optional) Clear adjacencies for the specified neighbor only. |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• show esis adjacency on page 80 |
| List of Sample Output | clear esis adjacency on page 78 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear esis adjacency user@host> clear esis adjacency

clear esis statistics

| | |
|---------------------------------|--|
| Syntax | clear esis statistics <instance <i>instance-name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Clear End System-to-Intermediate System (ES-IS) packet statistics. |
| Options | none —Clear ES-IS packet statistics for all routing instances. instance <i>instance-name</i> —(Optional) Clear ES-IS packet statistics for the specified routing instance only. |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• show esis statistics on page 84 |
| List of Sample Output | clear esis statistics on page 79 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear esis statistics user@host> clear esis statistics

show esis adjacency

| | |
|---------------------------------|---|
| Syntax | show esis adjacency <brief detail extensive> <esis-neighbor-id> <instance <i>instance-name</i> > <interface <i>interface-name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | (J Series and PTX Series routers only) Display End System-to-Intermediate System (ES-IS) adjacencies. |
| Options | <p>none—(Same as brief) Display all ES-IS adjacencies.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>esis-neighbor-id—(Optional) Display adjacencies for the specified neighbor's network service access point (NSAP) only.</p> <p>instance <i>instance-name</i>—(Optional) Display adjacencies for the specified routing instance only.</p> <p>interface <i>interface-name</i>—(Optional) Display adjacencies for the specified interface only.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> clear esis adjacency on page 78 |
| List of Sample Output | show esis adjacency on page 81 show esis adjacency brief on page 81 show esis adjacency detail on page 81 show esis adjacency extensive on page 81 |
| Output Fields | Table 18 on page 80 describes the output fields for the show esis adjacency command. Output fields are listed in the approximate order in which they appear. |

Table 18: show esis adjacency Output Fields

| Field Name | Field Description | Level of Output |
|--------------------|---|--------------------------------|
| Nbr Type | Type of network service access point (NSAP) of this neighbor. | brief none |
| NSAP/NET | NSAP of this neighbor. | All levels |
| Type | Type of NSAP of this neighbor. | detail extensive |
| Hold (secs) | Holdtime interval advertised by this neighbor. | brief none |
| Interface | Interface through which the neighbor is reachable. | All levels |

Table 18: show esis adjacency Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------|---|------------------|
| Advertised holdtime | Holdtime interval advertised by this neighbor. | detail extensive |
| Expires in | How long until the adjacency expires, in seconds. | detail extensive |
| SNPA | Subnetwork point of attachment (MAC address of the neighbor). | detail extensive |
| Transition log | List of recent transitions. <ul style="list-style-type: none"> • When—Time of advertisement from this neighbor. • State—State of the adjacency: Up, Down, New, One-way, Initializing, or Rejected. • Event—Event causing the state. • Down reason—Reason the adjacency is down. | extensive |

Sample Output

| | |
|--------------------------------------|--|
| show esis adjacency | <pre> user@host> show esis adjacency Nbr NSAP/NET Hold Interface Type (secs) IS 47.0005.80ff.f800.0000.0108.0001.0102.5501.6008 135 fe-0/0/0.0 </pre> |
| show esis adjacency brief | The output for the show esis adjacency brief command is identical to that for the show esis adjacency command. For sample output, see show esis adjacency on page 81 . |
| show esis adjacency detail | <pre> user@host> show esis adjacency detail NSAP/NET: 47.0005.80ff.f800.0000.0108.0001.0102.5501.6008, Type: IS Interface: fe-0/0/0.0, Advertised hold time: 180 secs, Expires in: 173 secs SNPA: 0:5:85:c1:73:71 </pre> |
| show esis adjacency extensive | <pre> user@host> show esis adjacency extensive NSAP/NET: 47.0005.80ff.f800.0000.0108.0001.0102.5501.6008, Type: IS Interface: fe-0/0/0.0, Advertised hold time: 180 secs, Expires in: 167 secs SNPA: 0:5:85:c1:73:71 Transition log: When State Event Down reason Sun Nov 26 22:07:35 Up Received ISH </pre> |

show esis interface

| | |
|---------------------------------|--|
| Syntax | show esis interface <brief detail extensive> <instance <i>instance-name</i> > <interface <i>interface-name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | (J Series and PTX Series routers only) Display End System-to-Intermediate System (ES-IS) interface information. |
| Options | <p>none—(Same as brief) Display information for all configured ES-IS interfaces.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>instance <i>instance-name</i>—(Optional) Display configured interfaces for the specified routing instance only.</p> <p>interface <i>interface-name</i>—(Optional) Display information about the specified interface only.</p> |
| Required Privilege Level | view |
| List of Sample Output | show esis interface on page 83 show esis interface brief on page 83 show esis interface detail on page 83 show esis interface extensive on page 83 |
| Output Fields | Table 19 on page 82 describes the output fields for the show esis interface command. Output fields are listed in the approximate order in which they appear. |

Table 19: show esis interface Output Fields

| Field Name | Field Description | Level of Output |
|-------------------------------|--|------------------|
| Interface | Interface through which the adjacency is made. | All levels |
| Receives | Types of hello messages that are received. | All levels |
| Sends | Types of hello messages that are sent. | All levels |
| Hello interval | Interface's hello interval, in seconds. | All levels |
| Adjacencies or Num Adj | Number of adjacencies established on this interface. | All levels |
| Holdtime | Interface's hold time, in seconds. | detail extensive |
| State | Internal implementation information. | detail extensive |

Table 19: show esis interface Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------------------|---|------------------|
| End system configuration timer | Time, in seconds, for the end system to configure itself for ES-IS. | detail extensive |
| Interface index | Index value. | detail extensive |
| NET used in hello | Network entity title used in hello messages. | detail extensive |

Sample Output

show esis interface

```
user@host> show esis interface
Interface          Receives    Sends    Hello Interval    Num Adj
fe-0/0/0.0         ISH         ISH         60.00             1
lo0.0              ISH         -          60.00             0
```

show esis interface brief

The output for the **show esis interface brief** command is identical to that for the **show esis interface** command. For sample output, see [show esis interface on page 83](#).

show esis interface detail

```
user@host> show esis interface detail
Interface: fe-0/0/0.0
  Receives: ISH, Sends: ISH, Hello interval: 60.00
  Adjacencies: 1, Holdtime: 180, End system configuration timer: 180
  Interface index: 68, State: 0x2
  NET used in hello: 47.0005.80ff.f800.0000.0108.0001.0102.5501.6007

Interface: lo0.0
  Receives: ISH, Sends: - , Hello interval: 60.00
  Adjacencies: 0, Holdtime: 180, End system configuration timer: 180
  Interface index: 64, State: 0x2
  NET used in hello: 47.0005.80ff.f800.0000.0108.0001.0102.5501.6007
```

show esis interface extensive

The output for the **show esis interface extensive** command is identical to that for the **show esis interface detail** command. For sample output, see [show esis interface detail on page 83](#).

show esis statistics

| | |
|---------------------------------|--|
| Syntax | <code>show esis statistics</code> <code><instance <i>instance-name</i>></code> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | (J Series and PTX Series routers only) Display End System-to-Intermediate System (ES-IS) packet statistics. |
| Options | <p>none—Display ES-IS packet statistics for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Display ES-IS statistics for the specified routing instance only.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> clear esis statistics on page 79 |
| List of Sample Output | show esis statistics on page 85 |
| Output Fields | Table 20 on page 84 describes the output fields for the show esis statistics command. Output fields are listed in the approximate order in which they appear. |

Table 20: show esis statistics Output Fields

| Field Name | Field Description |
|------------------------------------|---|
| PDU type | Protocol data unit type. |
| Received | Number of PDUs received since IS-IS started or since the statistics were set to zero. |
| Processed | Number of PDUs received less the number dropped. |
| Drops | Number of PDUs dropped. |
| Sent | Number of PDUs transmitted since IS-IS started or since the statistics were set to zero. |
| Total packets received/sent | Total number of PDUs received and transmitted since IS-IS started or since the statistics were set to zero. |

Sample Output

`show esis statistics`

```
user@host> show esis statistics
PDU type  Received Processed Drops   Sent
ESH              3         3     0      8
ISH             11        10     1      4
RD              0         0     0      0
Unknown         0         0     0      0
Totals         14        13     1     12
Total packets received: 14 sent: 0
```


CHAPTER 5

IP Multicast Operational Mode Commands

Table 21 on page 87 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot IP multicast. In the table, the commands are listed in alphabetical order.

Table 21: IP Multicast Operational Mode Commands

| Task | Command |
|---|--|
| Clear Automatic Multicast Tunneling (AMT) protocol statistics. | <code>clear amt statistics</code> |
| Clear Automatic Multicast Tunneling (AMT) protocol state. | <code>clear amt tunnel</code> |
| Clear Internet Group Management Protocol (IGMP) group members. | <code>clear igmp membership</code> |
| Clear IGMP snooping membership information. | <code>clear igmp snooping membership</code> |
| Clear IGMP snooping statistics. | <code>clear igmp snooping statistics</code> |
| Clear IGMP statistics. | <code>clear igmp statistics</code> |
| Clear Multicast Listener Discovery (MLD) group members. | <code>clear mld membership</code> |
| Clear MLD statistics. | <code>clear mld statistics</code> |
| Clear Multicast Source Discovery Protocol (MSDP) source active cache. | <code>clear msdp cache</code> |
| Clear MSDP statistics. | <code>clear msdp statistics</code> |
| Clear multicast bandwidth admissions. | <code>clear multicast bandwidth-admission</code> |
| Clear IP multicast forwarding cache entries. | <code>clear multicast forwarding-cache</code> |
| Clear multicast scope. | <code>clear multicast scope</code> |

Table 21: IP Multicast Operational Mode Commands (*continued*)

| Task | Command |
|---|---|
| Clear multicast sessions. | <code>clear multicast sessions</code> |
| Clear multicast snooping statistics. | <code>clear multicast snooping statistics</code> |
| Clear multicast statistics. | <code>clear multicast statistics</code> |
| Clear Pragmatic General Multicast (PGM) negative acknowledgments (NAKs). | <code>clear pgm negative-acknowledgments</code> |
| Clear PGM source-path messages. | <code>clear pgm source-path-messages</code> |
| Clear PGM statistics. | <code>clear pgm statistics</code> |
| Clear the Protocol Independent Multicast (PIM) join and prune states. | <code>clear pim join</code> |
| Redistribute PIM joins among available links. | <code>clear pim join-distribution</code> |
| Clear PIM register message counters. | <code>clear pim register</code> |
| Clear PIM snooping joins. | <code>clear pim snooping join</code> |
| Clear Pim snooping statistics. | <code>clear pim snooping statistics</code> |
| Clear PIM statistics. | <code>clear pim statistics</code> |
| Rebalance multicast tunnel (MT) interfaces. | <code>request pim multicast-tunnel rebalance</code> |
| Display Automatic Multicast Tunneling (AMT) protocol tunnel statistics. | <code>show amt statistics</code> |
| Display summary information about the Automatic Multicast Tunneling (AMT) protocol. | <code>show amt summary</code> |
| Display information about the Automatic Multicast Tunneling (AMT) dynamic tunnels. | <code>show amt tunnel</code> |
| Display the status of interfaces on which Distance Vector Multicast Routing Protocol (DVMRP) is configured. | <code>show dvmrp interfaces</code> |
| Display DVMRP neighbors. | <code>show dvmrp neighbors</code> |
| Display DVMRP prefixes. | <code>show dvmrp prefix</code> |
| Display DVMRP prunes. | <code>show dvmrp prunes</code> |
| Display members of IGMP groups. | <code>show igmp group</code> |

Table 21: IP Multicast Operational Mode Commands (*continued*)

| Task | Command |
|---|---|
| Display members of IGMP groups by interface. | <code>show igmp interface</code> |
| Display IGMP snooping interface information. | <code>show igmp snooping interface</code> |
| Display IGMP snooping membership information. | <code>show igmp snooping membership</code> |
| Display IGMP snooping statistics. | <code>show igmp snooping statistics</code> |
| Display IGMP statistics. | <code>show igmp statistics</code> |
| Display members of MLD groups. | <code>show mld group</code> |
| Display members of MLD groups by interface. | <code>show mld interface</code> |
| Display MLD statistics. | <code>show mld statistics</code> |
| Display MSDP peers. | <code>show msdp</code> |
| Display multicast sources learned from MSDP. | <code>show msdp source</code> |
| Display the MSDP source-active cache. | <code>show msdp source-active</code> |
| Display MSDP statistics. | <code>show msdp statistics</code> |
| Display backup PE router group information when ingress PE redundancy is configured. | <code>show multicast backup-pe-groups</code> |
| Display configuration information about IP multicast flow maps. | <code>show multicast flow-map</code> |
| Display IP multicast forwarding cache statistics. | <code>show multicast forwarding-cache statistics</code> |
| Display multicast interface bandwidth information. | <code>show multicast interface</code> |
| Display multicast network configuration. | <code>show multicast mrinfo</code> |
| Display entries in the multicast next-hop table. | <code>show multicast next-hops</code> |
| Display configuration information about PIM-to-IGMP message translation, also known as PIM-to-IGMP proxy. | <code>show multicast pim-to-igmp-proxy</code> |
| Display configuration information about PIM-to-MLD message translation, also known as PIM-to-MLD proxy. | <code>show multicast pim-to-mld-proxy</code> |
| Display entries in the multicast forwarding cache. | <code>show multicast route</code> |
| Display multicast reverse-path-forwarding calculations. | <code>show multicast rpf</code> |

Table 21: IP Multicast Operational Mode Commands (*continued*)

| Task | Command |
|---|---|
| Display administratively scoped addresses. | <code>show multicast scope</code> |
| Display multicast snooping next-hops | <code>show multicast snooping next-hops</code> |
| Display announced multicast sessions. | <code>show multicast sessions</code> |
| Display multicast snooping route. | <code>show multicast snooping route</code> |
| Display multicast snooping statistics. | <code>show multicast snooping statistics</code> |
| Display multicast statistics. | <code>show multicast statistics</code> |
| Display most active multicast groups. | <code>show multicast usage</code> |
| Display sent or received NAKs. | <code>show pgm negative-acknowledgments</code> |
| Display PGM source-path messages. | <code>show pgm source-path-messages</code> |
| Display PGM statistics. | <code>show pgm statistics</code> |
| Display bootstrap routers. | <code>show pim bootstrap</code> |
| Display the status of interfaces on which PIM is configured. | <code>show pim interfaces</code> |
| Display PIM (*,RP) join and prune states. | <code>show pim join</code> |
| Display PIM data-driven multicast distribution trees (MDTs). | <code>show pim mdt</code> |
| Display the information cached from multicast distribution tree (MDT) join TLV packets received by all PE routers in a PIM-enabled VPN routing and forwarding (VRF)-instance. | <code>show pim mdt data-mdt-joins</code> |
| Display the maximum number configured and the currently active data multicast distribution trees (MDTs) for a specific VPN routing and forwarding (VRF) instance. | <code>show pim mdt data-mdt-limit</code> |
| Display PIM neighbors. | <code>show pim neighbors</code> |
| Display rendezvous points. | <code>show pim rps</code> |
| Display information about PIM snooping interfaces. | <code>show pim snooping interfaces</code> |
| Display PIM snooping joins. | <code>show pim snooping join</code> |
| Display information about PIM snooping neighbors. | <code>show pim snooping neighbors</code> |

Table 21: IP Multicast Operational Mode Commands (*continued*)

| Task | Command |
|--|---|
| Display PIM snooping statistics. | <code>show pim snooping statistics</code> |
| Display PIM source RPF state. | <code>show pim source</code> |
| Display PIM statistics. | <code>show pim statistics</code> |
| Display Session Announcement Protocol (SAP) addresses. | <code>show sap listen</code> |
| Test MSDP peers. | <code>test msdp</code> |



NOTE: For information about the `mtrace` commands used to monitor IP multicast traffic in real time, see the *Junos System Basics and Services Command Reference*. For information about how to configure IP multicast, see the *Junos Multicast Protocols Configuration Guide*.

clear igmp membership

| | |
|---|--|
| Syntax | clear igmp membership <group <i>address-range</i> > <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switch and the QFX Series) | clear igmp membership <group <i>address-range</i> > <interface <i>interface-name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | Clear Internet Group Management Protocol (IGMP) group members. |
| Options | none —Clear all IGMP members on all interfaces and for all address ranges. group <i>address-range</i> —(Optional) Clear all IGMP members that are in a particular address range. An example of a range is 224.2/16 . If you omit the destination prefix length, the default is /32 . interface <i>interface-name</i> —(Optional) Clear all IGMP group members on an interface. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• show igmp group on page 134• show igmp interface on page 138 |
| List of Sample Output | clear igmp membership on page 93 clear igmp membership interface on page 93 clear igmp membership group on page 94 |
| Output Fields | See show igmp group for an explanation of output fields. |

Sample Output

clear igmp membership

The following sample output displays IGMP group information before and after the **clear igmp membership** command is entered:

```
user@host> show igmp group
Interface      Group           Last Reported   Timeout
so-0/0/0       224.2.127.253  10.1.128.1      186
so-0/0/0       224.2.127.254  10.1.128.1      186
so-0/0/0       239.255.255.255 10.1.128.1      187
so-0/0/0       224.1.127.255   10.1.128.1      188
local          224.0.0.6       (null)          0
local          224.0.0.5       (null)          0
local          224.2.127.254   (null)          0
local          239.255.255.255 (null)          0
local          224.0.0.2       (null)          0
local          224.0.0.13      (null)          0
```

```
user@host> clear igmp membership
Clearing Group Membership Info for so-0/0/0
Clearing Group Membership Info for so-1/0/0
Clearing Group Membership Info for so-2/0/0
```

```
user@host> show igmp group
Interface      Group           Last Reported   Timeout
local          224.0.0.6       (null)          0
local          224.0.0.5       (null)          0
local          224.2.127.254   (null)          0
local          239.255.255.255 (null)          0
local          224.0.0.2       (null)          0
local          224.0.0.13      (null)          0
```

clear igmp membership interface

The following sample output displays IGMP group information before and after the **clear igmp membership interface** command is issued:

```
user@host> show igmp group
Interface      Group           Last Reported   Timeout
so-0/0/0       224.2.127.253  10.1.128.1      210
so-0/0/0       239.255.255.255 10.1.128.1      210
so-0/0/0       224.1.127.255   10.1.128.1      215
so-0/0/0       224.2.127.254   10.1.128.1      216
local          224.0.0.6       (null)          0
local          224.0.0.5       (null)          0
local          224.2.127.254   (null)          0
local          239.255.255.255 (null)          0
local          224.0.0.2       (null)          0
local          224.0.0.13      (null)          0
```

```
user@host> clear igmp membership interface so-0/0/0
Clearing Group Membership Info for so-0/0/0
```

```
user@host> show igmp group
Interface      Group           Last Reported   Timeout
local          224.0.0.6       (null)          0
local          224.0.0.5       (null)          0
local          224.2.127.254   (null)          0
```

| | | | |
|-------|-----------------|--------|---|
| local | 239.255.255.255 | (null) | 0 |
| local | 224.0.0.2 | (null) | 0 |
| local | 224.0.0.13 | (null) | 0 |

clear igmp membership group

The following sample output displays IGMP group information before and after the **clear igmp membership group** command is entered:

```
user@host> show igmp group
```

| Interface | Group | Last Reported | Timeout |
|-----------|-----------------|---------------|---------|
| so-0/0/0 | 224.2.127.253 | 10.1.128.1 | 210 |
| so-0/0/0 | 239.255.255.255 | 10.1.128.1 | 210 |
| so-0/0/0 | 224.1.127.255 | 10.1.128.1 | 215 |
| so-0/0/0 | 224.2.127.254 | 10.1.128.1 | 216 |
| local | 224.0.0.6 | (null) | 0 |
| local | 224.0.0.5 | (null) | 0 |
| local | 224.2.127.254 | (null) | 0 |
| local | 239.255.255.255 | (null) | 0 |
| local | 224.0.0.2 | (null) | 0 |
| local | 224.0.0.13 | (null) | 0 |

```
user@host> clear igmp membership group 239.225/16
Clearing Group Membership Range 239.225.0.0/16 on so-0/0/0
Clearing Group Membership Range 239.225.0.0/16 on so-1/0/0
Clearing Group Membership Range 239.225.0.0/16 on so-2/0/0
```

```
user@host> show igmp group
```

| Interface | Group | Last Reported | Timeout |
|-----------|-----------------|---------------|---------|
| so-0/0/0 | 224.1.127.255 | 10.1.128.1 | 231 |
| so-0/0/0 | 224.2.127.254 | 10.1.128.1 | 233 |
| so-0/0/0 | 224.2.127.253 | 10.1.128.1 | 236 |
| local | 224.0.0.6 | (null) | 0 |
| local | 224.0.0.5 | (null) | 0 |
| local | 224.2.127.254 | (null) | 0 |
| local | 239.255.255.255 | (null) | 0 |
| local | 224.0.0.2 | (null) | 0 |
| local | 224.0.0.13 | (null) | 0 |

clear igmp snooping membership

| | |
|---------------------------------|--|
| Syntax | clear igmp snooping membership <group source <i>address</i> > <instance <i>instance-name</i> > <interface <i>interface-name</i> > <learning-domain <i>learning-domain-name</i> > <vlan-id <i>vlan-identifier</i> > |
| Release Information | Command introduced in Junos OS Release 8.5. |
| Description | Clear IP IGMP snooping membership information. |
| Options | <p>none—Clear IGMP snooping membership for all supported address families on all interfaces.</p> <p>group source <i>address</i>—(Optional) Clear IGMP snooping membership for the specified multicast group or source address.</p> <p>instance <i>instance-name</i>—(Optional) Clear IGMP snooping membership for the specified instance.</p> <p>interface <i>interface-name</i>—(Optional) Clear IGMP snooping membership on a specific interface.</p> <p>learning-domain <i>learning-domain-name</i>—(Optional) Perform this operation on all learning domains or on a particular learning domain.</p> <p>vlan-id <i>vlan-identifier</i>—(Optional) Perform this operation on a particular VLAN.</p> |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none"> • show igmp snooping membership on page 145 |
| List of Sample Output | clear igmp snooping membership on page 95 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

```
clear igmp snooping membership  user@host> clear igmp snooping membership
```

clear igmp snooping statistics

| | |
|---------------------------------|--|
| Syntax | <code>clear igmp snooping statistics</code> <code><instance <i>instance-name</i>></code> <code><interface <i>interface-name</i>></code> <code><learning-domain (all <i>learning-domain-name</i>)></code> <code><logical-system (all <i>logical-system-name</i>)></code> |
| Release Information | Command introduced in Junos OS Release 8.5. |
| Description | Clear IP IGMP snooping statistics. |
| Options | <p>none—Clear IGMP snooping statistics for all supported address families on all interfaces.</p> <p>instance <i>instance-name</i>—(Optional) Clear IGMP snooping statistics for the specified instance.</p> <p>interface <i>interface-name</i>—(Optional) Clear IGMP snooping statistics on a specific interface.</p> <p>learning-domain (all <i>learning-domain-name</i>)—(Optional) Perform this operation on all learning domains or on a particular learning domain.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• show igmp snooping statistics on page 149 |
| List of Sample Output | clear igmp snooping statistics on page 96 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

| | |
|---|---|
| <code>clear igmp snooping statistics</code> | <code>user@host> clear igmp snooping statistics</code> |
|---|---|

clear igmp statistics

| | |
|------------------------------------|--|
| Syntax | clear igmp statistics <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switches) | clear igmp statistics <interface <i>interface-name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | Clear Internet Group Management Protocol (IGMP) statistics. |
| Options | <p>none—Clear IGMP statistics on all interfaces.</p> <p>interface <i>interface-name</i>—(Optional) Clear IGMP statistics for the specified interface only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none"> • show igmp statistics on page 152 |
| List of Sample Output | clear igmp statistics on page 98 |
| Output Fields | See show igmp statistics for an explanation of output fields. |

Sample Output

clear igmp statistics

The following sample output displays IGMP statistics information before and after the **clear igmp statistics** command is entered:

```

user@host> show igmp statistics
IGMP packet statistics for all interfaces
IGMP Message type      Received      Sent  Rx errors
Membership Query        8883         459      0
V1 Membership Report    0            0        0
DVMRP                   19784        35476    0
PIM V1                  18310         0        0
Cisco Trace             0            0        0
V2 Membership Report    0            0        0
Group Leave             0            0        0
Mtrace Response         0            0        0
Mtrace Request          0            0        0
Domain Wide Report      0            0        0
V3 Membership Report    0            0        0
Other Unknown types     0            0        0
IGMP v3 unsupported type 0            0        0
IGMP v3 source required for SSM 0            0
IGMP v3 mode not applicable for SSM 0            0

IGMP Global Statistics
Bad Length              0
Bad Checksum            0
Bad Receive If          0
Rx non-local            1227

user@host> clear igmp statistics
user@host> show igmp statistics
IGMP packet statistics for all interfaces
IGMP Message type      Received      Sent  Rx errors
Membership Query        0            0      0
V1 Membership Report    0            0      0
DVMRP                   0            0      0
PIM V1                  0            0      0
Cisco Trace             0            0      0
V2 Membership Report    0            0      0
Group Leave             0            0      0
Mtrace Response         0            0      0
Mtrace Request          0            0      0
Domain Wide Report      0            0      0
V3 Membership Report    0            0      0
Other Unknown types     0            0      0
IGMP v3 unsupported type 0            0      0
IGMP v3 source required for SSM 0            0
IGMP v3 mode not applicable for SSM 0            0

IGMP Global Statistics
Bad Length              0
Bad Checksum            0
Bad Receive If          0
Rx non-local            0

```

clear mld membership

| | |
|---------------------------------|---|
| Syntax | clear mld membership <group <i>group-name</i> > <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Clear Multicast Listener Discovery (MLD) group membership. |
| Options | <p>none—Clear all MLD memberships.</p> <p>group <i>group-name</i>—(Optional) Clear MLD membership for the specified group.</p> <p>interface <i>interface-name</i>—(Optional) Clear MLD group membership for the specified interface.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • show mld group on page 155 |
| List of Sample Output | clear mld membership on page 99 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear mld membership user@host> clear mld membership

clear mld statistics

| | |
|---------------------------------|---|
| Syntax | <code>clear mld statistics</code> <code><interface <i>interface-name</i>></code> <code><logical-system (all <i>logical-system-name</i>)></code> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Clear Multicast Listener Discovery (MLD) statistics. |
| Options | none —(Same as logical-system all) Clear MLD statistics for all interfaces. interface <i>interface-name</i> —(Optional) Clear MLD statistics for the specified interface. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• show mld statistics on page 164 |
| List of Sample Output | clear mld statistics on page 100 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

`clear mld statistics` `user@host> clear mld statistics`

clear msdp cache

| | |
|---------------------------------|---|
| Syntax | clear msdp cache <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> <peer <i>peer-address</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 12.1 for the QFX Series. |
| Description | Clear the entries in the Multicast Source Discovery Protocol (MSDP) source-active cache. |
| Options | <p>none—Clear entries in the MSDP source-active cache for all instances, logical systems, and peers.</p> <p>instance <i>instance-name</i>—(Optional) Clear entries for a specific MSDP instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>peer <i>peer-address</i>—(Optional) Clear the MSDP source-active cache entries learned from a specific peer.</p> |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none"> • show msdp source-active on page 171 |
| List of Sample Output | clear msdp cache on page 101 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear msdp cache user@host> clear msdp cache

clear msdp statistics

| | |
|---------------------------------|---|
| Syntax | <code>clear msdp statistics</code> <code><instance <i>instance-name</i>></code> <code><logical-system (all <i>logical-system-name</i>)></code> <code><peer <i>peer-address</i>></code> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 12.1 for the QFX Series. |
| Description | Clear Multicast Source Discovery Protocol (MSDP) peer statistics. |
| Options | none —Clear MSDP statistics for all peers. instance <i>instance-name</i> —(Optional) Clear statistics for the specified instance. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. peer <i>peer-address</i> —(Optional) Clear the statistics for the specified peer. |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• show msdp statistics on page 174 |
| List of Sample Output | clear msdp statistics on page 102 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

`clear msdp statistics` user@host> clear msdp statistics

clear multicast bandwidth-admission

| | |
|---------------------------------|---|
| Syntax | <pre>clear multicast bandwidth-admission <group <i>group-address</i>> <inet inet6> <instance <i>instance-name</i>> <interface <i>interface-name</i>> <source <i>source-address</i>></pre> |
| Release Information | <p>Command introduced in Junos OS Release 8.3.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> |
| Description | Reapply IP multicast bandwidth admissions. |
| Options | <p>none—Reapply multicast bandwidth admissions for all IPv4 forwarding entries in the master routing instance.</p> <p>group <i>group-address</i>—(Optional) Reapply multicast bandwidth admissions for the specified group.</p> <p>inet—(Optional) Reapply multicast bandwidth admission settings for IPv4 flows.</p> <p>inet6—(Optional) Reapply multicast bandwidth admission settings for IPv6 flows.</p> <p>instance <i>instance-name</i>—(Optional) Reapply multicast bandwidth admission settings for the specified instance. If you do not specify an instance, the command applies to the master routing instance.</p> <p>interface <i>interface-name</i>—(Optional) Examines the corresponding outbound interface in the relevant entries and acts as follows:</p> <ul style="list-style-type: none"> • If the interface is congested, and it was admitted previously, it is removed. • If the interface was rejected previously, the clear multicast bandwidth-admission command enables the interface to be admitted as long as enough bandwidth exists on the interface. • If you do not specify an interface, issuing the clear multicast bandwidth-admission command readmits any previously rejected interface for the relevant entries as long as enough bandwidth exists on the interface. <p>To manually reject previously admitted outbound interfaces, you must specify the interface.</p> <p>source <i>source-address</i>—(Optional) Use with the group option to reapply multicast bandwidth admission settings for the specified (source, group) entry.</p> |
| Required Privilege Level | clear |

- Related Documentation**
- [show multicast interface on page 184](#)
- List of Sample Output** [clear multicast bandwidth-admission on page 104](#)
- Output Fields** When you enter this command, you are provided feedback on the status of your request.

Sample Output

`clear multicast
bandwidth-admission` `user@host> clear multicast bandwidth-admission`

clear multicast forwarding-cache

| | |
|---------------------------------|--|
| Syntax | clear multicast forwarding-cache <inet inet6> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced in Junos OS Release 12.2. |
| Description | Clear IP multicast forwarding cache entries. |
| Options | <p>none—(Same as logical-system all) Clear multicast forwarding cache entries.</p> <p>inet—(Optional) Clear multicast forwarding cache entries for IPv4 family addresses.</p> <p>inet6—(Optional) Clear multicast forwarding cache entries for IPv6 family addresses.</p> <p>interface <i>interface-name</i>—(Optional) Clear multicast forwarding cache entries on a specific interface.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none"> • show multicast forwarding-cache statistics on page 182 |
| List of Sample Output | clear multicast forwarding-cache on page 105 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear multicast forwarding-cache

```
user@host> clear multicast forwarding-cache
```

clear multicast scope

| | |
|---|---|
| Syntax | clear multicast scope <inet inet6> <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switch and the QFX Series) | clear multicast scope <inet inet6> <interface <i>interface-name</i> > |
| Release Information | Command introduced in Junos OS Release 7.6. Command introduced in Junos OS Release 9.0 for EX Series switches. inet6 option introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | Clear IP multicast scope statistics. |
| Options | none —(Same as logical-system all) Clear multicast scope statistics. inet —(Optional) Clear multicast scope statistics for IPv4 family addresses. inet6 —(Optional) Clear multicast scope statistics for IPv6 family addresses. interface <i>interface-name</i> —(Optional) Clear multicast scope statistics on a specific interface. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• show multicast scope on page 205 |
| List of Sample Output | clear multicast scope on page 106 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear multicast scope user@host> clear multicast scope

clear multicast sessions

| | |
|---|--|
| Syntax | clear multicast sessions <logical-system (all <i>logical-system-name</i>)> < <i>regular-expression</i> > |
| Syntax (EX Series Switch and the QFX Series) | clear multicast sessions < <i>regular-expression</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | Clear IP multicast sessions. |
| Options | <p>none—(Same as logical-system all) Clear multicast sessions.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><i>regular-expression</i>—(Optional) Clear only multicast sessions that contain the specified regular expression.</p> |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none"> • show multicast sessions on page 207 |
| List of Sample Output | clear multicast sessions on page 107 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear multicast sessions user@host> clear multicast sessions

clear multicast snooping statistics

| | |
|---------------------------------|---|
| Syntax | clear multicast snooping statistics <instance <i>instance-name</i> > <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced in Junos OS Release 8.5. |
| Description | Clear IP multicast snooping statistics. |
| Options | <p>none—Clear multicast snooping statistics for all supported address families on all interfaces.</p> <p>instance <i>instance-name</i>—(Optional) Clear multicast snooping statistics for the specified instance.</p> <p>interface <i>interface-name</i>—(Optional) Clear multicast snooping statistics on a specific interface.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• show multicast snooping statistics on page 215 |
| List of Sample Output | clear multicast snooping statistics on page 108 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

| | |
|--|--|
| <code>clear multicast snooping statistics</code> | <code>user@host> clear multicast snooping statistics</code> |
|--|--|

clear multicast statistics

| | |
|---|---|
| Syntax | clear multicast statistics <inet inet6> <instance <i>instance-name</i> > <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switch and the QFX Series) | clear multicast statistics <inet inet6> <instance <i>instance-name</i> > <interface <i>interface-name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | Clear IP multicast statistics. |
| Options | <p>none—Clear multicast statistics for all supported address families on all interfaces.</p> <p>inet—(Optional) Clear multicast statistics for IPv4 family addresses.</p> <p>inet6—(Optional) Clear multicast statistics for IPv6 family addresses.</p> <p>instance <i>instance-name</i>—(Optional) Clear multicast statistics for the specified instance.</p> <p>interface <i>interface-name</i>—(Optional) Clear multicast statistics on a specific interface.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none"> • show multicast statistics on page 218 |
| List of Sample Output | clear multicast statistics on page 109 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

```
clear multicast statistics      user@host> clear multicast statistics
```

clear pgm negative-acknowledgments

| | |
|---------------------------------|---|
| Syntax | clear pgm negative-acknowledgments |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Clear the Pragmatic General Multicast (PGM) negative acknowledgment (NAK) state received. |
| Options | This command has no options. |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• show pgm negative-acknowledgments on page 224 |
| List of Sample Output | clear pgm negative-acknowledgments on page 110 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

| | |
|---|---|
| <code>clear pgm negative-acknowledgments</code> | <code>user@host> clear pgm negative-acknowledgments</code> |
|---|---|

clear pgm source-path-messages

| | |
|---------------------------------|---|
| Syntax | clear pgm source-path-messages |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Clear Pragmatic General Multicast (PGM) source-path messages. |
| Options | This command has no options. |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• show pgm source-path-messages on page 226 |
| List of Sample Output | clear pgm source-path-messages on page 111 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

| | |
|-----------------------------------|---|
| clear pgm source-path-messages | user@host> clear pgm source-path-messages |
|-----------------------------------|---|

clear pgm statistics

| | |
|---------------------------------|---|
| Syntax | clear pgm statistics |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Clear Pragmatic General Multicast (PGM) statistics. |
| Options | This command has no options. |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• show pgm statistics on page 227 |
| List of Sample Output | clear pgm statistics on page 112 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

`clear pgm statistics` user@host> clear pgm statistics

clear pim join

| | |
|---|--|
| Syntax | clear pim join <group-address> <inet inet6> <instance instance-name> <logical-system (all logical-system-name)> |
| Syntax (EX Series Switch and the QFX Series) | clear pim join <group-address> <inet inet6> <instance instance-name> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | Clear the Protocol Independent Multicast (PIM) join and prune states. |
| Options | <p>none—Clear the PIM join and prune states for all groups, family addresses, and instances.</p> <p>group-address—(Optional) Clear the PIM join and prune states for a group address.</p> <p>inet inet6—(Optional) Clear the PIM join and prune states for IPv4 or IPv6 family addresses, respectively.</p> <p>instance instance-name—(Optional) Clear the join and prune states for a specific PIM-enabled routing instance.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Additional Information | The clear pim join command cannot be used to clear the PIM join and prune state on a backup Routing Engine when nonstop active routing is enabled. |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none"> • show pim join on page 235 |
| List of Sample Output | clear pim join on page 113 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear pim join user@host> clear pim join

clear pim join-distribution

| | |
|---------------------------------|--|
| Syntax | <code>clear pim join-distribution</code> <code><instance <i>instance-name</i>></code> <code><logical-system (all <i>logical-system-name</i>)></code> |
| Release Information | Command introduced in Junos OS Release 10.0. |
| Description | <p>Redistribute the Protocol Independent Multicast (PIM) join states.</p> <p>You can find out if there are multiple paths available for a source (for example, an RP) with the output of the show pim source command.</p> <p>When you include the join-load-balance statement in the configuration, the PIM join states are distributed evenly on available equal-cost multipath links. When an upstream neighbor link fails, Junos OS redistributes the PIM join states to the remaining links. However, when new links are added or the failed link is restored, the existing PIM joins are not redistributed to the new link. New flows will be distributed to the new links. However, in a network without new joins and prunes, the new link is not used for multicast traffic. The clear pim join-distribution command redistributes the existing flows to the new upstream neighbors. Redistributing the existing flows causes traffic to be disrupted, so we recommend that you run the clear pim join-distribution command during a maintenance window.</p> |
| Options | <p>none—Redistribute the PIM join states for the default master instance.</p> <p>instance <i>instance-name</i>—(Optional) Redistribute the join states for a specific PIM-enabled routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Additional Information | The clear pim join-distribution command cannot be used to redistribute the PIM join states on a backup Routing Engine when nonstop active routing is enabled. |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• show pim neighbors on page 253• show pim join on page 235• join-load-balance in the <i>Multicast Protocols Configuration Guide</i> |
| List of Sample Output | clear pim join-distribution on page 115 |
| Output Fields | When you enter this command, you are provided no feedback on the status of your request. You can enter the show pim join command before and after distributing the join state to verify the operation. |

Sample Output

clear pim
join-distribution

user@host> clear pim join-distribution

clear pim register

| | |
|---|---|
| Syntax | clear pim register <inet inet6> <instance <i>instance-name</i> > <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switch and the QFX Series) | clear pim register <inet inet6> <instance <i>instance-name</i> > <interface <i>interface-name</i> > |
| Syntax (PTX Series) | clear pim register <inet inet6> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced in Junos OS Release 7.6. Command introduced in Junos OS Release 9.0 for EX Series switches. inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | Clear Protocol Independent Multicast (PIM) register message counters. |
| Options | <p>none—Clear PIM register message counters for all family addresses, instances, and interfaces.</p> <p>inet inet6—(Optional) Clear PIM register message counters for IPv4 or IPv6 family addresses, respectively.</p> <p>instance <i>instance-name</i>—(Optional) Clear register message counters for a specific PIM-enabled routing instance.</p> <p>interface <i>interface-name</i>—(Optional) Clear PIM register message counters for a specific interface.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Additional Information | The clear pim register command cannot be used to clear the PIM register state on a backup Routing Engine when nonstop active routing is enabled. |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none"> • show pim statistics on page 285 |
| List of Sample Output | clear pim register on page 117 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

`clear pim register` `user@host> clear pim register`

clear pim snooping join

| | |
|---------------------------------|---|
| Syntax | clear pim snooping join <instance <i>instance-name</i> > <vlan-id <i>vlan-id</i> > |
| Release Information | Command introduced in Junos OS Release 12.3. |
| Description | Clear information about Protocol Independent Multicast (PIM) snooping joins. |
| Options | none —Display detailed information. instance <i>instance-name</i> —(Optional) Clear PIM snooping join information for the specified routing instance. vlan-id <i>vlan-identifier</i> —(Optional) Clear PIM snooping join information for the specified VLAN. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">• PIM Snooping for VPLS Use Cases• PIM Snooping for VPLS |
| List of Sample Output | clear pim snooping join on page 119 |
| Output Fields | See show pim snooping join for an explanation of the output fields. |

Sample Output

clear pim snooping join The following sample output displays information about PIM snooping joins before and after the **clear pim snooping join** command is entered:

```
user@host> show pim snooping join extensive
Instance: vpls1
Learning-Domain: vlan-id 10
Learning-Domain: vlan-id 20

Group: 225.1.1.2
Source: *
Flags: sparse,rptree,wildcard
Upstream state: None
Upstream neighbor: 20.0.120.5, port: ge-1/3/7.20
Downstream port: ge-1/3/1.20
Downstream neighbors:
20.0.120.2 State: Join Flags: SRW Timeout: 185

Group: 225.1.1.3
Source: *
Flags: sparse,rptree,wildcard
Upstream state: None
Upstream neighbor: 20.0.120.4, port: ge-1/3/5.20
Downstream port: ge-1/3/3.20
Downstream neighbors:
20.0.120.3 State: Join Flags: SRW Timeout: 175

user@host> clear pim snooping join
Clearing the Join/Prune state for 224.0.0.0/4
Clearing the Join/Prune state for 224.0.0.0/4

user@host> show pim snooping join extensive
Instance: vpls1
Learning-Domain: vlan-id 10
Learning-Domain: vlan-id 20
```

clear pim snooping statistics

| | |
|---------------------------------|--|
| Syntax | clear pim snooping statistics <instance <i>instance-name</i> > <interface <i>interface-name</i> > <vlan-id <i>vlan-id</i> > |
| Release Information | Command introduced in Junos OS Release 12.3. |
| Description | Clear Protocol Independent Multicast (PIM) snooping statistics. |
| Options | <p>none—Clear PIM snooping statistics for all family addresses, instances, and interfaces.</p> <p>instance <i>instance-name</i>—(Optional) Clear statistics for a specific PIM-snooping-enabled routing instance.</p> <p>interface <i>interface-name</i>—(Optional) Clear PIM snooping statistics for a specific interface.</p> <p>vlan-id <i>vlan-identifier</i>—(Optional) Clear PIM snooping statistics information for the specified VLAN.</p> |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• PIM Snooping for VPLS Use Cases• PIM Snooping for VPLS |
| List of Sample Output | clear pim snooping statistics on page 121 |
| Output Fields | See show pim snooping statistics for an explanation of the output fields. |

Sample Output

clear pim snooping statistics

The following sample output displays PIM snooping statistics before and after the **clear pim snooping statistics** command is entered:

```
user@host> show pim snooping statistics
Instance: vpls1
Learning-Domain: vlan-id 10
```

```
Tx J/P messages 0
RX J/P messages 660
Rx J/P messages -- seen 0
Rx J/P messages -- received 660
Rx Hello messages 1396
Rx Version Unknown 0
Rx Neighbor Unknown 0
Rx Upstream Neighbor Unknown 0
Rx Bad Length 0
Rx J/P Busy Drop 0
Rx J/P Group Aggregate 0
Rx Malformed Packet 0
```

```
Learning-Domain: vlan-id 20
```

```
user@host> clear pim snooping statistics
user@host> show pim snooping statistics
Instance: vpls1
Learning-Domain: vlan-id 10
```

```
Tx J/P messages 0
RX J/P messages 0
Rx J/P messages -- seen 0
Rx J/P messages -- received 0
Rx Hello messages 0
Rx Version Unknown 0
Rx Neighbor Unknown 0
Rx Upstream Neighbor Unknown 0
Rx Bad Length 0
Rx J/P Busy Drop 0
Rx J/P Group Aggregate 0
Rx Malformed Packet 0
```

```
Learning-Domain: vlan-id 20
```

clear pim statistics

| | |
|---|---|
| Syntax | clear pim statistics <inet inet6> <instance <i>instance-name</i> > <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switch and the QFX Series) | clear pim statistics <inet inet6> <instance <i>instance-name</i> > <interface <i>interface-name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | Clear Protocol Independent Multicast (PIM) statistics. |
| Options | none —Clear PIM statistics for all family addresses, instances, and interfaces. inet inet6 —(Optional) Clear PIM statistics for IPv4 or IPv6 family addresses, respectively. instance <i>instance-name</i> —(Optional) Clear statistics for a specific PIM-enabled routing instance. interface <i>interface-name</i> —(Optional) Clear PIM statistics for a specific interface. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Additional Information | The clear pim statistics command cannot be used to clear the PIM statistics on a backup Routing Engine when nonstop active routing is enabled. |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• show pim statistics on page 285 |
| List of Sample Output | clear pim statistics on page 123 |
| Output Fields | See show pim statistics for an explanation of output fields. |

Sample Output

clear pim statistics

The following sample output displays PIM statistics before and after the **clear pim statistics** command is entered:

```
user@host> show pim statistics
PIM statistics on all interfaces:
PIM Message type      Received      Sent  Rx errors
Hello                  0             0       0
Register               0             0       0
Register Stop         0             0       0
Join Prune             0             0       0
Bootstrap              0             0       0
Assert                0             0       0
Graft                  0             0       0
Graft Ack              0             0       0
Candidate RP           0             0       0
V1 Query               2111          4222       0
V1 Register            0             0       0
V1 Register Stop       0             0       0
V1 Join Prune          14200         13115       0
V1 RP Reachability     0             0       0
V1 Assert              0             0       0
V1 Graft               0             0       0
V1 Graft Ack           0             0       0
PIM statistics summary for all interfaces:
Unknown type           0
V1 Unknown type        0
Unknown Version        0
Neighbor unknown       0
Bad Length             0
Bad Checksum           0
Bad Receive If         0
Rx Intf disabled       2007
Rx V1 Require V2        0
Rx Register not RP      0
RP Filtered Source      0
Unknown Reg Stop        0
Rx Join/Prune no state  1040
Rx Graft/Graft Ack no state 0
...

user@host> clear pim statistics
user@host> show pim statistics
PIM statistics on all interfaces:
PIM Message type      Received      Sent  Rx errors
Hello                  0             0       0
Register               0             0       0
Register Stop         0             0       0
Join Prune             0             0       0
Bootstrap              0             0       0
Assert                0             0       0
Graft                  0             0       0
Graft Ack              0             0       0
Candidate RP           0             0       0
V1 Query               1             0       0
V1 Register            0             0       0
...
```


request pim multicast-tunnel rebalance

| | |
|------------------------------------|--|
| Syntax | request pim multicast-tunnel rebalance <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switches) | request pim multicast-tunnel rebalance <instance <i>instance-name</i> > |
| Release Information | Command introduced in Junos OS Release 10.2. Command introduced in Junos OS Release 10.2 for EX Series switches. |
| Description | Rebalance the assignment of multicast tunnel encapsulation interfaces across available tunnel-capable PICs or across a configured list of tunnel-capable PICs. You can determine whether a rebalance is necessary by running the show pim interfaces instance <i>instance-name</i> command. |
| Options | <p>none—Re-create and rebalance all tunnel interfaces for all routing instances.</p> <p>instance <i>instance-name</i>—Re-create and rebalance all tunnel interfaces for a specific instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | maintenance |
| Related Documentation | <ul style="list-style-type: none"> • show pim interfaces on page 232 • Load Balancing Multicast Tunnel Interfaces Among Available PICs in the <i>Junos Multicast Protocols Configuration Guide</i> |
| Output Fields | This command produces no output. To verify the operation of the command, run the show pim interface instance <i>instance-name</i> before and after running the request pim multicast-tunnel rebalance command. |

show dvmrp interfaces

| | |
|---------------------------------|--|
| Syntax | show dvmrp interfaces <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display information about Distance Vector Multicast Routing Protocol (DVMRP)–enabled interfaces. |
| Options | <p>none—(Same as logical-system all) Display information about DVMRP-enabled interfaces.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show dvmrp interfaces on page 127 |
| Output Fields | Table 22 on page 126 describes the output fields for the show dvmrp interfaces command. Output fields are listed in the approximate order in which they appear. |

Table 22: show dvmrp interfaces Output Fields

| Field Name | Field Description |
|------------------|--|
| Interface | Name of the interface. |
| State | State of the interface: up or down . |
| Leaf | Whether the interface is a leaf (that is, whether it has no neighbors) or whether it has neighbors. |
| Metric | Interface metric: a value from 1 through 31. |
| Announce | Number of routes the interface is announcing. |
| Mode | DVMRP mode: <ul style="list-style-type: none"> • Forwarding—DVMRP does both the routing and the multicast data forwarding. • Unicast-routing—DVMRP does only the routing. Forwarding of the multicast data packets can be done by enabling PIM on the interface. |

Sample Output

```
show dvmrp interfaces  user@host> show dvmrp interfaces
Interface State Leaf Metric Announce Mode
fxp0.0    Up   N   1   4 Forwarding
fxp1.0    Up   N   1   4 Forwarding
fxp2.0    Up   N   1   3 Forwarding
lo0.0     Up   Y   1   0 Unicast-routing
```

show dvmrp neighbors

| | |
|---------------------------------|---|
| Syntax | show dvmrp neighbors <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display information about Distance Vector Multicast Routing Protocol (DVMRP) neighbors. |
| Options | <p>none—(Same as logical-system all) Display information about DVMRP neighbors.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show dvmrp neighbors on page 129 |
| Output Fields | Table 23 on page 128 describes the output fields for the show dvmrp neighbors command. Output fields are listed in the approximate order in which they appear. |

Table 23: show dvmrp neighbors Output Fields

| Field Name | Field Description |
|--------------------|---|
| Neighbor | Address of the neighboring DVMRP router. |
| Interface | Interface through which the neighbor is reachable. |
| Version | Version of DVMRP that the neighbor is running, in the format <i>majorminor</i> . |
| Flags | <p>Information about the neighbor:</p> <ul style="list-style-type: none"> 1—One way. The local router has seen the neighbor, but the neighbor has not seen the local router. G—Neighbor supports generation ID. L—Neighbor is a leaf router. M—Neighbor supports mtrace. N—Neighbor supports netmask in prune messages and graft messages. P—Neighbor supports pruning. S—Neighbor supports SNMP. |
| Routes | Number of routes learned from the neighbor. |
| Timeout | How long until the DVMRP neighbor information times out, in seconds. |
| Transitions | Number of generation ID changes that have occurred since the local router learned about the neighbor. |

Sample Output

```
show dvmrp neighbors  user@host> show dvmrp neighbors
Neighbor      Interface      Version  Flags    Routes  Timeout  Transitions
192.168.1.1    ipip.0         3.255    PGM      3       28       1
```

show dvmrp prefix

| | |
|---------------------------------|---|
| Syntax | show dvmrp prefix <brief detail> <logical-system (all <i>logical-system-name</i>)> <prefix> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display information about Distance Vector Multicast Routing Protocol (DVMRP) prefixes. |
| Options | <p>none—Display standard information about all DVMRP prefixes.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>prefix—(Optional) Display information about specific prefixes.</p> |
| Required Privilege Level | view |
| List of Sample Output | show dvmrp prefix on page 131 show dvmrp prefix brief on page 131 show dvmrp prefix detail on page 131 |
| Output Fields | Table 24 on page 130 describes the output fields for the show dvmrp prefix command. Output fields are listed in the approximate order in which they appear. |

Table 24: show dvmrp prefix Output Fields

| Field Name | Field Description | Level of Output |
|------------------------|--|-----------------|
| Prefix | DVMRP route. | All levels |
| Next hop | Next hop from which the route was learned. | All levels |
| Age | Last time that the route was refreshed. | All levels |
| <i>multicast-group</i> | Multicast group address. | detail |
| Prunes sent | Number of prune messages sent to the multicast group. | detail |
| Grafts sent | Number of grafts sent to the multicast group. | detail |
| Cache lifetime | Lifetime of the group in the multicast cache, in seconds. | detail |
| Prune lifetime | Lifetime remaining and total lifetime of prune messages, in seconds. | detail |

Sample Output

show dvmrp prefix

```
user@host> show dvmrp prefix
Prefix          Next hop      Age
10.38.0.0       /30 10.38.0.1 00:06:17
10.38.0.4       /30 10.38.0.5 00:06:13
10.38.0.8       /30 10.38.0.2 00:00:04
10.38.0.12      /30 10.38.0.6 00:00:04
10.255.14.114   /32 10.255.14.114 00:06:17
10.255.14.142   /32 10.38.0.2 00:00:04
10.255.14.144   /32 10.38.0.2 00:00:04
10.255.70.15    /32 10.38.0.6 00:00:04
192.168.14.0    /24 192.168.14.114 00:06:17
192.168.195.40  /30 192.168.195.41 00:06:17
192.168.195.92  /30 10.38.0.2 00:00:04
```

show dvmrp prefix brief

The output for the **show dvmrp prefix brief** command is identical to that for the **show dvmrp prefix** command.

show dvmrp prefix detail

```
user@host> show dvmrp prefix detail
Prefix          Next hop      Age
10.38.0.0       /30 10.38.0.1 00:06:28
10.38.0.4       /30 10.38.0.5 00:06:24
10.38.0.8       /30 10.38.0.2 00:00:15
10.38.0.12      /30 10.38.0.6 00:00:15
10.255.14.114   /32 10.255.14.114 00:06:28
10.255.14.142   /32 10.38.0.2 00:00:15
10.255.14.144   /32 10.38.0.2 00:00:15
10.255.70.15    /32 10.38.0.6 00:00:15
192.168.14.0    /24 192.168.14.114 00:06:28
192.168.195.40  /30 192.168.195.41 00:06:28
192.168.195.92  /30 10.38.0.2 00:00:15
```

show dvmrp prunes

| | |
|---------------------------------|---|
| Syntax | show dvmrp prunes <all rx tx> <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display information about active Distance Vector Multicast Routing Protocol (DVMRP) prune messages. |
| Options | <p>none—Display received and transmitted DVMRP prune information.</p> <p>all—(Optional) Display information about all received and transmitted prune messages.</p> <p>rx—(Optional) Display information about received prune messages.</p> <p>tx—(Optional) Display information about transmitted prune messages.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show dvmrp prunes on page 133 |
| Output Fields | Table 25 on page 132 describes the output fields for the show dvmrp prunes command. Output fields are listed in the approximate order in which they appear. |

Table 25: show dvmrp prunes Output Fields

| Field Name | Field Description |
|----------------------|--|
| Group | Group address. |
| Source prefix | Prefix for the prune. |
| Timeout | How long until the prune message expires, in seconds. |
| Neighbor | Neighbor to which the prune was sent or from which the prune was received. |

Sample Output

`show dvmrp prunes`

```
user@host> show dvmrp prunes
Group          Source prefix      Timeout Neighbor
224.0.1.1      128.112.0.0       /12    7077 192.168.1.1
224.0.1.32     160.0.0.0         /3     7087 192.168.1.1
224.2.123.4    136.0.0.0         /5     6955 192.168.1.1
224.2.127.1    129.0.0.0         /8     7046 192.168.1.1
224.2.135.86   128.102.128.0     /17    7071 192.168.1.1
224.2.135.86   129.0.0.0         /8     7074 192.168.1.1
224.2.135.86   130.0.0.0         /7     7071 192.168.1.1
...
```

show igmp group

| | |
|---|--|
| Syntax | show igmp group <brief detail> <group-name> <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switch and the QFX Series) | show igmp group <brief detail> <group-name> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | Display Internet Group Management Protocol (IGMP) group membership information. |
| Options | <p>none—Display standard information about membership for all IGMP groups.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>group-name—(Optional) Display group membership for the specified IP address only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> clear igmp membership on page 92 |
| List of Sample Output | show igmp group (Include Mode) on page 136 show igmp group (Exclude Mode) on page 136 show igmp group brief on page 136 show igmp group detail on page 136 |
| Output Fields | <p>Table 26 on page 134 describes the output fields for the show igmp group command. Output fields are listed in the approximate order in which they appear.</p> |

Table 26: show igmp group Output Fields

| Field Name | Field Description | Level of Output |
|-------------------|---|-----------------|
| Interface | Name of the interface that received the IGMP membership report. A name of local indicates that the local routing device joined the group itself. | All levels |
| Group | Group address. | All levels |
| Group Mode | Mode the SSM group is operating in: Include or Exclude . | All levels |
| Source | Source address. | All levels |

Table 26: show igmp group Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|------------------|---|-----------------|
| Source timeout | Time remaining until the group traffic is no longer forwarded. The timer is refreshed when a listener in include mode sends a report. A group in exclude mode or configured as a static group displays a zero timer. | detail |
| Last reported by | Address of the host that last reported membership in this group. | All levels |
| Timeout | Time remaining until the group membership is removed. | brief none |
| Group timeout | Time remaining until a group in exclude mode moves to include mode. The timer is refreshed when a listener in exclude mode sends a report. A group in include mode or configured as a static group displays a zero timer. | detail |
| Type | Type of group membership: <ul style="list-style-type: none"> • Dynamic—Host reported the membership. • Static—Membership is configured. | All levels |

Sample Output

show igmp group (Include Mode)

```
user@host> show igmp group
Interface: t1-0/1/0.0
  Group: 232.1.1.1
    Group mode: Include
    Source: 10.0.0.2
    Last reported by: 10.9.5.2
    Timeout:      24 Type: Dynamic
  Group: 232.1.1.1
    Group mode: Include
    Source: 10.0.0.3
    Last reported by: 10.9.5.2
    Timeout:      24 Type: Dynamic
  Group: 232.1.1.1
    Group mode: Include
    Source: 10.0.0.4
    Last reported by: 10.9.5.2
    Timeout:      24 Type: Dynamic
  Group: 232.1.1.2
    Group mode: Include
    Source: 10.0.0.4
    Last reported by: 10.9.5.2
    Timeout:      24 Type: Dynamic
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
  Group: 224.0.0.2
    Source: 0.0.0.0
    Last reported by: Local
    Timeout:      0 Type: Dynamic
  Group: 224.0.0.22
    Source: 0.0.0.0
    Last reported by: Local
    Timeout:      0 Type: Dynamic
```

show igmp group (Exclude Mode)

```
user@host> show igmp group
Interface: t1-0/1/0.0
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
  Group: 224.0.0.2
    Source: 0.0.0.0
    Last reported by: Local
    Timeout:      0 Type: Dynamic
  Group: 224.0.0.22
    Source: 0.0.0.0
    Last reported by: Local
    Timeout:      0 Type: Dynamic
```

show igmp group brief

The output for the **show igmp group brief** command is identical to that for the **show igmp group** command.

show igmp group detail

```
user@host> show igmp group detail
```

```
Interface: t1-0/1/0.0
  Group: 232.1.1.1
    Group mode: Include
    Source: 10.0.0.2
    Source timeout: 12
    Last reported by: 10.9.5.2
    Group timeout: 0 Type: Dynamic
  Group: 232.1.1.1
    Group mode: Include
    Source: 10.0.0.3
    Source timeout: 12
    Last reported by: 10.9.5.2
    Group timeout: 0 Type: Dynamic
  Group: 232.1.1.1
    Group mode: Include
    Source: 10.0.0.4
    Source timeout: 12
    Last reported by: 10.9.5.2
    Group timeout: 0 Type: Dynamic
  Group: 232.1.1.2
    Group mode: Include
    Source: 10.0.0.4
    Source timeout: 12
    Last reported by: 10.9.5.2
    Group timeout: 0 Type: Dynamic
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
  Group: 224.0.0.2
    Group mode: Exclude
    Source: 0.0.0.0
    Source timeout: 0
    Last reported by: Local
    Group timeout: 0 Type: Dynamic
  Group: 224.0.0.22
    Group mode: Exclude
    Source: 0.0.0.0
    Source timeout: 0
    Last reported by: Local
    Group timeout: 0 Type: Dynamic
```

show igmp interface

| | |
|---|--|
| Syntax | show igmp interface <brief detail> <interface-name> <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switches and the QFX Series) | show igmp interface <brief detail> <interface-name> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | Display information about Internet Group Management Protocol (IGMP)-enabled interfaces. |
| Options | <p>none—Display standard information about all IGMP-enabled interfaces.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>interface-name—(Optional) Display information about the specified IGMP-enabled interface only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> clear igmp membership on page 92 |
| List of Sample Output | show igmp interface on page 141 show igmp interface brief on page 141 show igmp interface detail on page 141 show igmp interface <interface-name> on page 141 |
| Output Fields | Table 27 on page 138 describes the output fields for the show igmp interface command. Output fields are listed in the approximate order in which they appear. |

Table 27: show igmp interface Output Fields

| Field Name | Field Description | Level of Output |
|------------|---|-----------------|
| Interface | Name of the interface. | All levels |
| Querier | Address of the routing device that has been elected to send membership queries. | All levels |
| State | State of the interface: Up or Down . | All levels |

Table 27: show igmp interface Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------------|---|-----------------|
| SSM Map Policy | Name of the source-specific multicast (SSM) map policy that has been applied to the IGMP interface. | All levels |
| Timeout | How long until the IGMP querier is declared to be unreachable, in seconds. | All levels |
| Version | IGMP version being used on the interface: 1, 2, or 3. | All levels |
| Groups | Number of groups on the interface. | All levels |
| Group limit | Maximum number of groups allowed on the interface. Any joins requested after the limit is reached are rejected. | All levels |
| Group threshold | Configured threshold at which a warning message is generated. This threshold is based on a percentage of groups received on the interface. If the number of groups received reaches the configured threshold, the device generates a warning message. | All levels |
| Group log-interval | Time (in seconds) between consecutive log messages. | All levels |
| Immediate Leave | State of the immediate leave option: <ul style="list-style-type: none"> • On—Indicates that the router removes a host from the multicast group as soon as the router receives a leave group message from a host associated with the interface. • Off—Indicates that after receiving a leave group message, instead of removing a host from the multicast group immediately, the router sends a group query to determine if another receiver responds. | All levels |
| Promiscuous Mode | State of the promiscuous mode option: <ul style="list-style-type: none"> • On—Indicates that the router can accept IGMP reports from subnetworks that are not associated with its interfaces. • Off—Indicates that the router can accept IGMP reports only from subnetworks that are associated with its interfaces. | All levels |
| Passive | State of the passive mode option: <ul style="list-style-type: none"> • On—Indicates that the router can run IGMP on the interface but not send or receive control traffic such as IGMP reports, queries, and leaves. • Off—Indicates that the router can run IGMP on the interface and send or receive control traffic such as IGMP reports, queries, and leaves. <p>The passive statement enables you to selectively activate up to two out of a possible three available query or control traffic options. When enabled, the following options appear after the on state declaration:</p> <ul style="list-style-type: none"> • send-general-query—The interface sends general queries. • send-group-query—The interface sends group-specific and group-source-specific queries. • allow-receive—The interface receives control traffic. | All levels |
| OIF map | Name of the OIF map (if configured) associated with the interface. | All levels |

Table 27: show igmp interface Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|------------------------------|--|-----------------|
| SSM map | Name of the source-specific multicast (SSM) map (if configured) used on the interface. | All levels |
| Configured Parameters | Information configured by the user: <ul style="list-style-type: none"> • IGMP Query Interval—Interval (in seconds) at which this router sends membership queries when it is the querier. • IGMP Query Response Interval—Time (in seconds) that the router waits for a report in response to a general query. • IGMP Last Member Query Interval—Time (in seconds) that the router waits for a report in response to a group-specific query. • IGMP Robustness Count—Number of times the router retries a query. | All levels |
| Derived Parameters | Derived information: <ul style="list-style-type: none"> • IGMP Membership Timeout—Timeout period (in seconds) for group membership. If no report is received for these groups before the timeout expires, the group membership is removed. • IGMP Other Querier Present Timeout—Time (in seconds) that the router waits for the IGMP querier to send a query. | All levels |

Sample Output

show igmp interface

```

user@host> show igmp interface
Interface: at-0/3/1.0
  Querier: 10.111.30.1
  State:      Up Timeout:   None Version:  2 Groups:    4
  SSM Map Policy: ssm-policy-A
Interface: so-1/0/0.0
  Querier: 10.111.10.1
  State:      Up Timeout:   None Version:  2 Groups:    2
  SSM Map Policy: ssm-policy-B
Interface: so-1/0/1.0
  Querier: 10.111.20.1
  State:      Up Timeout:   None Version:  2 Groups:    4
  SSM Map Policy: ssm-policy-C
Immediate Leave: On
Promiscuous Mode: Off

Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2

Derived Parameters:
IGMP Membership Timeout: 260.0
IGMP Other Querier Present Timeout: 255.0

```

show igmp interface brief

The output for the **show igmp interface brief** command is identical to that for the **show igmp interface** command. For sample output, see [show igmp interface on page 141](#).

show igmp interface detail

The output for the **show igmp interface detail** command is identical to that for the **show igmp interface** command. For sample output, see [show igmp interface on page 141](#).

show igmp interface <interface-name>

```

user@host# show igmp interface ge-3/2/0.0
Interface: ge-3/2/0.0
  Querier: 20.1.1.1
  State: Up Timeout:   None Version:  3 Groups:    1
  Group limit: 8
  Group threshold: 60
  Group log-interval: 10
  Immediate leave: Off
  Promiscuous mode: Off

```

show igmp snooping interface

| | |
|---------------------------------|---|
| Syntax | show igmp snooping interface <i>interface-name</i> <brief detail> <bridge-domain <i>bridge-domain-name</i> > <virtual-switch <i>virtual-switch-name</i> > <vlan-id <i>vlan-identifier</i> > |
| Release Information | Command introduced in Junos OS Release 8.5. |
| Description | Display IGMP snooping interface information. |
| Options | <p>none—Display detailed information.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>bridge-domain <i>bridge-domain-name</i>—(Optional) Display information about a particular bridge domain.</p> <p>virtual-switch <i>virtual-switch-name</i>—(Optional) Display information about a particular virtual switch.</p> <p>vlan-id <i>vlan-identifier</i>—(Optional) Display information about a particular VLAN.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • show igmp snooping membership on page 145 • show igmp snooping statistics on page 149 |
| List of Sample Output | show igmp snooping interface on page 144 show igmp snooping interface (Group Limit Configured) on page 144 |
| Output Fields | <p>Table 28 on page 142 lists the output fields for the show igmp snooping interface command. Output fields are listed in the approximate order in which they appear.</p> |

Table 28: show igmp snooping interface Output Fields

| Field Name | Field Description | Level of Output |
|------------------------------|--|-----------------|
| Routing-instance | Routing instance for IGMP snooping. | All levels |
| Learning Domain | Learning domain for snooping. | All levels |
| IGMP Query Interval | Frequency (in seconds) with which this router sends membership queries when it is the querier. | detail |
| IGMP Query Response Interval | Time (in seconds) that the router waits for a response to a general query. | detail |

Table 28: show igmp snooping interface Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---|--|-----------------|
| IGMP Last Member Query Interval | Time (in seconds) that the router waits for a report in response to a group-specific query. | detail |
| IGMP Robustness Count | Number of times the router retries a query. | detail |
| immediate-leave | State of immediate leave: On or Off . | All levels |
| router-interface | Router interfaces that are part of this learning domain. | All levels |
| Group limit | Maximum number of (source,group) pairs allowed per interface. When a group limit is not configured, this field is not shown. | All levels |
| interface | Interfaces that are being snooped in this learning domain. | All levels |
| Groups | Number of groups on the interface. | none |
| State | State of the interface: Up or Down . | none |
| Up Groups | Number of active multicast groups attached to the logical interface. | All levels |
| IGMP Membeship Timeout | Timeout for group membership. If no report is received for these groups before the timeout expires, the group membership is removed. | none |
| IGMP Other Querier Present Timeout | Time that the router waits for the IGMP querier to send a query. | none |

Sample Output

show igmp snooping interface

```
user@host> show igmp snooping interface
Instance: bridge-domain bar

Learning-Domain: default
Interface: ge-0/1/0.200
  State:          Up Groups:      0
  Immediate leave: Off
  Router interface: yes
Interface: ge-0/1/2.200
  State:          Up Groups:      2
  Immediate leave: On
  Router interface: no
Interface: ge-0/1/3.200
  State:          Up Groups:      1
  Immediate leave: Off
  Router interface: no

Configured Parameters:
IGMP Query Interval: 130.0
IGMP Query Response Interval: 15.0
IGMP Last Member Query Interval: 2.0
IGMP Robustness Count: 3

Derived Parameters:
IGMP Membership Timeout: 405.0
IGMP Other Querier Present Timeout: 397.500
```

Sample Output

show igmp snooping interface (Group Limit Configured)

```
user@host> show igmp snooping interface instance vpls1
Instance: vpls1

Learning-Domain: default
Interface: ge-1/3/9.0
  State:          Up Groups:      0
  Immediate leave: Off
  Router interface: yes
Interface: ge-1/3/8.0
  State:          Up Groups:      0
  Immediate leave: Off
  Router interface: yes
  Group limit:    1000

Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2
```

show igmp snooping membership

| | |
|---------------------------------|---|
| Syntax | show igmp snooping membership <brief detail> <bridge-domain <i>bridge-domain-name</i> > <group <i>group-name</i> > <virtual-switch <i>virtual-switch-name</i> > <vlan-id <i>vlan-identifier</i> > |
| Release Information | Command introduced in Junos OS Release 8.5. |
| Description | Display IGMP snooping membership information. |
| Options | <p>none—Display detailed information.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>bridge-domain <i>bridge-domain-name</i>—(Optional) Display information about a particular bridge domain.</p> <p>group <i>group-name</i> —(Optional) Display information about this group address.</p> <p>virtual-switch <i>virtual-switch-name</i>—(Optional) Display information about a particular virtual switch.</p> <p>vlan-id <i>vlan-identifier</i>—(Optional) Display information about a particular VLAN.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • show igmp snooping interface on page 142 • show igmp snooping statistics on page 149 • clear igmp snooping membership on page 95 |
| List of Sample Output | show igmp snooping membership on page 147 show igmp snooping membership (Exclude Mode) on page 147 show igmp snooping membership interface ge-0/1/2.200 on page 147 show igmp snooping membership vlan-id 1 on page 148 |
| Output Fields | Table 29 on page 145 lists the output fields for the show igmp snooping membership command. Output fields are listed in the approximate order in which they appear. |

Table 29: show igmp snooping membership Output Fields

| Field Name | Field Description | Level of Output |
|-----------------|-------------------------------------|-----------------|
| Instance | Routing instance for IGMP snooping. | All levels |
| Learning Domain | Learning domain for snooping. | All levels |

Table 29: show igmp snooping membership Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-------------------------|--|-----------------|
| Interface | Interface on which this router is a proxy. | detail |
| Up Groups | Number of active multicast groups attached to the logical interface. | All levels |
| Group | Multicast group address in the membership database. | All levels |
| Group Mode | Mode the SSM group is operating in: Include or Exclude . | All levels |
| Source | Source address used on queries. | detail |
| Last reported by | Address of source last replying to the query. | detail |
| Group Timeout | Time remaining until a group in exclude mode moves to include mode. The timer is refreshed when a listener in exclude mode sends a report. A group in include mode or configured as a static group displays a zero timer. | All levels |
| Timeout | Length of time (in seconds) left until the entry is purged. | detail |
| Type | Way that the group membership information was learned: <ul style="list-style-type: none"> • Dynamic—Group membership was learned by the IGMP protocol. • Static—Group membership was learned by configuration. | detail |
| Include receiver | Source address of receiver included in membership with timeout (in seconds). | detail |

Sample Output

**show igmp snooping
membership**

```
user@host> show igmp snooping membership
Instance: vpls2

Learning-Domain: vlan-id 2
Interface: ge-3/0/0.2
Up Groups:      0
Interface: ge-3/1/0.2
Up Groups:      0
Interface: ge-3/1/5.2
Up Groups:      0

Instance: vpls1

Learning-Domain: vlan-id 1
Interface: ge-3/0/0.1
Up Groups:      0
Interface: ge-3/1/0.1
Up Groups:      0
Interface: ge-3/1/5.1
Up Groups:      1
  Group: 225.10.10.1
    Group mode: Exclude
    Source: 0.0.0.0
    Last reported by: 100.6.85.2
    Group timeout:    173 Type: Dynamic
```

**show igmp snooping
membership (Exclude
Mode)**

```
user@host> show igmp snooping membership
Instance: vpls2

Learning-Domain: vlan-id 2
Interface: ge-3/0/0.2
Up Groups:      0
Interface: ge-3/1/0.2
Up Groups:      0
Interface: ge-3/1/5.2
Up Groups:      0

Instance: vpls1

Learning-Domain: vlan-id 1
Interface: ge-3/0/0.1
Up Groups:      0
Interface: ge-3/1/0.1
Up Groups:      0
Interface: ge-3/1/5.1
Up Groups:      1
  Group: 225.10.10.1
    Group mode: Exclude
    Source: 0.0.0.0
    Last reported by: 100.6.85.2
    Group timeout:    173 Type: Dynamic
```

**show igmp snooping
membership interface
ge-0/1/2.200**

```
user@host> show igmp snooping membership interface ge-0/1/2.200
Instance: bridge-domain bar

Learning-Domain: default
```

```
Interface: ge-0/1/2.200
  Group: 225.1.1.1
    Source: 0.0.0.0
    Timeout: 391 Type: Static
  Group: 232.1.1.1
    Source: 192.168.1.1
    Timeout: 0 Type: Static
```

**show igmp snooping
membership vlan-id 1**

```
user@host> show igmp snooping membership vlan-id 1
Instance: vpls2
```

```
Instance: vpls1
```

```
Learning-Domain: vlan-id 1
Interface: ge-3/0/0.1
Up Groups: 0
Interface: ge-3/1/0.1
Up Groups: 0
Interface: ge-3/1/5.1
Up Groups: 1
  Group: 225.10.10.1
    Group mode: Exclude
    Source: 0.0.0.0
    Last reported by: 100.6.85.2
    Group timeout: 209 Type: Dynamic
```

show igmp snooping statistics

| | |
|---------------------------------|--|
| Syntax | show igmp snooping statistics <brief detail> <bridge-domain <i>bridge-domain-name</i> > <virtual-switch <i>virtual-switch-name</i> > <vlan-id <i>vlan-identifier</i> > |
| Release Information | Command introduced in Junos OS Release 8.5. |
| Description | Display IGMP snooping statistics. |
| Options | <p>none—(Optional) Display detailed information.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>bridge-domain <i>bridge-domain-name</i>—(Optional) Display information about a particular bridge domain.</p> <p>virtual-switch <i>virtual-switch-name</i>—(Optional) Display information about a particular virtual switch.</p> <p>vlan-id <i>vlan-identifier</i>—(Optional) Display information about a particular VLAN.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • show igmp snooping interface on page 142 • show igmp snooping membership on page 145 • clear igmp snooping statistics on page 96 |
| List of Sample Output | show igmp snooping statistics on page 151 |
| Output Fields | Table 30 on page 149 lists the output fields for the show igmp snooping statistics command. Output fields are listed in the approximate order in which they appear. |

Table 30: show igmp snooping statistics Output Fields

| Field Name | Field Description | Level of Output |
|------------------------|---|-----------------|
| Routing-instance | Routing instance for IGMP snooping. | All levels |
| IGMP packet statistics | Heading for IGMP snooping statistics for all interfaces or for the specified interface. | All levels |
| learning-domain | Appears at end of “IGMP packets statistics” line. | All levels |

Table 30: show igmp snooping statistics Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-------------------------------|--|-----------------|
| IGMP Message type | Summary of IGMP statistics: <ul style="list-style-type: none"> • Membership Query—Number of membership queries sent and received. • V1 Membership Report—Number of version 1 membership reports sent and received. • DVMRP—Number of DVMRP messages sent or received. • PIM V1—Number of PIM version 1 messages sent or received. • Cisco Trace—Number of Cisco trace messages sent or received. • V2 Membership Report—Number of version 2 membership reports sent or received. • Group Leave—Number of group leave messages sent or received. • Domain Wide Report—Number of domain-wide reports sent or received. • V3 Membership Report—Number of version 3 membership reports sent or received. • Other Unknown types—Number of unknown message types received. • IGMP v3 unsupported type—Number of messages received with unknown and unsupported IGMP version 3 message types. • IGMP v3 source required for SSM—Number of IGMP version 3 messages received that contained no source. • IGMP v3 mode not applicable for SSM—Number of IGMP version 3 messages received that did not contain a mode applicable for source-specific multicast (SSM). | All levels |
| Received | Number of messages received. | All levels |
| Sent | Number of messages sent. | All levels |
| Rx errors | Number of received packets that contained errors. | All levels |
| IGMP Global Statistics | Summary of IGMP snooping statistics for all interfaces. <ul style="list-style-type: none"> • Bad Length—Number of messages received with length errors so severe that further classification could not occur. • Bad Checksum—Number of messages received with a bad IP checksum. No further classification was performed. • Rx non-local—Number of messages received from senders that are not local. | All levels |

Sample Output

show igmp snooping
statistics

user@host> show igmp snooping statistics

Routing-instance foo

IGMP packet statistics for all interfaces in learning-domain vlan-100

| IGMP Message type | Received | Sent | Rx errors |
|-------------------------------------|----------|------|-----------|
| Membership Query | 89 | 51 | 0 |
| V1 Membership Report | 0 | 0 | 0 |
| DVMRP | 0 | 0 | 0 |
| PIM V1 | 0 | 0 | 0 |
| Cisco Trace | 0 | 0 | 0 |
| V2 Membership Report | 139 | 0 | 0 |
| Group Leave | 0 | 0 | 0 |
| Domain Wide Report | 0 | 0 | 0 |
| V3 Membership Report | 136 | 0 | 0 |
| Other Unknown types | | | 0 |
| IGMP v3 unsupported type | | | 0 |
| IGMP v3 source required for SSM | | | 23 |
| IGMP v3 mode not applicable for SSM | | | 0 |

IGMP Global Statistics

| | |
|--------------|---|
| Bad Length | 0 |
| Bad Checksum | 0 |
| Rx non-local | 0 |

Routing-instance bar

IGMP packet statistics for all interfaces in learning-domain vlan-100

| IGMP Message type | Received | Sent | Rx errors |
|-------------------------------------|----------|------|-----------|
| Membership Query | 89 | 51 | 0 |
| V1 Membership Report | 0 | 0 | 0 |
| DVMRP | 0 | 0 | 0 |
| PIM V1 | 0 | 0 | 0 |
| Cisco Trace | 0 | 0 | 0 |
| V2 Membership Report | 139 | 0 | 0 |
| Group Leave | 0 | 0 | 0 |
| Domain Wide Report | 0 | 0 | 0 |
| V3 Membership Report | 136 | 0 | 0 |
| Other Unknown types | | | 0 |
| IGMP v3 unsupported type | | | 0 |
| IGMP v3 source required for SSM | | | 23 |
| IGMP v3 mode not applicable for SSM | | | 0 |

IGMP Global Statistics

| | |
|--------------|---|
| Bad Length | 0 |
| Bad Checksum | 0 |
| Rx non-local | 0 |

show igmp statistics

| | |
|---|--|
| Syntax | show igmp statistics <brief detail> <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switch and the QFX Series) | show igmp statistics <brief detail> <interface <i>interface-name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | Display Internet Group Management Protocol (IGMP) statistics. |
| Options | <p>none—Display IGMP statistics for all interfaces.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>interface <i>interface-name</i>—(Optional) Display IGMP statistics about the specified interface only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> clear igmp statistics on page 97 |
| List of Sample Output | show igmp statistics on page 154 show igmp statistics interface on page 154 |
| Output Fields | Table 31 on page 152 describes the output fields for the show igmp statistics command. Output fields are listed in the approximate order in which they appear. |

Table 31: show igmp statistics Output Fields

| Field Name | Field Description |
|------------------------|--|
| IGMP packet statistics | Heading for IGMP packet statistics for all interfaces or for the specified interface name. |

Table 31: show igmp statistics Output Fields (*continued*)

| Field Name | Field Description |
|-------------------------------|---|
| IGMP Message type | <p>Summary of IGMP statistics:</p> <ul style="list-style-type: none"> • Membership Query—Number of membership queries sent and received. • V1 Membership Report—Number of version 1 membership reports sent and received. • DVMRP—Number of DVMRP messages sent or received. • PIM V1—Number of PIM version 1 messages sent or received. • Cisco Trace—Number of Cisco trace messages sent or received. • V2 Membership Report—Number of version 2 membership reports sent or received. • Group Leave—Number of group leave messages sent or received. • Mtrace Response—Number of Mtrace response messages sent or received. • Mtrace Request—Number of Mtrace request messages sent or received. • Domain Wide Report—Number of domain-wide reports sent or received. • V3 Membership Report—Number of version 3 membership reports sent or received. • Other Unknown types—Number of unknown message types received. • IGMP v3 unsupported type—Number of messages received with unknown and unsupported IGMP version 3 message types. • IGMP v3 source required for SSM—Number of IGMP version 3 messages received that contained no source. • IGMP v3 mode not applicable for SSM—Number of IGMP version 3 messages received that did not contain a mode applicable for source-specific multicast (SSM). |
| Received | Number of messages received. |
| Sent | Number of messages sent. |
| Rx errors | Number of received packets that contained errors. |
| IGMP Global Statistics | <p>Summary of IGMP statistics for all interfaces.</p> <ul style="list-style-type: none"> • Bad Length—Number of messages received with length errors so severe that further classification could not occur. • Bad Checksum—Number of messages received with a bad IP checksum. No further classification was performed. • Bad Receive If—Number of messages received on an interface not enabled for IGMP. • Rx non-local—Number of messages received from senders that are not local. • Timed out—Number of groups that timed out as a result of not receiving an explicit leave message. • Rejected Report—Number of reports dropped because of the IGMP group policy. • Total Interfaces—Number of interfaces configured to support IGMP. |

Sample Output

show igmp statistics

```

user@host> show igmp statistics
IGMP packet statistics for all interfaces
IGMP Message type      Received      Sent  Rx errors
Membership Query        8883         459      0
V1 Membership Report    0            0        0
DVMRP                   0            0        0
PIM V1                  0            0        0
Cisco Trace             0            0        0
V2 Membership Report    0            0        0
Group Leave             0            0        0
Mtrace Response         0            0        0
Mtrace Request          0            0        0
Domain Wide Report      0            0        0
V3 Membership Report    0            0        0
Other Unknown types     0            0        0
IGMP v3 unsupported type 0            0        0
IGMP v3 source required for SSM 0            0        0
IGMP v3 mode not applicable for SSM 0            0        0

IGMP Global Statistics
Bad Length              0
Bad Checksum            0
Bad Receive If          0
Rx non-local            1227
Timed out               0
Rejected Report         0
Total Interfaces        2

```

show igmp statistics interface

```

user@host> show igmp statistics interface fe-1/0/1.0
IGMP interface packet statistics for fe-1/0/1.0
IGMP Message type      Received      Sent  Rx errors
Membership Query        0            230      0
V1 Membership Report    0            0        0

```

show mld group

| | |
|---------------------------------|---|
| Syntax | show mld group <brief detail> <group-name> <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display information about Multicast Listener Discovery (MLD) group membership. |
| Options | <p>none—Display standard information about all MLD groups.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>group-name—(Optional) Display MLD information about the specified group.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> clear mld membership on page 99 |
| List of Sample Output | <p>show mld group (Include Mode) on page 157</p> <p>show mld group (Exclude Mode) on page 157</p> <p>show mld group brief on page 157</p> <p>show mld group detail (Include Mode) on page 158</p> <p>show mld group detail (Exclude Mode) on page 158</p> |
| Output Fields | Table 32 on page 155 describes the output fields for the show mld group command. Output fields are listed in the approximate order in which they appear. |

Table 32: show mld group Output Fields

| Field Name | Field Description | Level of Output |
|-------------------------|--|-----------------|
| Interface | Name of the interface that received the MLD membership report; local means that the local router joined the group itself. | All levels |
| Group | Group address. | All levels |
| Source | Source address. | All levels |
| Group Mode | Mode the SSM group is operating in: Include or Exclude . | All levels |
| Last reported by | Address of the host that last reported membership in this group. | All levels |

Table 32: show mld group Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|----------------|---|-----------------|
| Source timeout | Time remaining until the group traffic is no longer forwarded. The timer is refreshed when a listener in include mode sends a report. A group in exclude mode or configured as a static group displays a zero timer. | detail |
| Timeout | Time remaining until the group membership is removed. | brief none |
| Group timeout | Time remaining until a group in exclude mode moves to include mode. The timer is refreshed when a listener in exclude mode sends a report. A group in include mode or configured as a static group displays a zero timer. | detail |
| Type | Type of group membership: <ul style="list-style-type: none">• Dynamic—Host reported the membership.• Static—Membership is configured. | All levels |

Sample Output

show mld group (Include Mode)

```
user@host> show mld group
Interface: fe-0/1/2.0
  Group: ff02::1:ff05:1a67
    Group mode: Include
    Source: ::
    Last reported by: fe80::2e0:81ff:fe05:1a67
    Timeout:      245 Type: Dynamic
  Group: ff02::1:ffa8:c35e
    Group mode: Include
    Source: ::
    Last reported by: fe80::2e0:81ff:fe05:1a67
    Timeout:      241 Type: Dynamic
  Group: ff02::2:43e:d7f6
    Group mode: Include
    Source: ::
    Last reported by: fe80::2e0:81ff:fe05:1a67
    Timeout:      244 Type: Dynamic
  Group: ff05::2
    Group mode: Include
    Source: ::
    Last reported by: fe80::2e0:81ff:fe05:1a67
    Timeout:      244 Type: Dynamic
Interface: local
  Group: ff02::2
    Source: ::
    Last reported by: Local
    Timeout:      0 Type: Dynamic
  Group: ff02::16
    Source: ::
    Last reported by: Local
    Timeout:      0 Type: Dynamic
```

show mld group (Exclude Mode)

```
user@host> show mld group
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
  Group: ff02::6
    Source: ::
    Last reported by: fe80::21f:12ff:feb6:4b3a
    Timeout:      245 Type: Dynamic
  Group: ff02::16
    Source: ::
    Last reported by: fe80::21f:12ff:feb6:4b3a
    Timeout:      28 Type: Dynamic
Interface: local
  Group: ff02::2
    Source: ::
    Last reported by: Local
    Timeout:      0 Type: Dynamic
  Group: ff02::16
    Source: ::
    Last reported by: Local
    Timeout:      0 Type: Dynamic
```

show mld group brief

The output for the **show mld group brief** command is identical to that for the **show mld group** command. For sample output, see [show mld group \(Include Mode\) on page 157](#)

[show mld group \(Exclude Mode\) on page 157.](#)

**show mld group detail
(Include Mode)**

```
user@host> show mld group detail
Interface: fe-0/1/2.0
  Group: ff02::1:ff05:1a67
    Group mode: Include
    Source: ::
    Last reported by: fe80::2e0:81ff:fe05:1a67
    Timeout: 224 Type: Dynamic
  Group: ff02::1:ffa8:c35e
    Group mode: Include
    Source: ::
    Last reported by: fe80::2e0:81ff:fe05:1a67
    Timeout: 220 Type: Dynamic
  Group: ff02::2:43e:d7f6
    Group mode: Include
    Source: ::
    Last reported by: fe80::2e0:81ff:fe05:1a67
    Timeout: 223 Type: Dynamic
  Group: ff05::2
    Group mode: Include
    Source: ::
    Last reported by: fe80::2e0:81ff:fe05:1a67
    Timeout: 223 Type: Dynamic
Interface: so-1/0/1.0
  Group: ff02::2
    Group mode: Include
    Source: ::
    Last reported by: fe80::280:42ff:fe15:f445
    Timeout: 258 Type: Dynamic
Interface: local
  Group: ff02::2
    Group mode: Include
    Source: ::
    Last reported by: Local
    Timeout: 0 Type: Dynamic
  Group: ff02::16
    Source: ::
    Last reported by: Local
    Timeout: 0 Type: Dynamic
```

**show mld group detail
(Exclude Mode)**

```
user@host> show mld group detail
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
  Group: ff02::6
    Group mode: Exclude
    Source: ::
    Source timeout: 0
    Last reported by: fe80::21f:12ff:feb6:4b3a
    Group timeout: 226 Type: Dynamic
  Group: ff02::16
    Group mode: Exclude
    Source: ::
    Source timeout: 0
    Last reported by: fe80::21f:12ff:feb6:4b3a
    Group timeout: 246 Type: Dynamic
Interface: local
  Group: ff02::2
    Group mode: Exclude
```



```
Source: ::  
Source timeout: 0  
Last reported by: Local  
Group timeout:      0 Type: Dynamic  
Group: ff02::16  
Group mode: Exclude  
Source: ::  
Source timeout: 0  
Last reported by: Local  
Group timeout:      0 Type: Dynamic
```

show mld interface

| | |
|---------------------------------|---|
| Syntax | show mld interface <brief detail> <interface-name> <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display information about Multicast Listener Discovery (MLD)-enabled interfaces. |
| Options | <p>none—Display standard information about all MLD-enabled interfaces.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>interface-name—(Optional) Display information about the specified interface.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> clear mld membership on page 99 |
| List of Sample Output | show mld interface on page 163 show mld interface brief on page 163 show mld interface detail on page 163 show mld interface <interface-name> on page 163 |
| Output Fields | <p>Table 33 on page 160 describes the output fields for the show mld interface command. Output fields are listed in the approximate order in which they appear.</p> |

Table 33: show mld interface Output Fields

| Field Name | Field Description | Level of Output |
|-----------------------|--|-----------------|
| Interface | Name of the interface. | All levels |
| Querier | Address of the router that has been elected to send membership queries. | All levels |
| State | State of the interface: Up or Down . | All levels |
| SSM Map Policy | Name of the source-specific multicast (SSM) map policy that has been applied to the interface. | All levels |
| SSM Map Policy | Name of the source-specific multicast (SSM) map policy at the MLD interface. | All levels |
| Timeout | How long until the MLD querier is declared to be unreachable, in seconds. | All levels |
| Version | MLD version being used on the interface: 1 or 2. | All levels |

Table 33: show mld interface Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------------|---|-----------------|
| Groups | Number of groups on the interface. | All levels |
| Passive | <p>State of the passive mode option:</p> <ul style="list-style-type: none"> • On—Indicates that the router can run IGMP or MLD on the interface but not send or receive control traffic such as IGMP or MLD reports, queries, and leaves. • Off—Indicates that the router can run IGMP or MLD on the interface and send or receive control traffic such as IGMP or MLD reports, queries, and leaves. <p>The passive statement enables you to selectively activate up to two out of a possible three available query or control traffic options. When enabled, the following options appear after the on state declaration:</p> <ul style="list-style-type: none"> • send-general-query—The interface sends general queries. • send-group-query—The interface sends group-specific and group-source-specific queries. • allow-receive—The interface receives control traffic | All levels |
| OIF map | Name of the OIF map associated to the interface. | All levels |
| SSM map | Name of the source-specific multicast (SSM) map used on the interface, if configured. | All levels |
| Group limit | Maximum number of groups allowed on the interface. Any memberships requested after the limit is reached are rejected. | All levels |
| Group threshold | <p>Configured threshold at which a warning message is generated.</p> <p>This threshold is based on a percentage of groups received on the interface. If the number of groups received reaches the configured threshold, the device generates a warning message.</p> | All levels |
| Group log-interval | Time (in seconds) between consecutive log messages. | All levels |
| Immediate Leave | <p>State of the immediate leave option:</p> <ul style="list-style-type: none"> • On—Indicates that the router removes a host from the multicast group as soon as the router receives a multicast listener done message from a host associated with the interface. • Off—Indicates that after receiving a multicast listener done message, instead of removing a host from the multicast group immediately, the router sends a group query to determine if another receiver responds. | All levels |

Table 33: show mld interface Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|------------------------------|--|-----------------|
| Configured Parameters | Information configured by the user. <ul style="list-style-type: none">• MLD Query Interval (.1 secs)—Interval at which this router sends membership queries when it is the querier.• MLD Query Response Interval (.1 secs)—Time that the router waits for a report in response to a general query.• MLD Last Member Query Interval (.1 secs)—Time that the router waits for a report in response to a group-specific query.• MLD Robustness Count—Number of times the router retries a query. | All levels |
| Derived Parameters | Derived information. <ul style="list-style-type: none">• MLD Membership Timeout (.1 secs)—Timeout period for group membership. If no report is received for these groups before the timeout expires, the group membership will be removed.• MLD Other Querier Present Timeout (.1 secs)—Time that the router waits for the IGMP querier to send a query. | All levels |

Sample Output

show mld interface

```

user@host> show mld interface
Interface: fe-0/0/0
  Querier: None
  State: Up      Timeout:      0    Version:  1    Groups:    0
  SSM Map Policy: ssm-policy-A
Interface: at-0/3/1.0
  Querier: 8038::c0a8:c345
  State: Up      Timeout:    None   Version:  1    Groups:    0
  SSM Map Policy: ssm-policy-B
Interface: fe-1/0/1.0
  Querier: ::192.168.195.73
  State: Up      Timeout:    None   Version:  1    Groups:    3
  SSM Map Policy: ssm-policy-C
  SSM map: ipv6map1
Immediate Leave: On

Configured Parameters:
MLD Query Interval (.1 secs): 1250
MLD Query Response Interval (.1 secs): 100
MLD Last Member Query Interval (.1 secs): 10
MLD Robustness Count: 2

Derived Parameters:
MLD Membership Timeout (.1secs): 2600
MLD Other Querier Present Timeout (.1 secs): 2550

```

show mld interface brief

The output for the **show mld interface brief** command is identical to that for the **show mld interface** command. For sample output, see [show mld interface on page 163](#).

show mld interface detail

The output for the **show mld interface detail** command is identical to that for the **show mld interface** command. For sample output, see [show mld interface on page 163](#).

show mld interface <interface-name>

```

user@host# show mld interface ge-3/2/0.0
Interface: ge-3/2/0.0
  Querier: 20.1.1.1
  State: Up Timeout:      None Version:  3 Groups:    1
  Group limit: 8
  Group threshold: 60
  Group log-interval: 10
  Immediate leave: Off
  Promiscuous mode: Off

```

show mld statistics

| | |
|---------------------------------|--|
| Syntax | <code>show mld statistics</code> <code><interface <i>interface-name</i>></code> <code><logical-system (all <i>logical-system-name</i>)></code> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display information about Multicast Listener Discovery (MLD) statistics. |
| Options | none —Display MLD statistics for all interfaces. interface <i>interface-name</i> —(Optional) Display statistics about the specified interface. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">• clear mld statistics on page 100 |
| List of Sample Output | show mld statistics on page 166 show mld statistics interface on page 166 |
| Output Fields | Table 34 on page 164 describes the output fields for the show mld statistics command. Output fields are listed in the approximate order in which they appear. |

Table 34: show mld statistics Output Fields

| Field Name | Field Description |
|------------|---|
| Received | Number of received packets. |
| Sent | Number of transmitted packets. |
| Rx errors | Number of received packets that contained errors. |

Table 34: show mld statistics Output Fields (*continued*)

| Field Name | Field Description |
|------------------------------|--|
| MLD Message type | <p>Summary of MLD statistics.</p> <ul style="list-style-type: none"> • Listener Query (v1/v2)—Number of membership queries sent and received. • Listener Report (v1)—Number of version 1 membership reports sent and received. • Listener Done (v1/v2)—Number of Listener Done messages sent and received. • Listener Report (v2)—Number of version 2 membership reports sent and received. • Other Unknown types—Number of unknown message types received. • MLD v2 source required for SSM—Number of MLD version 2 messages received that contained no source. • MLD v2 mode not applicable for SSM—Number of MLD version 2 messages received that did not contain a mode applicable for source-specific multicast (SSM). |
| MLD Global Statistics | <p>Summary of MLD statistics for all interfaces.</p> <ul style="list-style-type: none"> • Bad Length—Number of messages received with length errors so severe that further classification could not occur. • Bad Checksum—Number of messages received with an invalid IP checksum. No further classification was performed. • Bad Receive If—Number of messages received on an interface not enabled for MLD. • Rx non-local—Number of messages received from nonlocal senders. • Timed out—Number of groups that timed out as a result of not receiving an explicit leave message. • Rejected Report—Number of reports dropped because of the MLD group policy. • Total Interfaces—Number of interfaces configured to support IGMP. |

Sample Output

show mld statistics

```

user@host> show mld statistics
MLD packet statistics for all interfaces
MLD Message type      Received      Sent  Rx errors
Listener Query (v1/v2)    0           2      0
Listener Report (v1)      0           0      0
Listener Done (v1/v2)     0           0      0
Listener Report (v2)      0           0      0
Other Unknown types       0           0      0
MLD v2 source required for SSM  2
MLD v2 mode not applicable for SSM 0

MLD Global Statistics
Bad Length                0
Bad Checksum              0
Bad Receive If            0
Rx non-local              0
Timed out                 0
Rejected Report           0
Total Interfaces          2

```

show mld statistics interface

```

user@host> show mld statistics interface fe-1/0/1.0
MLD interface packet statistics for fe-1/0/1.0
MLD Message type      Received      Sent  Rx errors
Listener Query (v1/v2)    0           2      0
Listener Report (v1)      0           0      0
Listener Done (v1/v2)     0           0      0
Listener Report (v2)      0           0      0
Other Unknown types       0           0      0
MLD v2 source required for SSM  2
MLD v2 mode not applicable for SSM 0

MLD Global Statistics
Bad Length                0
Bad Checksum              0
Bad Receive If            0
Rx non-local              0
Timed out                 0
Rejected Report           0
Total Interfaces          2

```


show msdp

| | |
|---------------------------------|---|
| Syntax | show msdp <brief detail> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> <peer <i>peer-address</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 12.1 for the QFX Series. |
| Description | Display Multicast Source Discovery Protocol (MSDP) information. |
| Options | <p>none—Display standard MSDP information for all routing instances.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>instance <i>instance-name</i>—(Optional) Display information for the specified instance only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>peer <i>peer-address</i>—(Optional) Display information about the specified peer only.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • show msdp source on page 169 • show msdp source-active on page 171 • show msdp statistics on page 174 |
| List of Sample Output | show msdp on page 168 show msdp brief on page 168 show msdp detail on page 168 |
| Output Fields | Table 35 on page 167 describes the output fields for the show msdp command. Output fields are listed in the approximate order in which they appear. |

Table 35: show msdp Output Fields

| Field Name | Field Description | Level of Output |
|---------------|--|-----------------|
| Peer address | IP address of the peer. | All levels |
| Local address | Local address of the peer. | All levels |
| State | Status of the MSDP connection: Listen , Established , or Inactive . | All levels |
| Last up/down | Time at which the most recent peer-state change occurred. | All levels |

Table 35: show msdp Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|----------------------|---|-----------------|
| Peer-Group | Peer group name. | All levels |
| SA Count | Number of source-active cache entries advertised by each peer that were accepted, compared to the number that were received, in the format <i>number-accepted/number-received</i> . | All levels |
| Peer Connect Retries | Number of peer connection retries. | detail |
| State timer expires | Number of seconds before another message is sent to a peer. | detail |
| Peer Times out | Number of seconds to wait for a response from the peer before the peer is declared unavailable. | detail |
| SA accepted | Number of entries in the source-active cache accepted from the peer. | detail |
| SA received | Number of entries in the source-active cache received by the peer. | detail |

Sample Output

show msdp

```

user@host> show msdp
Peer address    Local address  State          Last up/down Peer-Group SA Count
198.32.8.193    198.32.8.195  Established    5d 19:25:44 North23    120/150
198.32.8.194    198.32.8.195  Established    3d 19:27:27 North23    300/345
198.32.8.196    198.32.8.195  Established    5d 19:39:36 North23    10/13
198.32.8.197    198.32.8.195  Established    5d 19:32:27 North23    5/6
198.32.8.198    198.32.8.195  Established    3d 19:33:04 North23    2305/3000

```

show msdp brief

The output for the **show msdp brief** command is identical to that for the **show msdp** command. For sample output, see [show msdp on page 168](#).

show msdp detail

```

user@host> show msdp detail
Peer: 10.255.70.15
Local address: 10.255.70.19
State: Established
Peer Connect Retries: 0
State timer expires: 22
Peer Times out: 49
SA accepted: 0
SA received: 0

```

show msdp source

| | |
|---------------------------------|---|
| Syntax | <code>show msdp source</code> <code><instance <i>instance-name</i>></code> <code><logical-system (all <i>logical-system-name</i>)></code> <code><source-address></code> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 12.1 for the QFX Series. |
| Description | Display multicast sources learned from Multicast Source Discovery Protocol (MSDP). |
| Options | none —Display standard MSDP source information for all routing instances. instance <i>instance-name</i> —(Optional) Display information for the specified instance only. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. source-address —(Optional) IP address and optional prefix length. Display information for the specified source address only. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">• show msdp on page 167• show msdp source-active on page 171• show msdp statistics on page 174 |
| List of Sample Output | show msdp source on page 170 |

Output Fields Table 36 on page 170 describes the output fields for the **show msdp source** command. Output fields are listed in the approximate order in which they appear.

Table 36: show msdp source Output Fields

| Field Name | Field Description |
|----------------|---|
| Source address | IP address of the source. |
| /Len | Length of the prefix for this IP address. |
| Type | Discovery method for this multicast source: <ul style="list-style-type: none"> • Configured—Source-active limit explicitly configured for this source. • Dynamic—Source-active limit established when this source was discovered. |
| Maximum | Source-active limit applied to this source. |
| Threshold | Source-active threshold applied to this source. |
| Exceeded | Number of source-active messages received from this source exceeding the established maximum. |

Sample Output

show msdp source

```

user@host> show msdp source
Source address /Len  Type      Maximum  Threshold  Exceeded
0.0.0.0        /0    Configured    5         none        0
10.1.0.0       /16   Configured    500       none        0
10.1.1.1       /32   Configured    10000     none        0
10.1.1.2       /32   Dynamic       6936     none        0
10.1.5.5       /32   Dynamic       500       none        123
10.2.1.1       /32   Dynamic       2         none        0

```

show msdp source-active

| | |
|---------------------------------|---|
| Syntax | <pre>show msdp source-active <brief detail> <group <i>group</i>> <instance <i>instance-name</i>> <local> <logical-system (all <i>logical-system-name</i>)> <originator <i>originator</i>> <peer <i>peer-address</i>> <source <i>source-address</i>></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p> |
| Description | Display the Multicast Source Discovery Protocol (MSDP) source-active cache. |
| Options | <p>none—Display standard MSDP source-active cache information for all routing instances.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>group <i>group</i>—(Optional) Display source-active cache information for the specified group.</p> <p>instance <i>instance-name</i>—(Optional) Display information for the specified instance.</p> <p>local—(Optional) Display all source-active caches originated by this router.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>originator <i>originator</i>—(Optional) Display information about the peer that originated the source-active cache entries.</p> <p>peer <i>peer-address</i>—(Optional) Display the source-active cache of the specified peer.</p> <p>source <i>source-address</i>—(Optional) Display the source-active cache of the specified source.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • show msdp on page 167 • show msdp source on page 169 • show msdp statistics on page 174 |
| List of Sample Output | <p>show msdp source-active on page 173</p> <p>show msdp source-active brief on page 173</p> <p>show msdp source-active detail on page 173</p> <p>show msdp source-active source on page 173</p> |
| Output Fields | Table 37 on page 172 describes the output fields for the show msdp source-active command. Output fields are listed in the approximate order in which they appear. |

Table 37: show msdp source-active Output Fields

| Field Name | Field Description |
|---|--|
| Global active source limit exceeded | Number of times all peers have exceeded configured active source limits. |
| Global active source limit maximum | Configured number of active source messages accepted by the device. |
| Global active source limit threshold | Configured threshold for applying random early discard (RED) to drop some but not all MSDP active source messages. |
| Global active source limit log-warning | Threshold at which a warning message is logged (percentage of the number of active source messages accepted by the device). |
| Global active source limit log interval | Time (in seconds) between consecutive log messages. |
| Group address | Multicast address of the group. |
| Source address | IP address of the source. |
| Peer address | IP address of the peer. |
| Originator | Router ID configured on the source of the rendezvous point (RP) that originated the message, or the loopback address when the router ID is not configured. |
| Flags | Flags: Accept , Reject , or Filtered . |

Sample Output

show msdp source-active

```
user@host> show msdp source-active
Group address  Source address  Peer address  Originator  Flags
230.0.0.0      192.168.195.46  local        10.255.14.30 Accept
230.0.0.1      192.168.195.46  local        10.255.14.30 Accept
230.0.0.2      192.168.195.46  local        10.255.14.30 Accept
230.0.0.3      192.168.195.46  local        10.255.14.30 Accept
230.0.0.4      192.168.195.46  local        10.255.14.30 Accept
```

show msdp source-active brief

The output for the **show msdp source-active brief** command is identical to that for the **show msdp source-active** command. For sample output, see [show msdp source-active on page 173](#).

show msdp source-active detail

The output for the **show msdp source-active detail** command is identical to that for the **show msdp source-active** command. For sample output, see [show msdp source-active on page 173](#).

show msdp source-active source

```
user@host> show msdp source-active source 192.168.215.246
Global active source limit exceeded: 0
Global active source limit maximum: 25000
Global active source limit threshold: 24000
Global active source limit log-warning: 100
Global active source limit log interval: 0

Group address  Source address  Peer address  Originator  Flags
226.2.2.1      192.168.215.246 10.255.182.140 10.255.182.140 Accept
226.2.2.3      192.168.215.246 10.255.182.140 10.255.182.140 Accept
226.2.2.4      192.168.215.246 10.255.182.140 10.255.182.140 Accept
226.2.2.5      192.168.215.246 10.255.182.140 10.255.182.140 Accept
226.2.2.7      192.168.215.246 10.255.182.140 10.255.182.140 Accept
226.2.2.10     192.168.215.246 10.255.182.140 10.255.182.140 Accept
226.2.2.11     192.168.215.246 10.255.182.140 10.255.182.140 Accept
226.2.2.13     192.168.215.246 10.255.182.140 10.255.182.140 Accept
226.2.2.14     192.168.215.246 10.255.182.140 10.255.182.140 Accept
226.2.2.15     192.168.215.246 10.255.182.140 10.255.182.140 Accept
```

show msdp statistics

| | |
|---------------------------------|--|
| Syntax | show msdp statistics <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> <peer <i>peer-address</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 12.1 for the QFX Series. |
| Description | Display statistics about Multicast Source Discovery Protocol (MSDP) peers. |
| Options | <p>none—Display statistics about all MSDP peers for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Display statistics about a specific MSDP instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>peer <i>peer-address</i>—(Optional) Display statistics about a particular MSDP peer.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> clear msdp statistics on page 102 |
| List of Sample Output | show msdp statistics on page 177 show msdp statistics peer on page 177 |
| Output Fields | Table 38 on page 174 describes the output fields for the show msdp statistics command. Output fields are listed in the approximate order in which they appear. |

Table 38: show msdp statistics Output Fields

| Field Name | Field Description |
|---|---|
| Global active source limit exceeded | Number of times all peers have exceeded configured active source limits. |
| Global active source limit maximum | Configured number of active source messages accepted by the device. |
| Global active source limit threshold | Configured threshold for applying random early discard (RED) to drop some but not all MSDP active source messages. |
| Global active source limit log-warning | Threshold at which a warning message is logged (percentage of the number of active source messages accepted by the device). |
| Global active source limit log interval | Time (in seconds) between consecutive log messages. |
| Peer | Address of peer. |

Table 38: show msdp statistics Output Fields (*continued*)

| Field Name | Field Description |
|-------------------------------------|---|
| Last State Change | How long ago the peer state changed. |
| Last message received from the peer | How long ago the last message was received from the peer. |
| RPF Failures | Number of reverse path forwarding (RPF) failures. |
| Remote Closes | Number of times the remote peer closed. |
| Peer Timeouts | Number of peer timeouts. |
| SA messages sent | Number of source-active messages sent. |
| SA messages received | Number of source-active messages received. |
| SA request messages sent | Number of source-active request messages sent. |
| SA request messages received | Number of source-active request messages received. |
| SA response messages sent | Number of source-active response messages sent. |
| SA response messages received | Number of source-active response messages received. |
| Active source exceeded | Number of times this peer has exceeded configured source-active limits. |
| Active source Maximum | Configured number of active source messages accepted by this peer. |
| Active source threshold | Configured threshold on this peer for applying random early discard (RED) to drop some but not all MSDP active source messages. |
| Active source log-warning | Configured threshold on this peer at which a warning message is logged (percentage of the number of active source messages accepted by the device). |
| Active source log-interval | Time (in seconds) between consecutive log messages on this peer. |
| Keepalive messages sent | Number of keepalive messages sent. |
| Keepalive messages received | Number of keepalive messages received. |
| Unknown messages received | Number of unknown messages received. |

Table 38: show msdp statistics Output Fields (*continued*)

| Field Name | Field Description |
|--------------------------------|------------------------------------|
| Error messages received | Number of error messages received. |

Sample Output

show msdp statistics

```
user@host> show msdp statistics
Global active source limit exceeded: 0
Global active source limit maximum: 10
Global active source limit threshold: 8
Global active source limit log-warning: 60
Global active source limit log interval: 60

Peer: 10.255.245.39
Last State Change: 11:54:49 (00:24:59)
Last message received from peer: 11:53:32 (00:26:16)
RPF Failures: 0
Remote Closes: 0
Peer Timeouts: 0
SA messages sent: 376
SA messages received: 459
SA request messages sent: 0
SA request messages received: 0
SA response messages sent: 0
SA response messages received: 0
Active source exceeded: 0
Active source Maximum: 10
Active source threshold: 8
Active source log-warning: 60
Active source log-interval: 120
Keepalive messages sent: 17
Keepalive messages received: 19
Unknown messages received: 0
Error messages received: 0
```

**show msdp statistics
peer**

```
user@host> show msdp statistics peer 10.255.182.140
Peer: 10.255.182.140
Last State Change: 8:19:23 (00:01:08)
Last message received from peer: 8:20:05 (00:00:26)
RPF Failures: 0
Remote Closes: 0
Peer Timeouts: 0
SA messages sent: 17
SA messages received: 16
SA request messages sent: 0
SA request messages received: 0
SA response messages sent: 0
SA response messages received: 0
Active source exceeded: 20
Active source Maximum: 10
Active source threshold: 8
Active source log-warning: 60
Active source log-interval: 120
Keepalive messages sent: 0
Keepalive messages received: 0
Unknown messages received: 0
Error messages received: 0
```

show multicast backup-pe-groups

| | |
|---------------------------------|---|
| Syntax | show multicast backup-pe-groups <address <i>pe-address</i> > <group <i>group-name</i> > <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced in Junos OS Release 9.0. |
| Description | Display backup PE router group information when ingress PE redundancy is configured. Ingress PE redundancy provides a backup resource when point-to-multipoint LSPs are configured for multicast distribution. |
| Options | <p>none—Display standard information about all backup PE groups.</p> <p>address <i>pe-address</i>—(Optional) Display the groups that a PE address is associated with.</p> <p>group <i>group</i>—(Optional) Display the backup PE group information for a particular group.</p> <p>instance <i>instance-name</i>—(Optional) Display backup PE group information for a specific multicast instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show multicast backup-pe-groups on page 179 |
| Output Fields | Table 39 on page 178 describes the output fields for the show multicast backup-pe-groups command. Output fields are listed in the approximate order in which they appear. |

Table 39: show multicast backup-pe-groups Output Fields

| Field Name | Field Description |
|-----------------|--|
| Backup PE Group | Group name. |
| Designated PE | Primary PE router. Address of the PE router that is currently forwarding traffic on the static route. |
| Transitions | Number of times that the designated PE router has transitioned from the most eligible PE router to a backup PE router and back again to the most eligible PE router. |
| Last Transition | Time of the most recent transition. |
| Local Address | Address of the local PE router. |
| Backup PE List | List of PE routers that are configured to be backups for the group. |

Sample Output

`show multicast
backup-pe-groups`

```
user@host> show multicast backup-pe-groups
Instance: master

Backup PE group: b1
  Designated PE: 10.255.165.7
  Transitions: 1
  Last Transition: 03:15:01
  Local Address: 10.255.165.7
  Backup PE List:
    10.255.165.8

Backup PE group: b2
  Designated PE: 10.255.165.7
  Transitions: 2
  Last Transition: 02:58:20
  Local Address: 10.255.165.7
  Backup PE List:
    10.255.165.9
    10.255.165.8
```

show multicast flow-map

| | |
|---|--|
| Syntax | show multicast flow-map <brief detail> <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switch and the QFX Series) | show multicast flow-map <brief detail> |
| Release Information | Command introduced in Junos OS Release 8.2. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | Display configuration information about IP multicast flow maps. |
| Options | none —Display configuration information about IP multicast flow maps on all systems. brief detail —(Optional) Display the specified level of output. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | view |
| List of Sample Output | show multicast flow-map on page 181 show multicast flow-map detail on page 181 |
| Output Fields | Table 40 on page 180 describes the output fields for the show multicast flow-map command. Output fields are listed in the approximate order in which they appear. |

Table 40: show multicast flow-map Output Fields

| Field Name | Field Description | Levels of Output |
|---------------------------|---|------------------|
| Name | Name of the flow map. | All levels |
| Policy | Name of the policy associated with the flow map. | All levels |
| Cache-timeout | Cache timeout value assigned to the flow map. | All levels |
| Bandwidth | Bandwidth setting associated with the flow map. | All levels |
| Adaptive | Whether or not adaptive mode is enabled for the flow map. | none |
| Flow-map | Name of the flow map. | detail |
| Adaptive Bandwidth | Whether or not adaptive mode is enabled for the flow map. | detail |
| Redundant Sources | Redundant sources defined for the same destination group. | detail |

Sample Output

**show multicast
flow-map**

```
user@host> show multicast flow-map
Instance: master
Name          Policy          Cache timeout    Bandwidth Adaptive
map2          policy2         never            2000000 no
map1          policy1         60 seconds      2000000 no
```

Sample Output

**show multicast
flow-map detail**

```
user@host> show multicast flow-map detail
Instance: master
Flow-map: map1
  Policy:          policy1
  Cache Timeout:   600 seconds
  Bandwidth:       2000000
  Adaptive Bandwidth: yes
  Redundant Sources: 11.11.11.11
  Redundant Sources: 11.11.11.12
  Redundant Sources: 11.11.11.13
```

show multicast forwarding-cache statistics

| | |
|---------------------------------|---|
| Syntax | show multicast forwarding-cache statistics <inet inet6> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced in Junos OS Release 12.2. |
| Description | Display IP multicast forwarding cache statistics. |
| Options | <p>none—Display multicast forwarding cache statistics for all supported address families for all routing instances.</p> <p>inet inet6—(Optional) Display multicast forwarding cache statistics for IPv4 or IPv6 family addresses, respectively.</p> <p>instance <i>instance-name</i>—(Optional) Display multicast forwarding cache statistics for a specific routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> clear multicast forwarding-cache on page 105 |
| List of Sample Output | show multicast forwarding-cache statistics on page 183 show multicast forwarding-cache statistics instance on page 183 |
| Output Fields | Table 41 on page 182 describes the output fields for the show multicast forwarding-cache statistics command. Output fields are listed in the approximate order in which they appear. |

Table 41: show multicast forwarding-cache statistics Output Fields

| Field Name | Field Description |
|--------------------|---|
| Instance | Name of the routing instance for which multicast forwarding cache statistics are displayed. |
| Family | Protocol family for which multicast forwarding cache statistics are displayed: ALL , INET , or INET6 . |
| Suppress Threshold | Maximum number of multicast forwarding cache entries that can be added to the cache. When the number of entries reaches the configured threshold, the device suspends adding new multicast forwarding cache entries. |
| Reuse Value | Number of multicast forwarding cache entries that must be reached before the device creates new multicast forwarding cache entries. When the total number of multicast forwarding cache entries is below the reuse value, the device resumes adding new multicast forwarding cache entries. |

Table 41: show multicast forwarding-cache statistics Output Fields (*continued*)

| Field Name | Field Description |
|------------------------|--|
| Warning Threshold | Threshold at which a warning message is logged (percentage of the suppress threshold). |
| Currently Used Entries | Number of currently used multicast forwarding cache entries. |

Sample Output

**show multicast
forwarding-cache
statistics**

```
user@host> show multicast forwarding-cache statistics
Instance: master Family: INET
Suppress Threshold           100
Reuse Value                  80
Warning Threshold           90
Currently Used Entries       101
```

```
Instance: master Family: INET6
Suppress Threshold           50
Reuse Value                  50
Warning Threshold           80
Currently Used Entries        3
```

**show multicast
forwarding-cache
statistics instance**

```
user@host> show multicast forwarding-cache statistics instance VPN-A
Instance: VPN-A Family: ALL
Suppress Threshold           20
Reuse Value                  16
Warning Threshold           50
Currently Used Entries       17
```

show multicast interface

| | |
|---|--|
| Syntax | show multicast interface <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switch and the QFX Series) | show multicast interface |
| Release Information | Command introduced in Junos OS Release 8.3. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | Display bandwidth information about IP multicast interfaces. |
| Options | none —Display all interfaces that have multicast configured. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | view |
| List of Sample Output | show multicast interface on page 185 |
| Output Fields | Table 42 on page 184 describes the output fields for the show multicast interface command. Output fields are listed in the approximate order in which they appear. |

Table 42: show multicast interface Output Fields

| Field Name | Field Description |
|---|--|
| Interface | Name of the multicast interface. |
| Maximum bandwidth (bps) | Maximum bandwidth setting, in bits per second, for this interface. |
| Remaining bandwidth (bps) | Amount of bandwidth, in bits per second, remaining on the interface. |
| Mapped bandwidth deduction (bps) | <p>Amount of bandwidth, in bits per second, used by any flows that are mapped to the interface.</p> <p>NOTE: Adding the mapped bandwidth deduction value to the local bandwidth deduction value results in the total deduction value for the interface.</p> <p>This field does not appear in the output when the no QoS adjustment feature is disabled.</p> |

Table 42: show multicast interface Output Fields (*continued*)

| Field Name | Field Description |
|--|--|
| Local bandwidth deduction (bps) | <p>Amount of bandwidth, in bits per second, used by any mapped flows that are traversing the interface.</p> <p>NOTE: Adding the mapped bandwidth deduction value to the local bandwidth deduction value results in the total deduction value for the interface.</p> <p>This field does not appear in the output when the no QoS adjustment feature is disabled.</p> |
| Reverse OIF mapping | <p>State of the reverse OIF mapping feature (on or off).</p> <p>NOTE: This field does not appear in the output when the no QoS adjustment feature is disabled.</p> |
| Reverse OIF mapping no QoS adjustment | <p>State of the no QoS adjustment feature (on or off) for interfaces that are using reverse OIF mapping.</p> <p>NOTE: This field does not appear in the output when the no QoS adjustment feature is disabled.</p> |
| Leave timer | <p>Amount of time a mapped interface remains active after the last mapping ends.</p> <p>NOTE: This field does not appear in the output when the no QoS adjustment feature is disabled.</p> |
| No QoS adjustment | <p>State (on) of the no QoS adjustment feature when this feature is enabled.</p> <p>NOTE: This field does not appear in the output when the no QoS adjustment feature is disabled.</p> |

Sample Output

**show multicast
interface**

```

user@host> show multicast interface
Interface          Maximum bandwidth (bps) Remaining bandwidth (bps)
fe-0/0/3           100000000                0
fe-0/0/3.210       100000000                -2000000
fe-0/0/3.220       100000000                100000000
fe-0/0/3.230       200000000                18000000
fe-0/0/2.200       100000000                100000000

```

show multicast minfo

| | |
|---------------------------------|---|
| Syntax | <code>show multicast minfo</code> <code><host></code> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | Display configuration information about IP multicast networks, including neighboring multicast router addresses. |
| Options | none —Display configuration information about all multicast networks. host —(Optional) Display configuration information about a particular host. Replace <i>host</i> with a hostname or IP address. |
| Required Privilege Level | view |
| List of Sample Output | show multicast minfo on page 187 |
| Output Fields | Table 43 on page 186 describes the output fields for the show multicast minfo command. Output fields are listed in the approximate order in which they appear. |

Table 43: show multicast minfo Output Fields

| Field Name | Field Description |
|--------------------------------------|--|
| <i>source-address</i> | Query address, hostname (DNS name or IP address of the source address), and multicast protocol version or the software version of another vendor. |
| <i>ip-address-1—>ip-address-2</i> | Queried router interface address and directly attached neighbor interface address, respectively. |
| <i>(name or ip-address)</i> | Name or IP address of neighbor. |
| <i>[metric/threshold/type/flags]</i> | Neighbor's multicast profile: <ul style="list-style-type: none"> metric—Always has a value of 1, because minfo queries the directly connected interfaces of a device. threshold—Multicast threshold time-to-live (TTL). The range of values is 0 through 255. type—Multicast connection type: pim or tunnel. flags—Flags for this route: <ul style="list-style-type: none"> querier—Queried router is the designated router for the neighboring session. leaf—Link is a leaf in the multicast network. down—Link status indicator. |

Sample Output

```
show multicast mrinfo  user@host> show multicast mrinfo 10.35.4.1
10.35.4.1 (10.35.4.1) [version 12.0]:
  192.168.195.166 -> 0.0.0.0 (local) [1/0/pim/querier/leaf]
  10.38.20.1 -> 0.0.0.0 (local) [1/0/pim/querier/leaf]
  10.47.1.1 -> 10.47.1.2 (10.47.1.2) [1/5/pim]
  0.0.0.0 -> 0.0.0.0 (local) [1/0/pim/down]
```

show multicast next-hops

| | |
|---|--|
| Syntax | <code>show multicast next-hops</code> <code><brief detail></code> <code><identifier-number></code> <code><inet inet6></code> <code><logical-system (all <i>logical-system-name</i>)></code> |
| Syntax (EX Series Switch and the QFX Series) | <code>show multicast next-hops</code> <code><brief detail></code> <code><identifier-number></code> <code><inet inet6></code> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. inet6 option introduced in Junos OS Release 10.0 for EX Series switches. detail option display of next-hop ID number introduced in Junos OS Release 11.1 for M Series and T Series routers and EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. Support for bidirectional PIM added in Junos OS Release 12.1. |
| Description | Display the entries in the IP multicast next-hop table. |
| Options | none —Display standard information about all entries in the multicast next-hop table for all supported address families. brief detail —(Optional) Display the specified level of output. When you include the detail option on M Series and T Series routers and EX Series switches, the downstream interface name includes the next-hop ID number in parentheses, in the form fe-0/1/2.0-(1048574) where 1048574 is the next-hop ID number. identifier-number —(Optional) Show a particular next hop by ID number. The range of values is 1 through 65,535 . inet inet6 —(Optional) Display entries for IPv4 or IPv6 family addresses, respectively. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | view |
| List of Sample Output | show multicast next-hops on page 190 show multicast next-hops (Bidirectional PIM on page 190 show multicast next-hops brief on page 190 show multicast next-hops detail on page 190 |
| Output Fields | Table 44 on page 189 describes the output fields for the show multicast next-hops command. Output fields are listed in the approximate order in which they appear. |

Table 44: show multicast next-hops Output Fields

| Field Name | Field Description |
|--------------------------------|--|
| Family | Protocol family (such as INET). |
| ID | Next-hop identifier of the prefix. The identifier is returned by the routing device's Packet Forwarding Engine. |
| RefCount | Number of cache entries that are using this next hop. |
| KRefCount | Kernel reference count for the next hop. |
| Downstream interface | Interface names associated with each multicast next-hop ID. |
| Incoming interface list | List of interfaces that accept incoming traffic. Only shown for routes that do not use strict RPF-based forwarding, for example for bidirectional PIM. |

Sample Output

show multicast next-hops

```
user@host> show multicast next-hops
Family: INET
ID      Refcount  KRefCount Downstream interface
262142      4          2 so-1/0/0.0
262143      2          1 mt-1/1/0.49152
262148      2          1 mt-1/1/0.32769
```

show multicast next-hops (Bidirectional PIM)

```
user@host> show multicast next-hops
Family: INET
ID      Refcount  KRefCount Downstream interface
2097151      8          4 ge-0/0/1.0

Family: INET6
ID      Refcount  KRefCount Downstream interface
2097157      2          1 ge-0/0/1.0

Family: Incoming interface list
ID      Refcount  KRefCount Downstream interface
513      5          2 lo0.0
           ge-0/0/1.0
514      5          2 lo0.0
           ge-0/0/1.0
           xe-4/1/0.0
515      3          1 lo0.0
           ge-0/0/1.0
           xe-4/1/0.0
544      1          0 lo0.0
           xe-4/1/0.0
```

show multicast next-hops brief

The output for the **show multicast next-hops brief** command is identical to that for the **show multicast next-hops** command. For sample output, see [show multicast next-hops on page 190](#).

show multicast next-hops detail

```
user@host> show multicast next-hops detail
Family: INET
ID      Refcount  KRefCount Downstream interface
1048577      2          1 fe-0/1/2.0-(1048574)
           ge-0/2/3.0-(1048576)
```


show multicast pim-to-igmp-proxy

| | |
|---|---|
| Syntax | show multicast pim-to-igmp-proxy <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switch and the QFX Series) | show multicast pim-to-igmp-proxy <instance <i>instance-name</i> > |
| Release Information | Command introduced in Junos OS Release 9.6. Command introduced in Junos OS Release 9.6 for EX Series switches. instance option introduced in Junos OS Release 10.3. instance option introduced in Junos OS Release 10.3 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | Display configuration information about PIM-to-IGMP message translation, also known as PIM-to-IGMP proxy. |
| Options | <p>none—Display configuration information about PIM-to-IGMP message translation for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Display configuration information about PIM-to-IGMP message translation for a specific multicast instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> Configuring PIM-to-IGMP and PIM-to-MLD Message Translation |
| List of Sample Output | show multicast pim-to-igmp-proxy on page 192 show multicast pim-to-igmp-proxy instance on page 192 |
| Output Fields | Table 45 on page 191 describes the output fields for the show multicast pim-to-igmp-proxy command. Output fields are listed in the order in which they appear. |

Table 45: show multicast pim-to-igmp-proxy Output Fields

| Field Name | Field Description |
|------------------------------|---|
| Instance | Routing instance. Default instance is master (inet.0 routing table). |
| Proxy state | State of PIM-to-IGMP message translation, also known as PIM-to-IGMP proxy, on the configured upstream interfaces: enabled or disabled . |
| <i>interface-name</i> | Name of upstream interface (no more than two allowed) on which PIM-to-IGMP message translation is configured. |

Sample Output

**show multicast
pim-to-igmp-proxy**

```
user@host> show multicast pim-to-igmp-proxy
Instance: master Proxy state: enabled
ge-0/1/0.1
ge-0/1/0.2
```

**show multicast
pim-to-igmp-proxy
instance**

```
user@host> show multicast pim-to-igmp-proxy instance VPN-A
Instance: VPN-A Proxy state: enabled
ge-0/1/0.1
```

show multicast pim-to-mld-proxy

| | |
|---|---|
| Syntax | show multicast pim-to-mld-proxy <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switch and the QFX Series) | show multicast pim-to-mld-proxy <instance <i>instance-name</i> > |
| Release Information | Command introduced in Junos OS Release 9.6. Command introduced in Junos OS Release 9.6 for EX Series switches. instance option introduced in Junos OS Release 10.3. instance option introduced in Junos OS Release 10.3 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | Display configuration information about PIM-to-MLD message translation, also known as PIM-to-MLD proxy. |
| Options | none —Display configuration information about PIM-to-MLD message translation for all routing instances. instance <i>instance-name</i> —(Optional) Display configuration information about PIM-to-MLD message translation for a specific multicast instance. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | view |
| List of Sample Output | show multicast pim-to-mld-proxy on page 194 show multicast pim-to-mld-proxy instance on page 194 |
| Output Fields | Table 46 on page 193 describes the output fields for the show multicast pim-to-mld-proxy command. Output fields are listed in the order in which they appear. |

Table 46: show multicast pim-to-mld-proxy Output Fields

| Field Name | Field Description |
|-----------------------|---|
| Proxy state | State of PIM-to-MLD message translation, also known as PIM-to-MLD proxy, on the configured upstream interfaces: enabled or disabled . |
| <i>interface-name</i> | Name of upstream interface (no more than two allowed) on which PIM-to-MLD message translation is configured. |

Sample Output

**show multicast
pim-to-mld-proxy**

```
user@host> show multicast pim-to-mld-proxy
Instance: master Proxy state: enabled
ge-0/5/0.1
ge-0/5/0.2
```

**show multicast
pim-to-mld-proxy
instance**

```
user@host> show multicast pim-to-mld-proxy instance VPN-A
Instance: VPN-A Proxy state: enabled
ge-0/5/0.1
```

show multicast route

| | |
|---|--|
| Syntax | <pre>show multicast route <brief detail extensive summary> <active all inactive> <group group> <inet inet6> <instance instance name> <logical-system (all logical-system-name)> <regular-expression> <source-prefix source-prefix></pre> |
| Syntax (EX Series Switch and the QFX Series) | <pre>show multicast route <brief detail extensive summary> <active all inactive> <group group> <inet inet6> <instance instance name> <regular-expression> <source-prefix source-prefix></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional PIM added in Junos OS Release 12.1.</p> |
| Description | <p>Display the entries in the IP multicast forwarding table. You can display similar information with the show route table inet.1 command.</p> |
| Options | <p>none—Display standard information about all entries in the multicast forwarding table for all routing instances.</p> <p>brief detail extensive summary—(Optional) Display the specified level of output.</p> <p>active all inactive—(Optional) Display all active entries, all entries, or all inactive entries, respectively, in the multicast forwarding table.</p> <p>group group—(Optional) Display the cache entries for a particular group.</p> <p>inet inet6—(Optional) Display multicast forwarding table entries for IPv4 or IPv6 family addresses, respectively.</p> <p>instance instance-name—(Optional) Display entries in the multicast forwarding table for a specific multicast instance.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>regular-expression—(Optional) Display information about the multicast forwarding table entries that match a UNIX OS-style regular expression.</p> |

source-prefix *source-prefix*—(Optional) Display the cache entries for a particular source prefix.

Required Privilege Level view

List of Sample Output [show multicast route on page 198](#)
[show multicast route \(Bidirectional PIM\) on page 198](#)
[show multicast route brief on page 198](#)
[show multicast route detail on page 199](#)
[show multicast route extensive \(Bidirectional PIM\) on page 199](#)
[show multicast route instance <instance-name> on page 200](#)
[show multicast route summary on page 200](#)

Output Fields [Table 47 on page 196](#) describes the output fields for the **show multicast route** command. Output fields are listed in the approximate order in which they appear.

Table 47: show multicast route Output Fields

| Field Name | Field Description | Level of Output |
|--------------------------------------|---|-------------------------|
| family | IPv4 address family (INET) or IPv6 address family (INET6). | All levels |
| Group | Group address. For any-source multicast routes, for example for bidirectional PIM, the group address includes the prefix length. | All levels |
| Source | Prefix and length of the source as it is in the multicast forwarding table. | All levels |
| Incoming interface list | List of interfaces that accept incoming traffic. Only shown for routes that do not use strict RPF-based forwarding, for example for bidirectional PIM. | All levels |
| Upstream interface | Name of the interface on which the packet with this source prefix is expected to arrive. | All levels |
| Downstream interface list | List of interface names to which the packet with this source prefix is forwarded. | All levels |
| Number of outgoing interfaces | Total number of outgoing interfaces for each (S,G) entry. | extensive |
| Session description | Name of the multicast session. | detail extensive |
| Statistics | Rate at which packets are being forwarded for this source and group entry (in Kbps and pps), and number of packets that have been forwarded to this prefix. If one or more of the kilobits per second packet forwarding statistic queries fails or times out, the statistics field displays Forwarding statistics are not available . NOTE: On QFX Series switches, this field does not report valid statistics. | detail extensive |
| Next-hop ID | Next-hop identifier of the prefix. The identifier is returned by the routing device's Packet Forwarding Engine and is also displayed in the output of the show multicast nexthops command. | detail extensive |

Table 47: show multicast route Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---|--|--------------------------|
| Incoming interface list ID | For bidirectional PIM, incoming interface list identifier. Identifiers for interfaces that accept incoming traffic. Only shown for routes that do not use strict RPF-based forwarding, for example for bidirectional PIM. | detail extensive |
| Upstream protocol | Protocol running on the interface on which the packet with this source prefix is expected to arrive. | detail extensive |
| Route type | Type of multicast route. Values can be (S,G) or (*,G). | summary |
| Route state | Whether the group is Active or Inactive . | summary extensive |
| Route count | Number of multicast routes. | summary |
| Forwarding state | Whether the prefix is pruned or forwarding. | extensive |
| Cache lifetime/timeout | Number of seconds until the prefix is removed from the multicast forwarding table. A value of never indicates a permanent forwarding entry. A value of forever indicates routes that do not have keepalive times. | extensive |
| Wrong incoming interface notifications | Number of times that the upstream interface was not available. | extensive |
| Uptime | Time since the creation of a multicast route. | extensive |

Sample Output

```
show multicast route      user@host> show multicast route
                          Family: INET

                          Group: 228.0.0.0
                            Source: 10.255.14.144/32
                            Upstream interface: local
                            Downstream interface list:
                              so-1/0/0.0

                          Group: 239.1.1.1
                            Source: 10.255.14.144/32
                            Upstream interface: local
                            Downstream interface list:
                              so-1/0/0.0

                          Group: 239.1.1.1
                            Source: 10.255.70.15/32
                            Upstream interface: so-1/0/0.0
                            Downstream interface list:
                              mt-1/1/0.49152

                          Family: INET6
```

```
show multicast route      user@host> show multicast route
(Bidirectional PIM)      Family: INET

                          Group: 224.1.1.0/24
                            Source: *
                            Incoming interface list:
                              lo0.0 ge-0/0/1.0
                            Downstream interface list:
                              ge-0/0/1.0

                          Group: 224.1.3.0/24
                            Source: *
                            Incoming interface list:
                              lo0.0 ge-0/0/1.0 xe-4/1/0.0
                            Downstream interface list:
                              ge-0/0/1.0

                          Group: 225.1.1.0/24
                            Source: *
                            Incoming interface list:
                              lo0.0 ge-0/0/1.0
                            Downstream interface list:
                              ge-0/0/1.0

                          Group: 225.1.3.0/24
                            Source: *
                            Incoming interface list:
                              lo0.0 ge-0/0/1.0 xe-4/1/0.0
                            Downstream interface list:
                              ge-0/0/1.0

                          Family: INET6
```

The output for the **show multicast route brief** command is identical to that for the **show**

**show multicast route
brief**

multicast route command. For sample output, see [show multicast route on page 198](#) or [show multicast route \(Bidirectional PIM\) on page 198](#).

**show multicast route
detail**

```
user@host> show multicast route detail
Family: INET

Group: 228.0.0.0
  Source: 10.255.14.144/32
  Upstream interface: local
  Downstream interface list:
    so-1/0/0.0
  Session description: Unknown
  Statistics: 8 kbps, 100 pps, 45272 packets
  Next-hop ID: 262142
  Upstream protocol: PIM

Group: 239.1.1.1
  Source: 10.255.14.144/32
  Upstream interface: local
  Downstream interface list:
    so-1/0/0.0
  Session description: Administratively Scoped
  Statistics: 0 kbps, 0 pps, 13404 packets
  Next-hop ID: 262142
  Upstream protocol: PIM

Group: 239.1.1.1
  Source: 10.255.70.15/32
  Upstream interface: so-1/0/0.0
  Downstream interface list:
    mt-1/1/0.49152
  Session description: Administratively Scoped
  Statistics: 46 kbps, 1000 pps, 921077 packets

  Next-hop ID: 262143
  Upstream protocol: PIM

Family: INET6
```

**show multicast route
extensive
(Bidirectional PIM)**

```
user@host> show multicast route extensive
Family: INET

Group: 224.1.1.0/24
  Source: *
  Incoming interface list:
    lo0.0 ge-0/0/1.0
  Downstream interface list:
    ge-0/0/1.0
  Number of outgoing interfaces: 1
  Session description: NOB Cross media facilities
  Statistics: 0 kbps, 0 pps, 0 packets
  Next-hop ID: 2097153
  Incoming interface list ID: 585
  Upstream protocol: PIM
  Route state: Active
  Forwarding state: Forwarding
  Cache lifetime/timeout: forever
  Wrong incoming interface notifications: 0
```

```

Group: 224.1.3.0/24
Source: *
Incoming interface list:
  lo0.0 ge-0/0/1.0 xe-4/1/0.0
Downstream interface list:
  ge-0/0/1.0
Number of outgoing interfaces: 1
Session description: NOB Cross media facilities
Statistics: 0 kbps, 0 pps, 0 packets
Next-hop ID: 2097153
Incoming interface list ID: 589
Upstream protocol: PIM
Route state: Active
Forwarding state: Forwarding
Cache lifetime/timeout: forever
Wrong incoming interface notifications: 0

```

Family: INET6

**show multicast route
instance
<instance-name>**

```

user@host> show multicast route instance v1 extensive
Instance: v1 Family: INET

```

```

Group: 224.1.1.1
Source: (null)/0
Upstream interface: fe-1/3/0.111
Downstream interface list:
  lt-0/3/0.42 lt-0/3/0.46 lt-0/3/0.43
Number of outgoing interfaces: 3

```

```

Group: 224.1.1.2
Source: (null)/0
Upstream interface: fe-1/3/0.111
Downstream interface list:
  lt-0/3/0.42 lt-0/3/0.46 lt-0/3/0.43
Number of outgoing interfaces: 3

```

```

Group: 224.1.1.3
Source: (null)/0
Upstream interface: fe-1/3/0.111
Downstream interface list:
  lt-0/3/0.42 lt-0/3/0.46 lt-0/3/0.43
Number of outgoing interfaces: 3

```

Instance: v1 Family: INET6

**show multicast route
summary**

```

user@host> show multicast route summary
Instance: master Family: INET

```

| Route type | Route state | Route count |
|------------|-------------|-------------|
| (S,G) | Active | 2 |
| (S,G) | Inactive | 3 |

Instance: master Family: INET6

show multicast rpf

| | |
|---|---|
| Syntax | <pre>show multicast rpf <inet inet6> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <prefix> <summary></pre> |
| Syntax (EX Series Switch and the QFX Series) | <pre>show multicast rpf <inet inet6> <instance <i>instance-name</i>> <prefix> <summary></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> |
| Description | Display information about multicast reverse-path-forwarding (RPF) calculations. |
| Options | <p>none—Display RPF calculation information for all supported address families.</p> <p>inet inet6—(Optional) Display the RPF calculation information for IPv4 or IPv6 family addresses, respectively.</p> <p>instance <i>instance-name</i>—(Optional) Display information about multicast RPF calculations for a specific multicast instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>prefix—(Optional) Display the RPF calculation information for the specified prefix.</p> <p>summary—(Optional) Display a summary of all multicast RPF information.</p> |
| Required Privilege Level | view |
| List of Sample Output | <p>show multicast rpf on page 203</p> <p>show multicast rpf inet6 on page 203</p> <p>show multicast rpf prefix on page 204</p> <p>show multicast rpf summary on page 204</p> |

Output Fields Table 48 on page 202 describes the output fields for the **show multicast rpf** command. Output fields are listed in the approximate order in which they appear.

Table 48: show multicast rpf Output Fields

| Field Name | Field Description |
|----------------------|---|
| Instance | Name of the routing instance. (Displayed when multicast is configured within a routing instance.) |
| Source prefix | Prefix and length of the source as it exists in the multicast forwarding table. |
| Protocol | How the route was learned. |
| Interface | Upstream RPF interface. NOTE: The displayed interface information does not apply to bidirectional PIM RP addresses. This is because the show multicast rpf command does not take into account equal-cost paths or the designated forwarder. For accurate upstream RPF interface information, always use the show pim join extensive command when bidirectional PIM is configured. |
| Neighbor | Upstream RPF neighbor. NOTE: The displayed neighbor information does not apply to bidirectional PIM. This is because the show multicast rpf command does not take into account equal-cost paths or the designated forwarder. For accurate upstream RPF neighbor information, always use the show pim join extensive command when bidirectional PIM is configured. |

Sample Output

```

show multicast rpf      user@host> show multicast rpf

Multicast RPF table: inet.0, 12 entries

0.0.0.0/0
  Protocol: Static

10.255.14.132/32
  Protocol: Direct
  Interface: lo0.0

10.255.245.91/32
  Protocol: IS-IS
  Interface: so-1/1/1.0
  Neighbor: 192.168.195.21

127.0.0.1/32
Inactive
172.16.0.0/12
  Protocol: Static
  Interface: fxp0.0
  Neighbor: 192.168.14.254

192.168.0.0/16
  Protocol: Static
  Interface: fxp0.0
  Neighbor: 192.168.14.254

192.168.14.0/24
  Protocol: Direct
  Interface: fxp0.0

192.168.14.132/32
  Protocol: Local

192.168.195.20/30
  Protocol: Direct
  Interface: so-1/1/1.0

192.168.195.22/32
  Protocol: Local

192.168.195.36/30
  Protocol: IS-IS
  Interface: so-1/1/1.0
  Neighbor: 192.168.195.21

show multicast rpf     user@host> show multicast rpf inet6
inet6

Multicast RPF table: inet6.0, 12 entries

::10.255.14.132/128
  Protocol: Direct
  Interface: lo0.0

::10.255.245.91/128
  Protocol: IS-IS

```

```
Interface: so-1/1/1.0
Neighbor: fe80::2a0:a5ff:fe28:2e8c
```

```
::192.168.195.20/126
Protocol: Direct
Interface: so-1/1/1.0
```

```
::192.168.195.22/128
Protocol: Local
```

```
::192.168.195.36/126
Protocol: IS-IS
Interface: so-1/1/1.0
Neighbor: fe80::2a0:a5ff:fe28:2e8c
```

```
::192.168.195.76/126
Protocol: Direct
Interface: fe-2/2/0.0
```

```
::192.168.195.77/128
Protocol: Local
```

```
fe80::/64
Protocol: Direct
Interface: so-1/1/1.0
```

```
fe80::290:69ff:fe0c:993a/128
Protocol: Local
```

```
fe80::2a0:a5ff:fe12:84f/128
Protocol: Direct
Interface: lo0.0
```

```
ff02::2/128
Protocol: PIM
```

```
ff02::d/128
Protocol: PIM
```

show multicast rpf prefix

```
user@host> show multicast rpf ff02::/16

Multicast RPF table: inet6.0, 13 entries

ff02::2/128
  Protocol: PIM

ff02::d/128
  Protocol: PIM

...
```

show multicast rpf summary

```
user@host> show multicast rpf summary

Multicast RPF table: inet.0, 16 entries
Multicast RPF table: inet6.0, 12 entries
```

show multicast scope

| | |
|---|--|
| Syntax | show multicast scope <inet inet6> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switch and the QFX Series) | show multicast scope <inet inet6> <instance <i>instance-name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | Display administratively scoped IP multicast information. |
| Options | <p>none—Display standard information about administratively scoped multicast information for all supported address families in all routing instances.</p> <p>inet inet6—(Optional) Display scoped multicast information for IPv4 or IPv6 family addresses, respectively.</p> <p>instance <i>instance-name</i>—(Optional) Display administratively scoped information for a specific multicast instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show multicast scope on page 206 show multicast scope inet on page 206 show multicast scope inet6 on page 206 |
| Output Fields | Table 49 on page 205 describes the output fields for the show multicast scope command. Output fields are listed in the approximate order in which they appear. |

Table 49: show multicast scope Output Fields

| Field Name | Field Description |
|-----------------|---|
| Scope name | Name of the multicast scope. |
| Group Prefix | Range of multicast groups that are scoped. |
| Interface | Interface that is the boundary of the administrative scope. |
| Resolve Rejects | Number of kernel resolve rejects. |

Sample Output

```
show multicast scope user@host> show multicast scope
```

| Scope name | Group Prefix | Interface | Resolve Rejects |
|------------|----------------|------------|--------------------|
| 232-net | 232.232.0.0/16 | fe-0/0/0.1 | 0 |
| local | 239.255.0.0/16 | fe-0/0/0.1 | 0 |
| local | ff05::/16 | fe-0/0/0.1 | 0 |
| larry | ff05::1234/128 | fe-0/0/0.1 | 0 |

```
show multicast scope user@host> show multicast scope inet  
inet
```

| Scope name | Group Prefix | Interface | Resolve Rejects |
|------------|----------------|------------|--------------------|
| 232-net | 232.232.0.0/16 | fe-0/0/0.1 | 0 |
| local | 239.255.0.0/16 | fe-0/0/0.1 | 0 |

```
show multicast scope user@host> show multicast scope inet6  
inet6
```

| Scope name | Group Prefix | Interface | Resolve Rejects |
|------------|----------------|------------|--------------------|
| local | ff05::/16 | fe-0/0/0.1 | 0 |
| larry | ff05::1234/128 | fe-0/0/0.1 | 0 |

show multicast sessions

| | |
|---|--|
| Syntax | show multicast sessions <brief detail extensive> <logical-system (all <i>logical-system-name</i>)> < <i>regular-expression</i> > |
| Syntax (EX Series Switch and the QFX Series) | show multicast sessions <brief detail extensive> < <i>regular-expression</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | Display information about announced IP multicast sessions. |
| Options | <p>none—Display standard information about all multicast sessions for all routing instances.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><i>regular-expression</i>—(Optional) Display information about announced sessions that match a UNIX-style regular expression.</p> |
| Required Privilege Level | view |
| List of Sample Output | show multicast sessions on page 208 show multicast sessions <i>regular-expression</i> detail on page 208 |
| Output Fields | <p>Table 50 on page 207 describes the output fields for the show multicast sessions command. Output fields are listed in the approximate order in which they appear.</p> |

Table 50: show multicast sessions Output Fields

| Field Name | Field Description |
|---------------------|---|
| <i>session-name</i> | Name of the known announced multicast sessions. |

Sample Output

**show multicast
sessions**

```
user@host> show multicast sessions
1-Department of Biological Sciences, LSU
...
Monterey Bay - DockCam
Monterey Bay - JettyCam
Monterey Bay - StandCam
Monterey DockCam
Monterey DockCam / ROV cam
...
NASA TV (MPEG-1)
...
UO Broadcast - NASA Videos - 25 Years of Progress
UO Broadcast - NASA Videos - Journey through the Solar System
UO Broadcast - NASA Videos - Life in the Universe
UO Broadcast - NASA Videos - Nasa and the Airplane
UO Broadcasts OPB's Oregon Story
UO DOD News Clips
UO Medical Management of Biological Casualties (1)
UO Medical Management of Biological Casualties (2)
UO Medical Management of Biological Casualties (3)
...
376 active sessions.
```

**show multicast
sessions**

```
user@host> show multicast sessions "NASA TV" detail
SDP Version: 0  Originated by: -@128.223.83.33
Session: NASA TV (MPEG-1)
```

regular-expression detail

```

Description: NASA television in MPEG-1 format, provided by Private University.
Please contact the UO if you have problems with this feed.
Email: Your Name Here <multicast@lists.private.edu>
Phone: Your Name Here <888/555-1212>
Bandwidth: AS:1000
Start time: permanent
Stop time: none
Attribute: type:broadcast
Attribute: tool:IP/TV Content Manager 3.4.14
Attribute: live:capture:1
Attribute: x-iptv-capture:mp1s
Media: video 54302 RTP/AVP 32 31 96 97
Connection Data: 224.2.231.45 ttl 127
Attribute: quality:8
Attribute: framerate:30
Attribute: rtpmap:96 WBIH/90000
Attribute: rtpmap:97 MP4V-ES/90000
Attribute: x-iptv-svr:video 128.223.91.191 live
Attribute: fmp:32 type=mpeg1
Media: audio 28848 RTP/AVP 14 0 96 3 5 97 98 99 100 101 102 10 11 103 104 105 106
Connection Data: 224.2.145.37 ttl 127
Attribute: rtpmap:96 X-WAVE/8000
Attribute: rtpmap:97 L8/8000/2
Attribute: rtpmap:98 L8/8000
Attribute: rtpmap:99 L8/22050/2
Attribute: rtpmap:100 L8/22050
Attribute: rtpmap:101 L8/11025/2
Attribute: rtpmap:102 L8/11025
Attribute: rtpmap:103 L16/22050/2
Attribute: rtpmap:104 L16/22050

```

1 matching sessions.

show multicast snooping next-hops

| | |
|---------------------------------|--|
| Syntax | show multicast snooping next-hops <brief detail> |
| Release Information | Command introduced in Junos OS Release 11.2. |
| Description | Display information about the IP multicast snooping next-hops. |
| Options | brief detail —(Optional) Display the specified level of output. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">• show multicast snooping statistics on page 215• clear multicast snooping statistics on page 108 |
| List of Sample Output | show multicast snooping next-hops on page 211 |
| Output Fields | Table 51 on page 210 describes the output fields for the show multicast snooping next-hops command. Output fields are listed in the approximate order in which they appear. |

Table 51: show multicast snooping next-hops Output Fields

| Field Name | Field Description |
|-----------------------------|---|
| Family | Protocol family for which multicast snooping next hops are displayed: INET or INET6 . |
| Refcount | Number of cache entries that are using this next hop. |
| KRefcount | Kernel reference count for the next hop. |
| Downstream interface | Interface names associated with each multicast next-hop ID. |
| Nexthop Id | Identifier for the next-hop. |

Sample Output

show multicast
snooping next-hops

```
user@host> show multicast snooping next-hops
Family: INET
ID          Refcount KRefCount Downstream interface Nexthop Id
1048574      4          1 ge-0/1/0.1000
              ge-0/1/2.1000
              ge-0/1/3.1000
1048574      4          1 ge-0/1/0.1000-(2000)
              1048575
              1048576
1048575      2          0 ge-0/1/2.1000-(2001)
              ge-0/1/3.1000-(2002)
1048576      2          0 lsi.1048578-(2003)
              lsi.1048579-(2004)
```

show multicast snooping route

| | |
|---------------------------------|--|
| Syntax | <pre>show multicast snooping route <brief detail extensive> <active all inactive> <bridge-domain <i>bridge-domain-name</i>> <group <i>group</i>> <instance <i>instance-name</i>> <mesh-group <i>mesh-group-name</i>> <<i>regular-expression</i>> <source-prefix <i>source-prefix</i>></pre> |
| Release Information | Command introduced in Junos OS Release 8.5. |
| Description | Display the entries in the IP multicast snooping forwarding table. You can display some of this information with the show route table inet.1 command. |
| Options | <p>none—Display standard information about all entries in the multicast snooping table for all virtual switches and all bridge domains.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>active all inactive—(Optional) Display all active entries, all entries, or all inactive entries, respectively, in the multicast snooping table.</p> <p>bridge-domain <i>bridge-domain</i>—(Optional) Display the entries for a particular bridge domain.</p> <p>group <i>group</i>—(Optional) Display the entries for a particular group.</p> <p>instance <i>instance-name</i>—(Optional) Display the entries for a multicast instance.</p> <p>mesh-group <i>mesh-group-name</i>—(Optional) Display the entries for a particular mesh group.</p> <p><i>regular-expression</i>—(Optional) Display information about the multicast forwarding table entries that match a UNIX-style regular expression.</p> <p>source-prefix <i>source-prefix</i>—(Optional) Display the entries for a particular source prefix.</p> |
| Required Privilege Level | view |
| List of Sample Output | <p>show multicast snooping route bridge-domain on page 214</p> <p>show multicast snooping route instance vs on page 214</p> |
| Output Fields | Table 52 on page 213 describes the output fields for the show multicast snooping route command. Output fields are listed in the approximate order in which they appear. |

Table 52: show multicast snooping route Output Fields

| Field Name | Field Description | Level of Output |
|-------------------------------|---|-------------------------|
| Nexthop Bulking | Displays whether next-hop bulk updating is ON or OFF (only for routing-instance type of virtual switch or vpls). | All levels |
| Family | IPv4 address family (INET) or IPv6 address family (INET6). | All levels |
| Group | Group address. | All levels |
| Source | Prefix and length of the source as it is in the multicast forwarding table. | All levels |
| Routing-instance | Name of the routing instance to which this routing information applies. (Displayed when multicast is configured within a routing instance.) | All levels |
| Learning Domain | Name of the learning domain to which this routing information applies. | detail extensive |
| Statistics | Rate at which packets are being forwarded for this source and group entry (in Kbps and pps), and number of packets that have been forwarded to this prefix. | detail extensive |
| Next-hop ID | Next-hop identifier of the prefix. The identifier is returned by the router's Packet Forwarding Engine and is also displayed in the output of the show multicast nexthops command. | detail extensive |
| Route state | Whether the group is Active or Inactive . | extensive |
| Forwarding state | Whether the prefix is Pruned or Forwarding . | extensive |
| Cache lifetime/timeout | Number of seconds until the prefix is removed from the multicast forwarding table. A value of never indicates a permanent forwarding entry. | extensive |

Sample Output

**show multicast
snooping route
bridge-domain**

```
user@host> show multicast snooping route bridge-domain br-dom-1 extensive
Family: INET

Group: 232.1.1.1
  Source: 192.168.3.100/32
  Downstream interface list:
    ge-0/1/0.200
  Statistics: 0 kbps, 0 pps, 1 packets
  Next-hop ID: 1048577
  Route state: Active
  Forwarding state: Forwarding
  Cache lifetime/timeout: 240 seconds
```

**show multicast
snooping route
instance vs**

```
user@host> show multicast snooping route instance vs
Nexthop Bulking: ON

Family: INET

Group: 224.0.0.0
  Bridge-domain: vsid500

Group: 225.1.0.1
  Bridge-domain: vsid500
  Downstream interface list: vsid500
    ge-0/3/8.500 ge-1/1/9.500 ge1/2/5.500
```


show multicast snooping statistics

| | |
|---------------------------------|---|
| Syntax | show multicast snooping statistics <instance <i>instance-name</i> > <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced in Junos OS Release 8.5. |
| Description | Display IP multicast snooping statistics. |
| Options | <p>none—Display multicast snooping statistics for all supported address families for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Display statistics for a specific routing instance.</p> <p>interface <i>interface-name</i>—(Optional) Display statistics for a specific interface.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Additional Information | The input and output interface multicast snooping statistics are consistent, but not timely. They are constructed from the forwarding statistics, which are gathered at 30-second intervals. Therefore, the output from this command always lags the true count by up to 30 seconds. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> clear multicast snooping statistics on page 108 |
| List of Sample Output | show multicast snooping statistics on page 217 |
| Output Fields | Table 53 on page 215 describes the output fields for the show multicast snooping statistics command. Output fields are listed in the approximate order in which they appear. |

Table 53: show multicast snooping statistics Output Fields

| Field Name | Field Description |
|------------------|---|
| Routing-instance | Name of the routing instance. (Displayed when multicast is configured within a routing instance.) |
| Family | Protocol family for which multicast statistics are displayed: INET or INET6 . |
| Interface | Name of the interface for which statistics are being reported. |
| Routing Protocol | Primary multicast protocol on the interface: PIM , DVMRP for INET , or PIM for INET6 . |
| Mismatch | Number of multicast packets that did not arrive on the correct upstream interface. |

Table 53: show multicast snooping statistics Output Fields (*continued*)

| Field Name | Field Description |
|--------------------------|---|
| Kernel Resolve | Number of resolve requests processed by the primary multicast protocol on the interface. |
| Resolve No Route | Number of resolve requests that were ignored because there was no route to the source. |
| In Kbytes | Total accumulated incoming packets (in KB) since the last time the clear multicast snooping statistics command was issued. |
| Out Kbytes | Total accumulated outgoing packets (in KB) since the last time the clear multicast snooping statistics command was issued. |
| Mismatch error | Number of mismatches that were ignored because of internal errors. |
| Mismatch No Route | Number of mismatches that were ignored because there was no route to the source. |
| Routing Notify | Number of times that the multicast routing system has been notified of a new multicast source by a multicast routing protocol. |
| Resolve Error | Number of resolve requests that were ignored because of internal errors. |
| In packets | Total number of incoming packets since the last time the clear multicast snooping statistics command was issued. |
| Out packets | Total number of outgoing packets since the last time the clear multicast snooping statistics command was issued. |

Sample Output

**show multicast
snooping statistics**

```

user@host> show multicast snooping statistics
Routing-instance: foo
Family: INET
Interface: fe-0/0/2.200
  Routing protocol:      PIM  Mismatch error:      0
  Mismatch:              0   Mismatch no route:    0
  Kernel resolve:        22   Routing notify:        0
  Resolve no route:      0   Resolve error:         0
  Resolve filtered:      0   Notify filtered:       0
  In kbytes:              0   In packets:            0
  Out kbytes:             0   Out packets:           0

Routing-instance: bar
Family: INET
Interface: fe-0/1/2.200
  Routing protocol:      PIM  Mismatch error:      0
  Mismatch:              0   Mismatch no route:    0
  Kernel resolve:        22   Routing notify:        0
  Resolve no route:      0   Resolve error:         0
  Resolve filtered:      0   Notify filtered:       0
  In kbytes:              0   In packets:            0
  Out kbytes:             0   Out packets:           0

```

show multicast statistics

| | |
|---------------------------------|--|
| Syntax | show multicast statistics <inet inet6> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display IP multicast statistics. |
| Options | <p>none—Display multicast statistics for all supported address families for all routing instances.</p> <p>inet inet6—(Optional) Display multicast statistics for IPv4 or IPv6 family addresses, respectively.</p> <p>instance <i>instance-name</i>—(Optional) Display statistics for a specific routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Additional Information | The input and output interface multicast statistics are consistent, but not timely. They are constructed from the forwarding statistics, which are gathered at 30-second intervals. Therefore, the output from this command always lags the true count by up to 30 seconds. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • clear multicast statistics on page 109 |
| List of Sample Output | show multicast statistics on page 220 |
| Output Fields | Table 54 on page 218 describes the output fields for the show multicast statistics command. Output fields are listed in the approximate order in which they appear. |

Table 54: show multicast statistics Output Fields

| Field Name | Field Description |
|------------------|---|
| Family | Protocol family for which multicast statistics are displayed: INET or INET6 . |
| Interface | Name of the interface for which statistics are being reported. |
| Routing Protocol | Primary multicast protocol on the interface: PIM , DVMRP for INET , or PIM for INET6 . |
| Mismatch | Number of multicast packets that did not arrive on the correct upstream interface. |
| Kernel Resolve | Number of resolve requests processed by the primary multicast protocol on the interface. |

Table 54: show multicast statistics Output Fields (*continued*)

| Field Name | Field Description |
|---|--|
| Resolve No Route | Number of resolve requests that were ignored because there was no route to the source. |
| In Kbytes | Total accumulated incoming packets (in KB) since the last time the clear multicast statistics command was issued. |
| Out Kbytes | Total accumulated outgoing packets (in KB) since the last time the clear multicast statistics command was issued. |
| Mismatch error | Number of mismatches that were ignored because of internal errors. |
| Mismatch No Route | Number of mismatches that were ignored because there was no route to the source. |
| Routing Notify | Number of times that the multicast routing system has been notified of a new multicast source by a multicast routing protocol . |
| Resolve Error | Number of resolve requests that were ignored because of internal errors. |
| In Packets | Total number of incoming packets since the last time the clear multicast statistics command was issued. |
| Out Packets | Total number of outgoing packets since the last time the clear multicast statistics command was issued. |
| Resolve requests on interfaces not enabled for multicast <i>n</i> | Number of resolve requests on interfaces that are not enabled for multicast that have accumulated since the clear multicast statistics command was last issued. |
| Resolve requests with no route to source <i>n</i> | Number of resolve requests with no route to the source that have accumulated since the clear multicast statistics command was last issued. |
| Routing notifications on interfaces not enabled for multicast <i>n</i> | Number of routing notifications on interfaces not enabled for multicast that have accumulated since the clear multicast statistics command was last issued. |
| Routing notifications with no route to source <i>n</i> | Number of routing notifications with no route to the source that have accumulated since the clear multicast statistics command was last issued. |
| Interface Mismatches on interfaces not enabled for multicast <i>n</i> | Number of interface mismatches on interfaces not enabled for multicast that have accumulated since the clear multicast statistics command was last issued. |
| Group Membership on interfaces not enabled for multicast <i>n</i> | Number of group memberships on interfaces not enabled for multicast that have accumulated since the clear multicast statistics command was last issued. |

Sample Output

show multicast statistics

```

user@host> show multicast statistics
Address family: INET
Interface: fe-0/0/0
  Routing Protocol:      PIM  Mismatch error:      0
  Mismatch:              0    Mismatch No Route:    0
  Kernel Resolve:       10    Routing Notify:       0
  Resolve No Route:     0     Resolve Error:        0
  In Kbytes:            4641   In Packets:           50454
  Out Kbytes:           0     Out Packets:          0
Interface: so-0/1/1.0
  Routing Protocol:      PIM  Mismatch error:      0
  Mismatch:              0    Mismatch No Route:    0
  Kernel Resolve:       0     Routing Notify:       0
  Resolve No Route:     0     Resolve Error:        0
  In Kbytes:            0     In Packets:           0
  Out Kbytes:          4641   Out Packets:          50454

Resolve requests on interfaces not enabled for multicast 0
Resolve requests with no route to source 0
Routing notifications on interfaces not enabled for multicast 0
Routing notifications with no route to source 0
Interface Mismatches on interfaces not enabled for multicast 0
Group Membership on interfaces not enabled for multicast 25

Address family: INET6
Interface: fe-0/0/0.0
  Routing Protocol:      PIM  Mismatch error:      0
  Mismatch:              0    Mismatch No Route:    0
  Kernel Resolve:       0     Routing Notify:       0
  Resolve No Route:     0     Resolve Error:        0
  In Kbytes:            0     In Packets:           0
  Out Kbytes:           0     Out Packets:          0
Interface: so-0/1/1.0
  Routing Protocol:      PIM  Mismatch error:      0
  Mismatch:              0    Mismatch No Route:    0
  Kernel Resolve:       0     Routing Notify:       0
  Resolve No Route:     0     Resolve Error:        0
  In Kbytes:            0     In Packets:           0
  Out Kbytes:           0     Out Packets:          0

Resolve requests on interfaces not enabled for multicast 0
Resolve requests with no route to source 0
Routing notifications on interfaces not enabled for multicast 0
Routing notifications with no route to source 0
Interface Mismatches on interfaces not enabled for multicast 0
Group Membership on interfaces not enabled for multicast 0

```

show multicast usage

| | |
|---|--|
| Syntax | show multicast usage <brief detail> <inet inet6> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switch and the QFX Series) | show multicast usage <brief detail> <inet inet6> <instance <i>instance-name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | Display usage information about the 10 most active Distance Vector Multicast Routing Protocol (DVMRP) or Protocol Independent Multicast (PIM) groups. |
| Options | <p>none—Display multicast usage information for all supported address families for all routing instances.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>inet inet6—(Optional) Display usage information for IPv4 or IPv6 family addresses, respectively.</p> <p>instance <i>instance-name</i>—(Optional) Display information about the most active DVMRP or PIM groups for a specific multicast instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show multicast usage on page 223 show multicast usage brief on page 223 show multicast usage instance on page 223 show multicast usage detail on page 223 |
| Output Fields | Table 55 on page 221 describes the output fields for the show multicast usage command. Output fields are listed in the approximate order in which they appear. |

Table 55: show multicast usage Output Fields

| Field Name | Field Description |
|-----------------|---|
| Instance | Name of the routing instance. (Displayed when multicast is configured within a routing instance.) |

Table 55: show multicast usage Output Fields (*continued*)

| Field Name | Field Description |
|----------------|--|
| Group | Group address. |
| Sources | Number of sources. |
| Packets | Number of packets that have been forwarded to this prefix. If one or more of the packets forwarded statistic queries fails or times out, the packets field displays unavailable . |
| Bytes | Number of bytes that have been forwarded to this prefix. If one or more of the packets forwarded statistic queries fails or times out, the bytes field displays unavailable . |
| Prefix | IP address. |
| /len | Prefix length. |
| Groups | Number of multicast groups. |

Sample Output

show multicast usage

```
user@host> show multicast usage
Group          Sources Packets      Bytes
228.0.0.0      1         52847      4439148
239.1.1.1      2         13450      1125530

Prefix         /len Groups Packets      Bytes
10.255.14.144  /32  2       66254      5561304
10.255.70.15   /32  1        43       3374...
```

show multicast usage brief

The output for the **show multicast usage brief** command is identical to that for the **show multicast usage** command. For sample output, see [show multicast usage on page 223](#).

show multicast usage instance

```
user@host> show multicast usage instance VPN-A
Group          Sources Packets      Bytes
224.2.127.254  1         5538      509496
224.0.1.39     1          13         624
224.0.1.40     1          13         624

Prefix         /len Groups Packets      Bytes
192.168.195.34 /32  1       5538      509496
10.255.14.30   /32  1        13         624
10.255.245.91  /32  1        13         624
...
```

show multicast usage detail

```
user@host> show multicast usage detail
Group          Sources Packets      Bytes
228.0.0.0      1         53159      4465356
  Source: 10.255.14.144 /32 Packets: 53159 Bytes: 4465356
239.1.1.1      2         13450      1125530
  Source: 10.255.14.144 /32 Packets: 13407 Bytes: 1122156
  Source: 10.255.70.15  /32 Packets: 43 Bytes: 3374

Prefix         /len Groups Packets      Bytes
10.255.14.144  /32  2       66566      5587512
  Group: 228.0.0.0      Packets: 53159 Bytes: 4465356
  Group: 239.1.1.1      Packets: 13407 Bytes: 1122156
10.255.70.15   /32  1        43       3374
  Group: 239.1.1.1      Packets: 43 Bytes: 3374
```

show pgm negative-acknowledgments

| | |
|---------------------------------|---|
| Syntax | show pgm negative-acknowledgments |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display the sent or received Pragmatic General Multicast (PGM) negative acknowledgments (NAKs), the source-path message (SPM) sequence number being negatively acknowledged, and the current state of repair. |
| Options | This command has no options. |
| Required Privilege Level | view |
| List of Sample Output | show pgm negative-acknowledgments on page 225 |
| Output Fields | Table 56 on page 224 describes the output fields for the show pgm negative-acknowledgments command. Output fields are listed in the approximate order in which they appear. |

Table 56: show pgm negative-acknowledgments Output Fields

| Field Name | Field Description |
|---|---|
| Global source id | Global source identifier (GSI), which combines with the source port to determine the transport session identifier (TSI). |
| Network layer address | Network layer address of the local system. |
| Source port | Source port number, which is combined with the GSI to determine the TSI. |
| SPM sequence number | Numeric sequence identifier of the source-path message. |
| Window (trailing/leading sequence) | Range of sequence numbers used by the source for sequentially numbering and transmitting the most recent packets. The trailing (or left) edge of the transmit window is the sequence number of the oldest data packet available for repair from a source. The leading (or right) edge of the transmit window is defined as the sequence number of the most recent data packet a source has transmitted. |
| Outstanding NAKS | <p>Total number of outstanding negative acknowledgments sent or received by the local system. NAK packets indicate that a packet in the expected original data sequence has been detected as missing.</p> <ul style="list-style-type: none"> • Sequence number—Numeric sequence identifier of the source-path message. • Group—Group address. • Source—Multicast source. • Interface—Interface name. • Receiver—IP address receiving the multicast. |

Sample Output

`show pgm negative-acknowledgments`

```
user@host> show pgm negative-acknowledgments
Global source ID: 010203040506 Source port: 1111
  Network layer address: 10.38.0.1
  SPM sequence number: 1
  Window (trailing/leading sequence): 0/1
  Outstanding NAKs:
    Sequence number: 1
    Group: 225.1.1.1
    Source: 192.168.195.121
    Interface: t3-0/2/0:0 Receiver: 10.38.0.10
```

show pgm source-path-messages

| | |
|---------------------------------|--|
| Syntax | show pgm source-path-messages |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display the Pragmatic General Multicast (PGM) source-path messages received. |
| Options | This command has no options. |
| Required Privilege Level | view |
| List of Sample Output | show pgm source-path-messages on page 226 |
| Output Fields | Table 57 on page 226 describes the output fields for the show pgm source-path-messages command. Output fields are listed in the approximate order in which they appear. |

Table 57: show pgm source-path-messages Output Fields

| Field Name | Field Description |
|------------------------------|--|
| Global source ID | Global source identifier (GSI), which combines with the source port to determine the transport session identifier (TSI). |
| Port | Source port number, which combines with the GSI to determine the TSI. |
| SPM number | Numeric sequence identifier of the source-path message. |
| Trail number | Sequence number of the oldest data packet available for repair from a source. |
| Lead number | Sequence number of the most recent data packet a source has transmitted. |
| Network layer address | Network layer address of the local system. |

Sample Output

```

show pgm
source-path-messages
user@host> show pgm source-path-messages
Global source ID  Port  SPM number  Trail number  Lead number  Network layer address
010203040506     1111         1           0             1    10.38.0.1

```

show pgm statistics

| | |
|---------------------------------|--|
| Syntax | show pgm statistics |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display Pragmatic General Multicast (PGM) packet statistics, including general loss and repair statistics. |
| Options | This command has no options. |
| Required Privilege Level | view |
| List of Sample Output | show pgm statistics on page 229 |
| Output Fields | Table 58 on page 227 describes the output fields for the show pgm statistics command. Output fields are listed in the approximate order in which they appear. |

Table 58: show pgm statistics Output Fields

| Field Name | Field Description |
|--|---|
| PGM type, # received, # sent | <p>Number of packets received and sent for the following PGM packet types:</p> <ul style="list-style-type: none"> • SPM—Number of total source path messages received and sent by the local system. Source path messages (SPMs) are sent by a source to establish the source path state in network elements and to provide the transmit-window state to receivers. • POLL—Total number of poll requests received and sent by the local system. • POLR—Total number of poll responses received and sent by the local system. • ODATA—Total number of original data packets received and sent by the local system. • RDATA—Total number of repair data packets received and sent by the local system. RDATA packets are generated in response to negative acknowledgments (NAKs), which indicate a missing packet from the original data sequence. • NAK—Total number of negative acknowledgments received and sent by the local system. NAK packets indicate that a packet in the expected original data sequence has been detected as missing. • NULLNAK—Total number of null negative acknowledgments received and sent by the local system. NULLNAKs are transmitted by a designated local repairer that receives NAKs redirected to it by either receivers or network elements to provide flow-control feedback to a source. • NCF—Total number of NAK confirmations received and sent by the local system. NAK confirmations are generated in response to NAK packets that are received. • SPMR—Total number of source path message requests (SPMRs) received and sent by the local system. SPMRs are used to solicit a source path message from a source in a nonimplosive way. The typical application is for late-joining receivers to solicit source path messages directly from a source in order to be able to send NAKs for missing packets, without having to wait for a regularly scheduled source path message from that source. • OTHER—Total number of other PGM packets received and sent by the local system. |
| packets shorter than minimum PGM header length | Total number of packets received with headers that are shorter than the minimum required PGM header length. |

Table 58: show pgm statistics Output Fields (*continued*)

| Field Name | Field Description |
|--|---|
| packets received with incorrect check sum | Total number of packets received with an incorrect checksum. The checksum field is the 1's complement of the 1's complement sum of the entire PGM packet, including the header. |
| packets received with zero check sum | Total number of packets received with a zero checksum. If the computed checksum is zero, it is transmitted as all ones. A value of zero in this field means that the transmitter generated no checksum. |
| packets received with TSDU length incorrect | Total number of packets received with an incorrect Transport Service Data Unit (TSDU) length (16 bits). |
| packets received with SPM length incorrect | Total number of packets received with an incorrect source path message length. |
| packets received with unknown SPM address family | Total number of packets received with an unknown source path message address family indicator (AFI). |
| packets received with NAK length incorrect | Total number of packets received with an incorrect NAK length. |
| packets received with unknown NAK address family | Total number of packets received with an unknown NAK address family indicator (AFI). |
| packets received with NAK for unknown TSI | Total number of NAK packets received with an unknown transport session identifier (TSI). |
| packets received when NAK throttled | Total number of packets received when NAK is throttled. |
| packets received with NCF length incorrect | Total number of packets received with an incorrect NAK confirmation length. |
| packets received with unknown NCF address family | Total number of packets received with an unknown NAK confirmation address family indicator (AFI). |
| packets received with NCF for unknown TSI | Total number of NAK confirmation packets received with an unknown transport session identifier (TSI). |
| packets received with RDATA length incorrect | Total number of packets received with an incorrect RDATA length. |
| packets received with RDATA for unknown TSI | Total number of RDATA packets received with an unknown transport session identifier (TSI). |

Sample Output

show pgm statistics

```
user@host> show pgm statistics
PGM type      # received  # sent
SPM            0          0
POLL           0          0
POLR           0          0
ODATA          0          0
RDATA          0          0
NAK            0          0
NULLNAK        0          0
NCF            0          0
SPMR           0          0
OTHER          0          0

packets shorter than minimum PGM header length :      0
packets received with incorrect check sum       :      0
packets received with zero check sum            :      0
packets received with TSDU length incorrect     :      0
packets received with SPM length incorrect      :      0
packets received with unknown SPM address family:      0
packets received with NAK length incorrect      :      0
packets received with unknown NAK address family:      0
packets received with NAK for unknown TSI       :      0
packets received when NAK throttled             :      0
packets received with NCF length incorrect      :      0
packets received with unknown NCF address family:      0
packets received with NCF for unknown TSI       :      0
packets received with RDATA length incorrect    :      0
packets received with RDATA for unknown TSI     :      0
```

show pim bootstrap

| | |
|---|--|
| Syntax | show pim bootstrap <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switch and the QFX Series) | show pim bootstrap <instance <i>instance-name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. instance option introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | For sparse mode only, display information about Protocol Independent Multicast (PIM) bootstrap routers. |
| Options | <p>none—Display PIM bootstrap router information for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Display information about bootstrap routers for a specific PIM-enabled routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show pim bootstrap on page 231 show pim bootstrap instance on page 231 |
| Output Fields | Table 59 on page 230 describes the output fields for the show pim bootstrap command. Output fields are listed in the approximate order in which they appear. |

Table 59: show pim bootstrap Output Fields

| Field Name | Field Description |
|----------------------|---|
| Instance | Name of the routing instance. |
| BSR | Bootstrap router. |
| Pri | Priority of the routing device as elected to be the bootstrap router. |
| Local address | Local routing device address. |
| Pri | Local routing device address priority to be elected as the bootstrap router. |
| State | Local routing device election state: Candidate , Elected , or Ineligible . |

Table 59: show pim bootstrap Output Fields (*continued*)

| Field Name | Field Description |
|----------------|--|
| Timeout | How long until the local routing device declares the bootstrap router to be unreachable, in seconds. |

Sample Output

show pim bootstrap

```
user@host> show pim bootstrap
Instance: PIM.master
```

| BSR | Pri | Local address | Pri | State | Timeout |
|-------------------------|-----|-------------------------|-----|------------|---------|
| None | 0 | 10.255.71.46 | 0 | InEligible | 0 |
| feco:1:1:1:1:0:aff:785c | 34 | feco:1:1:1:1:0:aff:7c12 | 0 | InEligible | 0 |

show pim bootstrap instance

```
user@host> show pim bootstrap instance VPN-A
Instance: PIM.VPN-A
```

| BSR | Pri | Local address | Pri | State | Timeout |
|------|-----|-----------------|-----|------------|---------|
| None | 0 | 192.168.196.105 | 0 | InEligible | 0 |

show pim interfaces

| | |
|---|---|
| Syntax | show pim interfaces <inet inet6> <instance (<i>instance-name</i> all)> <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switch and the QFX Series) | show pim interfaces <inet inet6> <instance (<i>instance-name</i> all)> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. Support for bidirectional PIM added in Junos OS Release 12.1. Support for the instance all option added in Junos OS Release 12.1. |
| Description | Display information about the interfaces on which Protocol Independent Multicast (PIM) is configured. |
| Options | <p>none—Display interface information for all family addresses for the main instance.</p> <p>inet inet6—(Optional) Display interface information for IPv4 or IPv6 family addresses, respectively.</p> <p>instance (<i>instance-name</i> all)—(Optional) Display information about interfaces for a specific PIM-enabled routing instance or for all routing instances.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show pim interfaces on page 234 |
| Output Fields | Table 60 on page 232 describes the output fields for the show pim interfaces command. Output fields are listed in the approximate order in which they appear. |

Table 60: show pim interfaces Output Fields

| Field Name | Field Description |
|------------|--|
| Instance | Name of the routing instance. |
| Name | Interface name. |
| State | State of the interface. The state also is displayed in the show interfaces command. |

Table 60: show pim interfaces Output Fields (*continued*)

| Field Name | Field Description |
|---------------------|--|
| Mode | <p>PIM mode running on the interface:</p> <ul style="list-style-type: none"> • B—In bidirectional mode, multicast groups are carried across the network over bidirectional shared trees. This type of tree minimizes PIM routing state, which is especially important in networks with numerous and dispersed senders and receivers. • S—In sparse mode, routing devices must join and leave multicast groups explicitly. Upstream routing devices do not forward multicast traffic to this routing device unless this device has sent an explicit request (using a join message) to receive multicast traffic. • Dense—Unlike sparse mode, where data is forwarded only to routing devices sending an explicit request, dense mode implements a flood-and-prune mechanism, similar to DVMRP (the first multicast protocol used to support the multicast backbone). (Not supported on QFX Series.) • Sparse-Dense—Sparse-dense mode allows the interface to operate on a per-group basis in either sparse or dense mode. A group specified as dense is not mapped to a rendezvous point (RP). Instead, data packets destined for that group are forwarded using PIM-Dense Mode (PIM-DM) rules. A group specified as sparse is mapped to an RP, and data packets are forwarded using PIM-Sparse Mode (PIM-SM) rules. (Not supported on QFX Series.) <p>When sparse-dense mode is configured, the output includes both S and D. When bidirectional-sparse mode is configured, the output includes S and B. When bidirectional-sparse-dense mode is configured, the output includes B, S, and D.</p> |
| IP | Version number of the address family on the interface: 4 (IPv4) or 6 (IPv6). |
| V | PIM version running on the interface: 1 or 2. |
| State | <p>State of PIM on the interface:</p> <ul style="list-style-type: none"> • Active—Bidirectional mode is enabled on the interface and on all PIM neighbors. • DR—Designated router. • NotCap—Bidirectional mode is not enabled on the interface. This can happen when bidirectional PIM is not configured locally, when one of the neighbors is not configured for bidirectional PIM, or when one of the neighbors has not implemented the bidirectional PIM protocol. • NotDR—Not the designated router. • P2P—Point to point. |
| NbrCnt | Number of neighbors that have been seen on the interface. |
| JoinCnt(sg) | Number of (s,g) join messages that have been seen on the interface. |
| JointCnt(*g) | Number of (*g) join messages that have been seen on the interface. |
| DR address | Address of the designated router. |

Sample Output

show pim interfaces

```
user@host> show pim interfaces
```

Stat = Status, V = Version, NbrCnt = Neighbor Count,

S = Sparse, D = Dense, B = Bidirectional,

DR = Designated Router, P2P = Point-to-point link,

Active = Bidirectional is active, NotCap = Not Bidirectional Capable

| Name | Stat | Mode | IP | V | State | NbrCnt | JoinCnt(sg/*g) | DR address |
|----------------|------|------|----|---|--------------|--------|----------------|------------|
| ge-0/3/0.0 | Up | S | 4 | 2 | NotDR,NotCap | 1 | 0/0 | 40.0.0.3 |
| ge-0/3/3.50 | Up | S | 4 | 2 | DR,NotCap | 1 | 9901/100 | 50.0.0.2 |
| ge-0/3/3.51 | Up | S | 4 | 2 | DR,NotCap | 1 | 0/0 | 51.0.0.2 |
| pe-1/2/0.32769 | Up | S | 4 | 2 | P2P,NotCap | 0 | 0/0 | |

show pim join

| | |
|---|--|
| Syntax | <pre>show pim join <brief detail extensive summary> <inet inet6> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <range></pre> |
| Syntax (EX Series Switch and the QFX Series) | <pre>show pim join <brief detail extensive summary> <inet inet6> <instance <i>instance-name</i>> <range></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>summary option introduced in Junos OS Release 9.6.</p> <p>inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Support for bidirectional PIM added in Junos OS Release 12.1.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> |
| Description | <p>Display information about Protocol Independent Multicast (PIM) groups for all PIM modes.</p> <p>For bidirectional PIM, display information about PIM group ranges (*G-range) for each active bidirectional RP group range, in addition to each of the joined (*G) routes.</p> |
| Options | <p>none—Display the standard information about PIM groups for all supported family addresses for all routing instances.</p> <p>brief detail extensive summary—(Optional) Display the specified level of output.</p> <p>inet inet6—(Optional) Display PIM group information for IPv4 or IPv6 family addresses, respectively.</p> <p>instance <i>instance-name</i>—(Optional) Display information about groups for the specified PIM-enabled routing instance only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>range—(Optional) Address range of the group, specified as <i>prefix/prefix-length</i>.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • clear pim join on page 113 |
| List of Sample Output | <p>show pim join summary on page 239</p> <p>show pim join (PIM Sparse Mode) on page 239</p> |

[show pim join \(Bidirectional PIM\) on page 239](#)
[show pim join instance <instance-name> on page 240](#)
[show pim join detail on page 240](#)
[show pim join extensive \(PIM Sparse Mode\) on page 240](#)
[show pim join extensive \(Bidirectional PIM\) on page 241](#)
[show pim join extensive \(Bidirectional PIM with a Directly Connected Phantom RP\) on page 242](#)
[show pim join instance <instance-name> extensive on page 243](#)

Output Fields [Table 61 on page 236](#) describes the output fields for the **show pim join** command. Output fields are listed in the approximate order in which they appear.

Table 61: show pim join Output Fields

| Field Name | Field Description | Level of Output |
|--|--|--|
| Instance | Name of the routing instance. | brief detail extensive summary none |
| Family | Name of the address family: inet (IPv4) or inet6 (IPv6). | brief detail extensive summary none |
| Route type | Type of multicast route: (S,G) or (*G). | summary |
| Route count | Number of (S,G) routes and number of (*G) routes. | summary |
| R | Rendezvous Point Tree. | brief detail extensive none |
| S | Sparse. | brief detail extensive none |
| W | Wildcard. | brief detail extensive none |
| Group | Group address. | brief detail extensive none |
| Bidirectional group prefix length | For bidirectional PIM, length of the IP prefix for RP group ranges. | All levels |
| Source | Multicast source: <ul style="list-style-type: none"> • * (wildcard value) • <i>ipv4-address</i> • <i>ipv6-address</i> | brief detail extensive none |
| RP | Rendezvous point for the PIM group. | brief detail extensive none |

Table 61: show pim join Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------------|--|------------------------------------|
| Flags | PIM flags: <ul style="list-style-type: none"> • bidirectional—Bidirectional mode entry. • dense—Dense mode entry. • rptree—Entry is on the rendezvous point tree. • sparse—Sparse mode entry. • spt—Entry is on the shortest-path tree for the source. • wildcard—Entry is on the shared tree. | brief detail extensive none |
| Upstream interface | RPF interface toward the source address for the source-specific state (S,G) or toward the rendezvous point (RP) address for the non-source-specific state (*,G). For bidirectional PIM, RP Link means that the interface is directly connected to a subnet that contains a phantom RP address. | brief detail extensive none |
| Upstream neighbor | Information about the upstream neighbor: Direct , Local , Unknown , or a specific IP address. For bidirectional PIM, Direct means that the interface is directly connected to a subnet that contains a phantom RP address. | extensive |
| Upstream state | Information about the upstream interface: <ul style="list-style-type: none"> • Join to RP—Sending a join to the rendezvous point. • Join to Source—Sending a join to the source. • Local RP—Sending neither join messages nor prune messages toward the RP, because this router is the rendezvous point. • Local Source—Sending neither join messages nor prune messages toward the source, because the source is locally attached to this routing device. • Prune to RP—Sending a prune to the rendezvous point. • Prune to Source—Sending a prune to the source. <p>NOTE: RP group range entries have None in the Upstream state field because RP group ranges do not trigger actual PIM join messages between routers.</p> | extensive |

Table 61: show pim join Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---|---|------------------|
| Downstream neighbors | <p>Information about downstream interfaces:</p> <ul style="list-style-type: none"> • Interface—Interface name for the downstream neighbor. <p>NOTE: A pseudo PIM-SM interface appears for all IGMP-only interfaces.</p> <ul style="list-style-type: none"> • Interface address—Address of the downstream neighbor. • State—Information about the downstream neighbor: join or prune. • Flags—PIM join flags: R (RPtree), S (Sparse), W (Wildcard), or zero. • Uptime—Time since the downstream interface joined the group. • Time since last Join—Time since the last join message was received from the downstream interface. • Time since last Prune—Time since the last prune message was received from the downstream interface. | extensive |
| Number of downstream interfaces | Total number of outgoing interfaces for each (S,G) entry. | extensive |
| Assert Timeout | Length of time between assert cycles on the downstream interface. Not displayed if the assert timer is null. | extensive |
| Keepalive timeout | Time remaining until the downstream join state is updated (in seconds). If the downstream join state is not updated before this keepalive timer reaches zero, the entry is deleted. If there is a directly connected host, Keepalive timeout is Infinity . | extensive |
| Uptime | Time since the creation of (S,G) or (*,G) state. The uptime is not refreshed every time a PIM join message is received for an existing (S,G) or (*,G) state. | extensive |
| Bidirectional accepting interfaces | <p>Interfaces on the router that forward bidirectional PIM traffic.</p> <p>The reasons for forwarding bidirectional PIM traffic are that the interface is the winner of the designated forwarder election (DF Winner), or the interface is the reverse path forwarding (RPF) interface toward the RP (RPF).</p> | extensive |

Sample Output

**show pim join
summary**

```
user@host> show pim join summary
Instance: PIM.master Family: INET

Route type      Route count
(S,g)           2
(*,g)           1

Instance: PIM.master Family: INET6
```

**show pim join (PIM
Sparse Mode)**

```
user@host> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
  Source: *
  RP: 10.255.14.144
  Flags: sparse,rptree,wildcard
  Upstream interface: Local

Group: 239.1.1.1
  Source: 10.255.14.144
  Flags: sparse,spt
  Upstream interface: Local

Group: 239.1.1.1
  Source: 10.255.70.15
  Flags: sparse,spt
  Upstream interface: so-1/0/0.0

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

**show pim join
(Bidirectional PIM)**

```
user@host> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.13.2
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0

Group: 224.1.3.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.1.3
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0 (RP Link)

Group: 225.1.1.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.13.2
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0
```

```
Group: 225.1.3.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.1.3
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0 (RP Link)

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join instance **<instance-name>**

```
user@host> show pim join instance VPN-A
Instance: PIM.VPN-A Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 235.1.1.2
  Source: *
  RP: 10.10.47.100
  Flags: sparse,rptree,wildcard
  Upstream interface: Local

Group: 235.1.1.2
  Source: 192.168.195.74
  Flags: sparse,spt
  Upstream interface: at-0/3/1.0

Group: 235.1.1.2
  Source: 192.168.195.169
  Flags: sparse
  Upstream interface: so-1/0/1.0

Instance: PIM.VPN-A Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join detail

```
user@host> show pim join detail
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
  Source: *
  RP: 10.255.14.144
  Flags: sparse,rptree,wildcard
  Upstream interface: Local

Group: 239.1.1.1
  Source: 10.255.14.144
  Flags: sparse,spt
  Upstream interface: Local

Group: 239.1.1.1
  Source: 10.255.70.15
  Flags: sparse,spt
  Upstream interface: so-1/0/0.0

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
```

extensive (PIM Sparse Mode)

R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```

Group: 239.1.1.1
Source: *
RP: 10.255.14.144
Flags: sparse,rptree,wildcard
Upstream interface: Local
Upstream neighbor: Local
Upstream state: Local RP
Uptime: 00:03:49
Downstream neighbors:
  Interface: so-1/0/0.0
    10.111.10.2 State: Join Flags: SRW Timeout: 174
    Uptime: 00:03:49 Time since last Join: 00:01:49
  Interface: mt-1/1/0.32768
    10.10.47.100 State: Join Flags: SRW Timeout: Infinity
    Uptime: 00:03:49 Time since last Join: 00:01:49
Number of downstream interfaces: 2

```

```

Group: 239.1.1.1
Source: 10.255.14.144
Flags: sparse,spt
Upstream interface: Local
Upstream neighbor: Local
Upstream state: Local Source, Local RP
Keepalive timeout: 344
Uptime: 00:03:49
Downstream neighbors:
  Interface: so-1/0/0.0
    10.111.10.2 State: Join Flags: S Timeout: 174
    Uptime: 00:03:49 Time since last Prune: 00:01:49
  Interface: mt-1/1/0.32768
    10.10.47.100 State: Join Flags: S Timeout: Infinity
    Uptime: 00:03:49 Time since last Prune: 00:01:49
Number of downstream interfaces: 2

```

```

Group: 239.1.1.1
Source: 10.255.70.15
Flags: sparse,spt
Upstream interface: so-1/0/0.0
Upstream neighbor: 10.111.10.2
Upstream state: Local RP, Join to Source
Keepalive timeout: 344
Uptime: 00:03:49
Downstream neighbors:
  Interface: Pseudo-GMP
    fe-0/0/0.0 fe-0/0/1.0 fe-0/0/3.0
  Interface: so-1/0/0.0 (pruned)
    10.111.10.2 State: Prune Flags: SR Timeout: 174
    Uptime: 00:03:49 Time since last Prune: 00:01:49
  Interface: mt-1/1/0.32768
    10.10.47.100 State: Join Flags: S Timeout: Infinity
    Uptime: 00:03:49 Time since last Prune: 00:01:49
Number of downstream interfaces: 3

```

Instance: PIM.master Family: INET6

R = Rendezvous Point Tree, S = Sparse, W = Wildcard

show pim join extensive

user@host> show pim join extensive

Instance: PIM.master Family: INET

R = Rendezvous Point Tree, S = Sparse, W = Wildcard

(Bidirectional PIM)

```

Group: 224.1.1.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.13.2
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0
  Upstream neighbor: 10.10.1.2
  Upstream state: None
  Uptime: 00:03:49
  Bidirectional accepting interfaces:
    Interface: ge-0/0/1.0      (RPF)
    Interface: lo0.0          (DF Winner)
  Number of downstream interfaces: 0

Group: 225.1.1.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.13.2
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0
  Upstream neighbor: 10.10.1.2
  Upstream state: None
  Uptime: 00:03:49
  Bidirectional accepting interfaces:
    Interface: ge-0/0/1.0      (RPF)
    Interface: lo0.0          (DF Winner)
  Downstream neighbors:
    Interface: lt-1/0/10.24
      10.0.24.4 State: Join   RW   Timeout: 185
    Interface: lt-1/0/10.23
      10.0.23.3 State: Join   RW   Timeout: 184
  Number of downstream interfaces: 2

Group: 225.1.3.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.1.3
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0 (RP Link)
  Upstream neighbor: Direct
  Upstream state: Local RP
  Uptime: 00:03:49
  Bidirectional accepting interfaces:
    Interface: ge-0/0/1.0      (RPF)
    Interface: lo0.0          (DF Winner)
    Interface: xe-4/1/0.0      (DF Winner)
  Number of downstream interfaces: 0

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```

**show pim join
extensive
(Bidirectional PIM with**

```

user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```

a Directly Connected Phantom RP)

```

Group: 224.1.3.0
Bidirectional group prefix length: 24
Source: *
RP: 10.10.1.3
Flags: bidirectional,rptree,wildcard
Upstream interface: ge-0/0/1.0 (RP Link)
Upstream neighbor: Direct
Upstream state: Local RP
Uptime: 00:03:49
Bidirectional accepting interfaces:
  Interface: ge-0/0/1.0      (RPF)
  Interface: lo0.0          (DF Winner)
  Interface: xe-4/1/0.0     (DF Winner)
Number of downstream interfaces: 0

```

show pim join instance <instance-name> extensive

```

user@host> show pim join instance VPN-A extensive
Instance: PIM.VPN-A Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```

```

Group: 235.1.1.2
Source: *
RP: 10.10.47.100
Flags: sparse,rptree,wildcard
Upstream interface: Local
Upstream neighbor: Local
Upstream state: Local RP
Uptime: 00:03:49
Downstream neighbors:
  Interface: mt-1/1/0.32768
    10.10.47.101 State: Join Flags: SRW Timeout: 156
    Uptime: 00:03:49 Time since last Join: 00:01:49
Number of downstream interfaces: 1

```

```

Group: 235.1.1.2
Source: 192.168.195.74
Flags: sparse,spt
Upstream interface: at-0/3/1.0
Upstream neighbor: 10.111.30.2
Upstream state: Local RP, Join to Source
Keepalive timeout: 156
Uptime: 00:14:52

```

```

Group: 235.1.1.2
Source: 192.168.195.169
Flags: sparse
Upstream interface: so-1/0/1.0
Upstream neighbor: 10.111.20.2
Upstream state: Local RP, Join to Source
Keepalive timeout: 156
Uptime: 00:14:52

```

show pim mdt

| | |
|---------------------------------|---|
| Syntax | <pre>show pim mdt instance <i>instance-name</i> <brief detail extensive> <incoming outgoing> <logical-system (all logical-system-name)> <range></pre> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display information about Protocol Independent Multicast (PIM) default multicast distribution tree (MDT) and the data MDTs in a Layer 3 VPN environment for a routing instance. |
| Options | <p>instance <i>instance-name</i>—Display information about data-MDTs for a specific PIM-enabled routing instance.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>incoming outgoing—(Optional) Display incoming or outgoing multicast data tunnels, respectively.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>range—(Optional) Display information about an IP address with optional prefix length representing a particular multicast group.</p> |
| Required Privilege Level | view |
| List of Sample Output | show pim mdt instance on page 246 show pim mdt instance detail on page 246 show pim mdt instance extensive on page 246 show pim mdt instance incoming on page 246 show pim mdt instance outgoing on page 247 show pim mdt instance (SSM Mode) on page 247 |
| Output Fields | Table 62 on page 244 describes the output fields for the show pim mdt command. Output fields are listed in the approximate order in which they appear. |

Table 62: show pim mdt Output Fields

| Field Name | Field Description | Level of Output |
|------------------|---|-----------------|
| Instance | Name of the routing instance. | All levels |
| Tunnel direction | Direction the tunnel faces, from the router's perspective: Outgoing or Incoming . | All levels |
| Tunnel mode | Mode the tunnel is operating in: PIM-SSM or PIM-ASM . | All levels |

Table 62: show pim mdt Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-----------------------------------|--|-----------------|
| Default group address | Default multicast group address using this tunnel. | All levels |
| Default source address | Default multicast source address using this tunnel. | All levels |
| Default tunnel interface | Default multicast tunnel interface. | All levels |
| Default tunnel source | Address used as the source address for outgoing PIM control messages. | All levels |
| C-Group | Customer-facing multicast group address using this tunnel. If you enable dynamic reuse of data MDT group addresses, more than one group address can use the same data MDT. | detail |
| C-Source | IP address of the multicast source in the customer's address space. If you enable dynamic reuse of data MDT group addresses, more than one source address can use the same data MDT. | detail |
| P-Group | Service provider-facing multicast group address using this tunnel. | detail |
| Data tunnel interface | Multicast data tunnel interface that set up the data-MDT tunnel. | detail |
| Last known forwarding rate | Last known rate, in kilobits per second, at which the tunnel was forwarding traffic. | detail |
| Configured threshold rate | Rate, in kilobits per second, above which a data-MDT tunnel is created and below which it is deleted. | detail |
| Tunnel uptime | Time that this data-MDT tunnel has existed. The format is <i>hours:minutes:seconds</i> . | detail |

Sample Output

```

show pim mdt instance user@host> show pim mdt instance VPN-A
Instance: PIM.VPN-A
Tunnel direction: Outgoing
Default group address: 239.1.1.1
Default tunnel interface: mt-1/1/0.32768
Default tunnel source: 192.168.7.1

C-group address   C-source address   P-group address   Data tunnel interface
235.1.1.2         192.168.195.74     228.0.0.0         mt-1/1/0.32769

Instance: PIM.VPN-A
Tunnel direction: Incoming
Default group address: 239.1.1.1
Default tunnel interface: mt-1/1/0.49152

```

```

show pim mdt instance user@host> show pim mdt instance VPN-A detail
detail
Instance: PIM.VPN-A
Tunnel direction: Outgoing
Default group address: 239.1.1.1
Default tunnel interface: mt-1/1/0.32768
Default tunnel source: 192.168.7.1

C-Group: 235.1.1.2
  C-Source: 192.168.195.74
  P-Group : 228.0.0.0
  Data tunnel interface      : mt-1/1/0.32769
  Last known forwarding rate : 48 kbps (6 kbps)
  Configured threshold rate  : 10 kbps
  Tunnel uptime              : 00:00:34

Instance: PIM.VPN-A
Tunnel direction: Incoming
Default group address: 239.1.1.1
Default tunnel interface: mt-1/1/0.49152

```

```

show pim mdt instance user@host> show pim mdt instance VPN-A extensive
extensive
Instance: PIM.VPN-A
Tunnel direction: Outgoing
Default group address: 239.1.1.1
Default tunnel interface: mt-1/1/0.32768
Default tunnel source: 192.168.7.1

C-Group: 235.1.1.2
  C-Source: 192.168.195.74
  P-Group : 228.0.0.0
  Data tunnel interface      : mt-1/1/0.32769
  Last known forwarding rate : 48 kbps (6 kbps)
  Configured threshold rate  : 10 kbps
  Tunnel uptime              : 00:00:41

Instance: PIM.VPN-A
Tunnel direction: Incoming
Default group address: 239.1.1.1
Default tunnel interface: mt-1/1/0.49152

```


show pim mdt instance incoming

```
user@host> show pim mdt instance VPN-A incoming
Instance: PIM.VPN-A
Tunnel direction: Incoming
Default group address: 239.1.1.1
Default tunnel interface: mt-1/1/0.49152
```

show pim mdt instance outgoing

```
user@host> show pim mdt instance VPN-A outgoing
Instance: PIM.VPN-A
Tunnel direction: Outgoing
Default group address: 239.1.1.1
Default tunnel interface: mt-1/1/0.32768
Default tunnel source: 192.168.7.1
```

| C-group address | C-source address | P-group address | Data tunnel interface |
|-----------------|------------------|-----------------|-----------------------|
| 235.1.1.2 | 192.168.195.74 | 228.0.0.0 | mt-1/1/0.32769 |

show pim mdt instance (SSM Mode)

```
user@host> show pim mdt instance vpn-a
Instance: PIM.vpn-a
Tunnel direction: Outgoing
Tunnel mode: PIM-SSM
Default group address: 232.1.1.1
Default source address: 10.255.14.216
Default tunnel interface: mt-1/3/0.32769
Default tunnel source: 192.168.7.1
```

```
Instance: PIM.vpn-a
Tunnel direction: Incoming
Tunnel mode: PIM-SSM
Default group address: 232.1.1.1
Default source address: 10.255.14.217
Default tunnel interface: mt-1/3/0.49153
```

```
Instance: PIM.vpn-a
Tunnel direction: Incoming
Tunnel mode: PIM-SSM
Default group address: 232.1.1.1
Default source address: 10.255.14.218
Default tunnel interface: mt-1/3/0.49153
```

show pim mdt data-mdt-joins

Syntax `show pim mdt data-mdt-joins`
`<logical-system (all | logical-system-name)> instance instance-name`

Release Information Command introduced in Junos OS Release 11.2.

Description In a draft-rosen Layer 3 multicast virtual private network (MVPN) configured with service provider tunnels, display the advertisements of new multicast distribution tree (MDT) group addresses cached by the provider edge (PE) routers in the specified VPN routing and forwarding (VRF) instance that is configured to use the Protocol Independent Multicast (PIM) protocol.

Options `instance instance-name`—Display data MDT join packets cached by PE routers in a specific PIM instance.

`logical-system (all | logical-system-name)`—(Optional) Perform this operation on all logical systems or on a particular logical system.



NOTE: Draft-rosen multicast VPNs are not supported in a logical system environment even though the configuration statements can be configured under the logical-systems hierarchy.

Required Privilege Level view

Related Documentation

- Understanding Data MDTs
- Example: Configuring Data MDTs and Provider Tunnels Operating in Source-Specific Multicast Mode
- Example: Configuring Data MDTs and Provider Tunnels Operating in Any-Source Multicast Mode

List of Sample Output [show pim mdt data-mdt-joins on page 249](#)

Output Fields [Table 63 on page 248](#) describes the output fields for the `show pim mdt data-mdt-joins` command. Output fields are listed in the approximate order in which they appear.

Table 63: show pim mdt data-mdt-joins Output Fields

| Field Name | Field Description |
|-----------------|---|
| C-Group | IPv4 group address in the address space of the customer's VPN-specific PIM-enabled routing instance of the multicast traffic destination. This 32-bit value is carried in the C-group field of the MDT join TLV packet. |
| C-Source | IPv4 address in the address space of the customer's VPN-specific PIM-enabled routing instance of the multicast traffic source. This 32-bit value is carried in the C-source field of the MDT join TLV packet. |

Table 63: show pim mdt data-mdt-joins Output Fields (*continued*)

| Field Name | Field Description |
|-----------------|---|
| P-Group | IPv4 group address in the service provider's address space of the new data MDT that the PE router will use to encapsulate the VPN multicast traffic flow (C-Source, C-Group). This 32-bit value is carried in the P-group field of the MDT join TLV packet. |
| P-Source | IPv4 address of the PE router. |
| Timeout | Timeout, in seconds, remaining for this cache entry. When the cache entry is created, this field is set to 180 seconds. After an entry times out, the PE router deletes the entry from its cache and prunes itself off the data MDT. |

Sample Output

```

show pim mdt data-mdt-joins      user@host show pim mdt data-mdt-joins instance VPN-A
C-Source      C-Group      P-Source      P-Group      Timeout
20.2.15.9     225.1.1.2     20.0.0.5     239.10.10.0  172
20.2.15.9     225.1.1.3     20.0.0.5     239.10.10.1  172

```

show pim mdt data-mdt-limit

Syntax `show pim mdt data-mdt-limit instance instance-name`
`<logical-system (all | logical-system-name)>`

Release Information Command introduced in Junos OS Release 12.2.

Description Display the maximum number configured and the currently active data multicast distribution trees (MDTs) for a specific VPN routing and forwarding (VRF) instance.

Options `instance instance-name`—Display data MDT information for the specified VRF instance.

`logical-system (all | logical-system-name)`—(Optional) Perform this operation on all logical systems or on a particular logical system.



NOTE: Draft-rosen multicast VPNs are not supported in a logical system environment even though the configuration statements can be configured under the logical-systems hierarchy.

Required Privilege Level view

Related Documentation

- Understanding Data MDTs
- Example: Configuring Data MDTs and Provider Tunnels Operating in Source-Specific Multicast Mode
- Example: Configuring Data MDTs and Provider Tunnels Operating in Any-Source Multicast Mode

List of Sample Output [show pim mdt data-mdt-limit on page 251](#)

Output Fields [Table 64 on page 250](#) describes the output fields for the `show pim mdt data-mdt-limit` command. Output fields are listed in the approximate order in which they appear.

Table 64: show pim mdt data-mdt-limit Output Fields

| Field Name | Field Description |
|----------------------|---|
| Maximum Data Tunnels | Maximum number of data MDTs created in this VRF instance. If the number is 0, no data MDTs are created for this VRF instance. |
| Active Data Tunnels | Active number of data MDTs in this VRF instance. |

Sample Output

`show pim mdt
data-mdt-limit`

```
user@host: show pim mdt data-mdt-limit instance VPN-A
Maximum Data Tunnels           10
Active Data Tunnels             2
```

show pim mvpn

| | |
|---------------------------------|--|
| Syntax | show pim mvpn <logical-system (all logical-system-name) > |
| Release Information | Command introduced in Junos OS Release 9.4. |
| Description | Display information about multicast virtual private network (MVPN) instances. |
| Options | logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | view |
| List of Sample Output | show pim mvpn on page 252 |
| Output Fields | Table 65 on page 252 describes the output fields for the show pim mvpn command. Output fields are listed in the approximate order in which they appear. |

Table 65: show pim mvpn Output Fields

| Field Name | Field Description | Level of Output |
|------------------|---|-----------------|
| Instance | Name of the routing instance. | All levels |
| VPN-Group | Multicast group address configured for the default multicast distribution tree. | All levels |
| Mode | Mode the tunnel is operating in: PIM-MVPN , NGEN-MVPN , NGEN-TRANSITION or None . | All levels |
| Tunnel | Type of tunnel: PIM-SSM , PIM-SM , NGEN PMSI , or None (VRF-only). If NGEN-PMSI is displayed, enter the show mvpn instance command for more information. | All levels |

Sample Output

```

show pim mvpn
user@host> show pim mvpn
Instance      VPN-Group      Mode      Tunnel
PIM.ce1       232.1.1.1      PIM-MVPN  PIM-SSM

```

show pim neighbors

| | |
|---|--|
| Syntax | <pre>show pim neighbors <brief detail> <inet inet6> <instance (instance-name all)> <logical-system (all logical-system-name)></pre> |
| Syntax (EX Series Switch and the QFX Series) | <pre>show pim neighbors <brief detail> <inet inet6> <instance (instance-name all)></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional PIM added in Junos OS Release 12.1.</p> <p>Support for the instance all option added in Junos OS Release 12.1.</p> |
| Description | Display information about Protocol Independent Multicast (PIM) neighbors. |
| Options | <p>none—(Same as brief) Display standard information about PIM neighbors for all supported family addresses for the main instance.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>inet inet6—(Optional) Display information about PIM neighbors for IPv4 or IPv6 family addresses, respectively.</p> <p>instance (instance-name all)—(Optional) Display information about neighbors for the specified PIM-enabled routing instance or for all routing instances.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | <p>show pim neighbors on page 256</p> <p>show pim neighbors brief on page 256</p> <p>show pim neighbors instance on page 256</p> <p>show pim neighbors detail on page 256</p> <p>show pim neighbors detail (With BFD) on page 257</p> |
| Output Fields | <p>Table 66 on page 254 describes the output fields for the show pim neighbors command. Output fields are listed in the approximate order in which they appear.</p> |

Table 66: show pim neighbors Output Fields

| Field Name | Field Description | Level of Output |
|--|--|-------------------|
| Instance | Name of the routing instance. | All levels |
| Interface | Interface through which the neighbor is reachable. | All levels |
| Neighbor addr | Address of the neighboring PIM routing device. | All levels |
| IP | IP version: 4 or 6. | All levels |
| V | PIM version running on the neighbor: 1 or 2. | All levels |
| Mode | PIM mode of the neighbor: Sparse , Dense , SparseDense , or Unknown . When the neighbor is running PIM version 2, this mode is always Unknown . | All levels |
| Option | Can be one or more of the following: <ul style="list-style-type: none"> • B—Bidirectional Capable. • H—Hello Option Holdtime. • G—Generation Identifier. • P—Hello Option DR Priority. • L—Hello Option LAN Prune Delay. | brief none |
| Uptime | Time the neighbor has been operational since the PIM process was last initialized, in the format dd:hh:mm:ss ago for less than a week and nwnd:hh:mm:ss ago for more than a week. | All levels |
| Address | Address of the neighboring PIM router. | detail |
| BFD | Status and operational state of the Bidirectional Forwarding Detection (BFD) protocol on the interface: Enabled , Operational state is up , or Disabled . | detail |
| Hello Option Holdtime | Time for which the neighbor is available, in seconds. The range of values is 0 through 65,535. | detail |
| Hello Default Holdtime | Default holdtime and the time remaining if the holdtime option is not in the received hello message. | detail |
| Hello Option DR Priority | Designated router election priority. The range of values is 0 through 255. | detail |
| Hello Option Generation ID | 9-digit or 10-digit number used to tag hello messages. | detail |
| Hello Option Bi-Directional PIM supported | Neighbor can process bidirectional PIM messages. | detail |
| Hello Option LAN Prune Delay | Time to wait before the neighbor receives prune messages, in the format delay nnn ms override nnnn ms . | detail |

Table 66: show pim neighbors Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|----------------------------|---|-----------------|
| Join Suppression supported | Neighbor is capable of join suppression. | detail |
| Rx Join | Information about joins received from the neighbor. <ul style="list-style-type: none">• Group—Group addresses in the join message.• Source—Address of the source in the join message.• Timeout—Time for which the join is valid. | detail |

Sample Output

show pim neighbors

```
user@host> show pim neighbors
Instance: PIM.master
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority

Interface          IP V Mode      Option      Uptime Neighbor addr
so-1/0/0.0         4 2           HPLG        00:07:10 10.111.10.2
```

show pim neighbors brief

The output for the **show pim neighbors brief** command is identical to that for the **show pim neighbors** command. For sample output, see [show pim neighbors on page 256](#).

show pim neighbors instance

```
user@host> show pim neighbors instance VPN-A
Instance: PIM.VPN-A
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority

Interface          IP V Mode      Option      Uptime Neighbor addr
at-0/3/1.0         4 2           HPLG        00:07:54 10.111.30.2
mt-1/1/0.32768     4 2           HPLG        00:07:22 10.10.47.101
so-1/0/1.0         4 2           HPLG        00:07:50 10.111.20.2
```

show pim neighbors detail

```
user@host> show pim neighbors detail
Instance: PIM.master
Interface: ge-0/0/1.0

    Address: 10.10.1.1, IPv4, PIM v2, Mode: SparseDense, sg Join Count: 0, ts
Join Count: 2
    Hello Option Holdtime: 65535 seconds
    Hello Option DR Priority: 1
    Hello Option Generation ID: 2053759302
    Hello Option Bi-Directional PIM supported
    Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
                                Join Suppression supported

    Address: 10.10.1.2, IPv4, PIM v2, sg Join Count: 0, ts
Join Count: 2
    BFD: Disabled
    Hello Option Holdtime: 105 seconds 93 remaining
    Hello Option DR Priority: 1
    Hello Option Generation ID: 1734018161
    Hello Option Bi-Directional PIM supported
    Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
                                Join Suppression supported

Interface: lo0.0

    Address: 10.255.179.246, IPv4, PIM v2, Mode: SparseDense, sg Join Count:
0, ts
Join Count: 0
    Hello Option Holdtime: 65535 seconds
    Hello Option DR Priority: 1
    Hello Option Generation ID: 1997462267
    Hello Option Bi-Directional PIM supported
    Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
                                Join Suppression supported
```

**show pim neighbors
detail (With BFD)**

```
user@host> show pim neighbors detail
Instance: PIM.master
Interface: fe-1/0/0.0
  Address: 192.168.11.1,      IPv4, PIM v2, Mode: Sparse
    Hello Option Holdtime: 65535 seconds
    Hello Option DR Priority: 1
    Hello Option Generation ID: 836607909
    Hello Option LAN Prune Delay: delay 500 ms override 2000 ms

  Address: 192.168.11.2,      IPv4, PIM v2
    BFD: Enabled, Operational state is up
    Hello Default Holdtime: 105 seconds 104 remaining
    Hello Option DR Priority: 1
    Hello Option Generation ID: 1907549685
    Hello Option LAN Prune Delay: delay 500 ms override 2000 ms

Interface: fe-1/0/1.0
  Address: 192.168.12.1,      IPv4, PIM v2
    BFD: Disabled
    Hello Default Holdtime: 105 seconds 80 remaining
    Hello Option DR Priority: 1
    Hello Option Generation ID: 1971554705
    Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

show pim rps

| | |
|---|---|
| Syntax | <code>show pim rps</code> <code><brief detail extensive></code> <code><group-address></code> <code><inet inet6></code> <code><instance instance-name></code> <code><logical-system (all logical-system-name)></code> |
| Syntax (EX Series Switch and the QFX Series) | <code>show pim rps</code> <code><brief detail extensive></code> <code><group-address></code> <code><inet inet6></code> <code><instance instance-name></code> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. Support for bidirectional PIM added in Junos OS Release 12.1. |
| Description | Display information about Protocol Independent Multicast (PIM) rendezvous points (RPs). |
| Options | none —Display standard information about PIM RPs for all groups and family addresses for all routing instances. brief detail extensive —(Optional) Display the specified level of output. group-address —(Optional) Display the RPs for a particular group. If you specify a group address, the output lists the routing device that is the RP for that group. inet inet6 —(Optional) Display information for IPv4 or IPv6 family addresses, respectively. instance instance-name —(Optional) Display information about RPs for a specific PIM-enabled routing instance. logical-system (all logical-system-name) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">• Example: Configuring Bidirectional PIM |
| List of Sample Output | show pim rps on page 261 show pim rps brief on page 261 show pim rps <group-address> (Bidirectional PIM) on page 261 show pim rps <group-address> (PIM Dense Mode) on page 261 |

[show pim rps <group-address> \(SSM Range Without asm-override-ssm Configured\) on page 261](#)
[show pim rps <group-address> \(SSM Range With asm-override-ssm Configured and a Sparse-Mode RP\) on page 262](#)
[show pim rps <group-address> \(SSM Range With asm-override-ssm Configured and a Bidirectional RP\) on page 262](#)
[show pim rps instance on page 263](#)
[show pim rps extensive \(PIM Sparse Mode\) on page 263](#)
[show pim rps extensive \(Bidirectional PIM\) on page 263](#)
[show pim rps extensive \(PIM Anycast RP in Use\) on page 263](#)

Output Fields Table 67 on page 259 describes the output fields for the **show pim rps** command. Output fields are listed in the approximate order in which they appear.

Table 67: show pim rps Output Fields

| Field Name | Field Description | Level of Output |
|---------------------------------|--|-------------------------|
| Instance | Name of the routing instance. | All levels |
| Family or Address family | Name of the address family: inet (IPv4) or inet6 (IPv6). | All levels |
| RP address | Address of the rendezvous point. | All levels |
| Type | Type of RP: <ul style="list-style-type: none"> auto-rp—Address of the RP known through the Auto-RP protocol. bootstrap—Address of the RP known through the bootstrap router protocol (BSR). embedded—Address of the RP known through an embedded RP (IPv6). static—Address of RP known through static configuration. | brief none |
| Holdtime | How long to keep the RP active, with time remaining, in seconds. | All levels |
| Timeout | How long until the local routing device determines the RP to be unreachable, in seconds. | All levels |
| Groups | Number of groups currently using this RP. | All levels |
| Group prefixes | Addresses of groups that this RP can span. | brief none |
| Learned via | Address and method by which the RP was learned. | detail extensive |
| Mode | The PIM mode of the RP: bidirectional or sparse. If a sparse and bidirectional RPs are configured with the same RP address, they appear as separate entries in both formats. | All levels |
| Time Active | How long the RP has been active, in the format <i>hh:mm:ss</i> . | detail extensive |

Table 67: show pim rps Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------------------------|--|---|
| Device Index | Index value of the order in which Junos OS finds and initializes the interface. For bidirectional RPs, the Device Index output field is omitted because bidirectional RPs do not require encapsulation and de-encapsulation interfaces. | detail extensive |
| Subunit | Logical unit number of the interface. For bidirectional RPs, the Subunit output field is omitted because bidirectional RPs do not require encapsulation and de-encapsulation interfaces. | detail extensive |
| Interface | Either the encapsulation or the de-encapsulation logical interface, depending on whether this routing device is a designated router (DR) facing an RP router, or is the local RP, respectively. For bidirectional RPs, the Interface output field is omitted because bidirectional RPs do not require encapsulation and de-encapsulation interfaces. | detail extensive |
| Group Ranges | Addresses of groups that this RP spans. | detail extensive <i>group-address</i> |
| Active groups using RP | Number of groups currently using this RP. | detail extensive |
| total | Total number of active groups for this RP. | detail extensive |
| Register State for RP | Current register state for each group: <ul style="list-style-type: none"> • Group—Multicast group address. • Source—Multicast source address for which the PIM register is sent or received, depending on whether this router is a designated router facing an RP router, or is the local RP, respectively: • First Hop—PIM-designated routing device that sent the Register message (the source address in the IP header). • RP Address—RP to which the Register message was sent (the destination address in the IP header). • State: On the designated router: <ul style="list-style-type: none"> • Send—Sending Register messages. • Probe—Sent a null register. If a Register-Stop message does not arrive in 5 seconds, the designated router resumes sending Register messages. • Suppress—Received a Register-Stop message. The designated router is waiting for the timer to resume before changing to Probe state. • On the RP: <ul style="list-style-type: none"> • Receive—Receiving Register messages. | extensive |
| Anycast-PIM rpset | If anycast RP is configured, the addresses of the RPs in the set. | extensive |
| Anycast-PIM local address used | If anycast RP is configured, the local address used by the RP. | extensive |

Table 67: show pim rps Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-----------------------------------|--|----------------------|
| Anycast-PIM Register State | <p>If anycast RP is configured, the current register state for each group:</p> <ul style="list-style-type: none"> • Group—Multicast group address. • Source—Multicast source address for which the PIM register is sent or received, depending on whether this routing device is a designated router facing an RP router, or is the local RP, respectively. • Origin—How the information was obtained: <ul style="list-style-type: none"> • DIRECT—From a local attachment • MSDP—From the Multicast Source Discovery Protocol (MSDP) • DR—From the designated router | extensive |
| RP selected | For sparse mode and bidirectional mode, the identity of the RP for the specified group address. | <i>group-address</i> |

Sample Output

show pim rps

```

user@host> show pim rps
Instance: PIM.master
Address family INET
RP address      Type      Mode      Holdtime Timeout Groups  Group prefixes
10.10.1.3       static    bidir     150      None     2  224.1.3.0/24
                225.1.3.0/24
10.10.13.2      static    bidir     150      None     2  224.1.1.0/24
                225.1.1.0/24

```

show pim rps brief

The output for the **show pim rps brief** command is identical to that for the **show pim rps** command. For sample output, see [show pim rps on page 261](#).

show pim rps <group-address> (Bidirectional PIM)

```

user@host> show pim rps 224.1.1.1
Instance: PIM.master

224.1.0.0/16
  11.4.12.75 (Bidirectional)

RP selected: 11.4.12.75

```

show pim rps <group-address> (PIM Dense Mode)

```

user@host> show pim rps 224.1.1.1
Instance: PIM.master

Dense Mode active for group 224.1.1.1

```

show pim rps <group-address> (SSM Range Without

```

user@host> show pim rps 224.1.1.1
Instance: PIM.master

Source-specific Mode (SSM) active for group 224.1.1.1

```

**asm-override-ssm
Configured)****show pim rps
<group-address>
(SSM Range With
asm-override-ssm
Configured and a
Sparse-Mode RP)****user@host> show pim rps 224.1.1.1**

Instance: PIM.master

Source-specific Mode (SSM) active with Sparse Mode ASM override for group 224.1.1.1

224.1.0.0/16

11.4.12.75

RP selected: 11.4.12.75

**show pim rps
<group-address>
(SSM Range With
asm-override-ssm****user@host> show pim rps 224.1.1.1**

Instance: PIM.master

Source-specific Mode (SSM) active with Sparse Mode ASM override for group 224.1.1.1

Configured and a Bidirectional RP)

224.1.0.0/16
11.4.12.75 (Bidirectional)

RP selected: (null)

show pim rps instance

user@host> show pim rps instance VPN-A

Instance: PIM.VPN-A

Address family INET

| RP address | Type | Holdtime | Timeout | Groups | Group prefixes |
|--------------|--------|----------|---------|--------|----------------|
| 10.10.47.100 | static | 0 | None | 1 | 224.0.0.0/4 |

Address family INET6

show pim rps extensive (PIM Sparse Mode)

user@host> show pim rps extensive

Instance: PIM.master

Family: INET

RP: 10.255.245.91

Learned via: static configuration

Time Active: 00:05:48

Holdtime: 45 with 36 remaining

Device Index: 122

Subunit: 32768

Interface: pd-6/0/0.32768

Group Ranges:

224.0.0.0/4, 36s remaining

Active groups using RP:

225.1.1.1

total 1 groups active

Register State for RP:

| Group | Source | FirstHop | RP Address | State | Timeout |
|-----------|----------------|---------------|---------------|---------|---------|
| 225.1.1.1 | 192.168.195.78 | 10.255.14.132 | 10.255.245.91 | Receive | 0 |

show pim rps extensive (Bidirectional PIM)

user@host> show pim rps extensive

Instance: PIM.master

Address family INET

RP: 10.10.1.3

Learned via: static configuration

Mode: Bidirectional

Time Active: 01:58:07

Holdtime: 150

Group Ranges:

224.1.3.0/24

225.1.3.0/24

RP: 10.10.13.2

Learned via: static configuration

Mode: Bidirectional

Time Active: 01:58:07

Holdtime: 150

Group Ranges:

224.1.1.0/24

225.1.1.0/24

user@host> show pim rps extensive

show pim rps extensive
(PIM Anycast RP in Use)

```

Instance: PIM.master
Family: INET
RP: 10.10.10.2
Learned via: static configuration
Time Active: 00:54:52
Holdtime: 0
Device Index: 130
Subunit: 32769
Interface: pimd.32769
Group Ranges:
    224.0.0.0/4
Active groups using RP:
    224.10.10.10

    total 1 groups active

Anycast-PIM rpset:
    10.100.111.34
    10.100.111.17
    10.100.111.55

Anycast-PIM local address used: 10.100.111.1
Anycast-PIM Register State:

```

| Group | Source | Origin |
|--------------|------------|--------|
| 224.1.1.1 | 10.10.95.2 | DIRECT |
| 224.1.1.2 | 10.10.95.2 | DIRECT |
| 224.10.10.10 | 10.10.70.1 | MSDP |
| 224.10.10.11 | 10.10.70.1 | MSDP |
| 224.20.20.1 | 10.10.71.1 | DR |

```

Address family INET6

Anycast-PIM rpset:
    ab::1
    ab::2
Anycast-PIM local address used: cd::1

Anycast-PIM Register State:

```

| Group | Source | Origin |
|---------------|--------------|--------|
| ::224.1.1.1 | ::10.10.95.2 | DIRECT |
| ::224.1.1.2 | ::10.10.95.2 | DIRECT |
| ::224.20.20.1 | ::10.10.71.1 | DR |

show pim snooping interfaces

| | |
|---------------------------------|--|
| Syntax | show pim snooping interfaces <brief detail> <instance <i>instance-name</i> > <interface <i>interface-name</i> > <vlan-id <i>vlan-identifier</i> > |
| Release Information | Command introduced in Junos OS Release 12.3. |
| Description | Display information about Protocol Independent Multicast (PIM) snooping interfaces. |
| Options | <p>none—Display detailed information.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>instance <i>instance-name</i>—(Optional) Display PIM snooping interface information for the specified routing instance.</p> <p>interface <i>interface-name</i>—(Optional) Display PIM snooping information for the specified interface only.</p> <p>vlan-id <i>vlan-identifier</i>—(Optional) Display PIM snooping interface information for the specified VLAN.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> PIM Snooping for VPLS Use Cases PIM Snooping for VPLS show pim interfaces on page 232 |
| List of Sample Output | show pim snooping interfaces on page 267 show pim snooping interfaces instance on page 267 show pim snooping interfaces interface on page 267 show pim snooping interfaces vlan-id on page 267 |
| Output Fields | <p>Table 68 on page 265 lists the output fields for the show pim snooping interface command. Output fields are listed in the approximate order in which they appear.</p> |

Table 68: show pim snooping interfaces Output Fields

| Field Name | Field Description | Level of Output |
|------------------------|--|-----------------|
| Instance | Routing instance for PIM snooping. | All levels |
| Learning-Domain | Learning domain for PIM snooping. | All levels |
| Name | Router interfaces that are part of this learning domain. | All levels |

Table 68: show pim snooping interfaces Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-------------------|--|-----------------|
| State | State of the interface: Up , or Down . | All levels |
| IP-Version | Version of IP used: 4 for IPv4, or 6 for IPv6. | All levels |
| NbrCnt | Number of neighboring routers connected through the specified interface. | All levels |
| DR address | IP address of the designated router. | All levels |

Sample Output

show pim snooping interfaces

```
user@host> show pim snooping interfaces
Instance: vpls1
Learning-Domain: vlan-id 10
Name State IP-Version NbrCnt
ge-1/3/1.10 Up 4 1
ge-1/3/3.10 Up 4 1
ge-1/3/5.10 Up 4 1
ge-1/3/7.10 Up 4 1
DR address: 20.0.110.5
DR flooding is ON

Learning-Domain: vlan-id 20
Name State IP-Version NbrCnt
ge-1/3/1.20 Up 4 1
ge-1/3/3.20 Up 4 1
ge-1/3/5.20 Up 4 1
ge-1/3/7.20 Up 4 1
DR address: 20.0.120.5
DR flooding is ON
```

show pim snooping interfaces instance

```
user@host> show pim snooping interfaces instance vpls1
Instance: vpls1

Learning-Domain: vlan-id 10
Name State IP-Version NbrCnt
ge-1/3/1.10 Up 4 1
ge-1/3/3.10 Up 4 1
ge-1/3/5.10 Up 4 1
ge-1/3/7.10 Up 4 1
DR address: 20.0.110.5
DR flooding is ON

Learning-Domain: vlan-id 20
Name State IP-Version NbrCnt
ge-1/3/1.20 Up 4 1
ge-1/3/3.20 Up 4 1
ge-1/3/5.20 Up 4 1
ge-1/3/7.20 Up 4 1
DR address: 20.0.120.5
DR flooding is ON
```

show pim snooping interfaces interface

```
user@host> show pim snooping interfaces interface ge-1/3/1.10
Instance: vpls1
Learning-Domain: vlan-id 10

Name State IP-Version NbrCnt
ge-1/3/1.10 Up 4 1
DR address: 20.0.110.5
DR flooding is ON

Learning-Domain: vlan-id 20
DR address: 20.0.120.5
DR flooding is ON

user@host> show pim snooping interfaces vlan-id 10
```

**show pim snooping
interfaces vlan-id**

Instance: vpls1
Learning-Domain: vlan-id 10

| Name | State | IP-Version | NbrCnt |
|-------------|-------|------------|--------|
| ge-1/3/1.10 | Up | 4 | 1 |
| ge-1/3/3.10 | Up | 4 | 1 |
| ge-1/3/5.10 | Up | 4 | 1 |
| ge-1/3/7.10 | Up | 4 | 1 |

DR address: 20.0.110.5
DR flooding is ON

show pim snooping join

| | |
|---------------------------------|--|
| Syntax | show pim snooping join <brief detail extensive> <instance <i>instance-name</i> > <vlan-id <i>vlan-id</i> > |
| Release Information | Command introduced in Junos OS Release 12.3. |
| Description | Display information about Protocol Independent Multicast (PIM) snooping joins. |
| Options | <p>none—Display detailed information.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>instance <i>instance-name</i>—(Optional) Display PIM snooping join information for the specified routing instance.</p> <p>vlan-id <i>vlan-identifier</i>—(Optional) Display PIM snooping join information for the specified VLAN.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • PIM Snooping for VPLS Use Cases • PIM Snooping for VPLS |
| List of Sample Output | show pim snooping join on page 271 show pim snooping join extensive on page 271 show pim snooping join instance on page 271 show pim snooping join vlan-id on page 272 |
| Output Fields | Table 69 on page 269 lists the output fields for the show pim snooping join command. Output fields are listed in the approximate order in which they appear. |

Table 69: show pim snooping join Output Fields

| Field Name | Field Description | Level of Output |
|-----------------|--|-----------------|
| Instance | Routing instance for PIM snooping. | All levels |
| Learning-Domain | Learning domain for PIM snooping. | All levels |
| Group | Multicast group address. | All levels |
| Source | Multicast source address: <ul style="list-style-type: none"> • * (wildcard value) • <ipv4-address> • <ipv6-address> | All levels |

Table 69: show pim snooping join Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-----------------------------|--|------------------|
| Flags | PIM flags: <ul style="list-style-type: none"> • bidirectional—Bidirectional mode entry. • dense—Dense mode entry. • rptree—Entry is on the rendezvous point tree. • sparse—Sparse mode entry. • spt—Entry is on the shortest-path tree for the source. • wildcard—Entry is on the shared tree. | All levels |
| Upstream state | Information about the upstream interface: <ul style="list-style-type: none"> • Join to RP—Sending a join to the rendezvous point. • Join to Source—Sending a join to the source. • Local RP—Sending neither join messages nor prune messages toward the RP, because this router is the rendezvous point. • Local Source—Sending neither join messages nor prune messages toward the source, because the source is locally attached to this routing device. • Prune to RP—Sending a prune to the rendezvous point. • Prune to Source—Sending a prune to the source. <p>NOTE: RP group range entries have None in the Upstream state field because RP group ranges do not trigger actual PIM join messages between routers.</p> | All levels |
| Upstream neighbor | Information about the upstream neighbor: Direct , Local , Unknown , or a specific IP address. For bidirectional PIM, Direct means that the interface is directly connected to a subnet that contains a phantom RP address. | All levels |
| Upstream port | RPF interface toward the source address for the source-specific state (S,G) or toward the rendezvous point (RP) address for the non-source-specific state (*G). For bidirectional PIM, RP Link means that the interface is directly connected to a subnet that contains a phantom RP address. | All levels |
| Downstream port | Information about downstream interfaces. | extensive |
| Downstream neighbors | Address of the downstream neighbor. | extensive |
| Timeout | Time remaining until the downstream join state is updated (in seconds). | extensive |

Sample Output

```

show pim snooping join user@host> show pim snooping join
Instance: vpls1

Learning-Domain: vlan-id 10
Group: 225.1.1.2
Source: *
Flags: sparse,rptree,wildcard
Upstream state: None
Upstream neighbor: 20.0.110.4, port: ge-1/3/5.10

Learning-Domain: vlan-id 20
Group: 225.1.1.3
Source: *
Flags: sparse,rptree,wildcard
Upstream state: None
Upstream neighbor: 20.0.120.4, port: ge-1/3/5.20

show pim snooping join extensive user@host> show pim snooping join extensive
extensive Instance: vpls1
Learning-Domain: vlan-id 10

Group: 225.1.1.2
Source: *
Flags: sparse,rptree,wildcard
Upstream state: None
Upstream neighbor: 20.0.110.4, port: ge-1/3/5.10
Downstream port: ge-1/3/1.10
Downstream neighbors:
20.0.110.2 State: Join Flags: SRW Timeout: 166

Learning-Domain: vlan-id 20
Group: 225.1.1.3
Source: *
Flags: sparse,rptree,wildcard
Upstream state: None
Upstream neighbor: 20.0.120.4, port: ge-1/3/5.20
Downstream port: ge-1/3/3.20
Downstream neighbors:
20.0.120.3 State: Join Flags: SRW Timeout: 168

show pim snooping join instance user@host> show pim snooping join instance vpls1
instance Instance: vpls1

Learning-Domain: vlan-id 10
Group: 225.1.1.2
Source: *
Flags: sparse,rptree,wildcard
Upstream state: None
Upstream neighbor: 20.0.110.4, port: ge-1/3/5.10

Learning-Domain: vlan-id 20
Group: 225.1.1.3
Source: *
Flags: sparse,rptree,wildcard
Upstream state: None

```

Upstream neighbor: 20.0.120.4, port: ge-1/3/5.20

**show pim snooping join
vlan-id**

```
user@host> show pim snooping join vlan-id 10
Instance: vpls1
Learning-Domain: vlan-id 10
Group: 225.1.1.2
Source: *
Flags: sparse,rptree,wildcard
Upstream state: None
Upstream neighbor: 20.0.110.4, port: ge-1/3/5.10
```

show pim snooping neighbors

| | | |
|---------------------------------|--|--|
| Syntax | <pre>show pim snooping neighbors <brief detail> <instance <i>instance-name</i>> <interface <i>interface-name</i>> <vlan-id <i>vlan-identifier</i>></pre> | |
| Release Information | Command introduced in Junos OS Release 12.3. | |
| Description | Display information about Protocol Independent Multicast (PIM) snooping neighbors. | |
| Options | <p>none—Display detailed information.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>instance <i>instance-name</i>—(Optional) Display PIM snooping neighbor information for the specified routing instance.</p> <p>interface <i>interface-name</i>—(Optional) Display information for the specified PIM snooping neighbor interface.</p> <p>vlan-id <i>vlan-identifier</i>—(Optional) Display PIM snooping neighbor information for the specified VLAN.</p> | |
| Required Privilege Level | view | |
| Related Documentation | <ul style="list-style-type: none"> Configuring Interface Priority for PIM Designated Router Selection Modifying the PIM Hello Interval PIM Snooping for VPLS Use Cases PIM Snooping for VPLS show pim neighbors on page 253 | |
| List of Sample Output | show pim snooping neighbors on page 275 show pim snooping neighbors detail on page 275 show pim snooping neighbors instance on page 276 show pim snooping neighbors interface on page 277 show pim snooping neighbors vlan-id on page 277 | |
| Output Fields | <p>Table 70 on page 273 lists the output fields for the show pim snooping neighbors command. Output fields are listed in the approximate order in which they appear.</p> | |

Table 70: show pim snooping neighbors Output Fields

| Field Name | Field Description | Level of Output |
|------------|------------------------------------|-----------------|
| Instance | Routing instance for PIM snooping. | All levels |

Table 70: show pim snooping neighbors Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-------------------------------------|---|-----------------|
| Learning-Domain | Learning domain for PIM snooping. | All levels |
| Interface | Router interface for which PIM snooping neighbor details are displayed. | All levels |
| Option | PIM snooping options available on the specified interface: <ul style="list-style-type: none"> • H = Hello Option Holdtime • P = Hello Option DR Priority • L = Hello Option LAN Prune Delay • G = Generation Identifier • T = Tracking Bit | All levels |
| Uptime | Time the neighbor has been operational since the PIM process was last initialized, in the format dd:hh:mm:ss ago for less than a week and nwnd:hh:mm:ss ago for more than a week. | All levels |
| Neighbor addr | IP address of the PIM snooping neighbor connected through the specified interface. | All levels |
| Address | IP address of the specified router interface. | All levels |
| Hello Option Holdtime | Time for which the neighbor is available, in seconds. The range of values is 0 through 65,535 . | detail |
| Hello Option DR Priority | Designated router election priority. The range of values is 0 through 255 . NOTE: By default, every PIM interface has the lowest probability (priority 0) of being selected as the DR. | detail |
| Hello Option Generation ID | 9-digit or 10-digit number used to tag hello messages. | detail |
| Hello Option LAN Prune Delay | Time to wait before the neighbor receives prune messages, in the format delay nnn ms override nnnn ms . | detail |

Sample Output

show pim snooping neighbors

```
user@host> show pim snooping neighbors
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority, T = Tracking Bit

Instance: vpls1
Learning-Domain: vlan-id 10

Interface Option Uptime Neighbor addr
ge-1/3/1.10 HPLGT 00:43:33 20.0.110.2
ge-1/3/3.10 HPLGT 00:43:33 20.0.110.3
ge-1/3/5.10 HPLGT 00:43:33 20.0.110.4
ge-1/3/7.10 HPLGT 00:43:33 20.0.110.5

Learning-Domain: vlan-id 20

Interface Option Uptime Neighbor addr
ge-1/3/1.20 HPLGT 00:43:33 20.0.120.2
ge-1/3/3.20 HPLGT 00:43:33 20.0.120.3
ge-1/3/5.20 HPLGT 00:43:33 20.0.120.4
ge-1/3/7.20 HPLGT 00:43:33 20.0.120.5
```

show pim snooping neighbors detail

```
user@host> show pim snooping neighbors detail
Instance: vpls1
Learning-Domain: vlan-id 10

Interface: ge-1/3/1.10
Address: 20.0.110.2
Uptime: 00:44:51
Hello Option Holdtime: 105 seconds 83 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 830908833
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Tracking is supported

Interface: ge-1/3/3.10
Address: 20.0.110.3
Uptime: 00:44:51
Hello Option Holdtime: 105 seconds 97 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 2056520742
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Tracking is supported

Interface: ge-1/3/5.10
Address: 20.0.110.4
Uptime: 00:44:51
Hello Option Holdtime: 105 seconds 81 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 1152066227
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Tracking is supported

Interface: ge-1/3/7.10
Address: 20.0.110.5
Uptime: 00:44:51
Hello Option Holdtime: 105 seconds 96 remaining
```

```
Hello Option DR Priority: 1
Hello Option Generation ID: 1113200338
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Tracking is supported
Learning-Domain: vlan-id 20
```

```
Interface: ge-1/3/1.20
Address: 20.0.120.2
Uptime: 00:44:51
Hello Option Holdtime: 105 seconds 81 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 963205167
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Tracking is supported
```

```
Interface: ge-1/3/3.20
Address: 20.0.120.3
Uptime: 00:44:51
Hello Option Holdtime: 105 seconds 104 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 166921538
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Tracking is supported
```

```
Interface: ge-1/3/5.20
Address: 20.0.120.4
Uptime: 00:44:51
Hello Option Holdtime: 105 seconds 88 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 789422835
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Tracking is supported
```

```
Interface: ge-1/3/7.20
Address: 20.0.120.5
Uptime: 00:44:51
Hello Option Holdtime: 105 seconds 88 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 1563649680
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Tracking is supported
```

show pim snooping neighbors instance

```
user@host> show pim snooping neighbors instance vpls1
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority, T = Tracking Bit
```

```
Instance: vpls1
Learning-Domain: vlan-id 10
```

```
Interface Option Uptime Neighbor addr
ge-1/3/1.10 HPLGT 00:46:03 20.0.110.2
ge-1/3/3.10 HPLGT 00:46:03 20.0.110.3
ge-1/3/5.10 HPLGT 00:46:03 20.0.110.4
ge-1/3/7.10 HPLGT 00:46:03 20.0.110.5
```

```
Learning-Domain: vlan-id 20
```

```
Interface Option Uptime Neighbor addr
ge-1/3/1.20 HPLGT 00:46:03 20.0.120.2
```

```

ge-1/3/3.20 HPLGT 00:46:03 20.0.120.3
ge-1/3/5.20 HPLGT 00:46:03 20.0.120.4
ge-1/3/7.20 HPLGT 00:46:03 20.0.120.5

```

show pim snooping neighbors interface

```

user@host> show pim snooping neighbors interface ge-1/3/1.20
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority, T = Tracking Bit

```

```

Instance: vpls1
Learning-Domain: vlan-id 10
Learning-Domain: vlan-id 20

```

```

Interface Option Uptime Neighbor addr
ge-1/3/1.20 HPLGT 00:48:04 20.0.120.2

```

show pim snooping neighbors vlan-id

```

user@host> show pim snooping neighbors vlan-id 10
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority, T = Tracking Bit

```

```

Instance: vpls1
Learning-Domain: vlan-id 10

```

```

Interface Option Uptime Neighbor addr
ge-1/3/1.10 HPLGT 00:49:12 20.0.110.2
ge-1/3/3.10 HPLGT 00:49:12 20.0.110.3
ge-1/3/5.10 HPLGT 00:49:12 20.0.110.4
ge-1/3/7.10 HPLGT 00:49:12 20.0.110.5

```

show pim snooping statistics

| | |
|---------------------------------|--|
| Syntax | show pim snooping statistics <instance <i>instance-name</i> > <interface <i>interface-name</i> > <vlan-id <i>vlan-id</i> > |
| Release Information | Command introduced in Junos OS Release 12.3. |
| Description | Display Protocol Independent Multicast (PIM) snooping statistics. |
| Options | <p>none—Display PIM statistics.</p> <p>instance <i>instance-name</i>—(Optional) Display statistics for a specific routing instance enabled by Protocol Independent Multicast (PIM) snooping.</p> <p>interface <i>interface-name</i>—(Optional) Display statistics about the specified interface for PIM snooping.</p> <p>vlan-id <i>vlan-identifier</i>—(Optional) Display PIM snooping statistics information for the specified VLAN.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> PIM Snooping for VPLS Use Cases PIM Snooping for VPLS clear pim snooping statistics on page 120 |
| List of Sample Output | show pim snooping statistics on page 280 show pim snooping statistics instance on page 280 show pim snooping statistics interface on page 281 show pim snooping statistics vlan-id on page 281 |
| Output Fields | <p>Table 71 on page 278 lists the output fields for the show pim snooping statistics command. Output fields are listed in the approximate order in which they appear.</p> |

Table 71: show pim snooping statistics Output Fields

| Field Name | Field Description | Level of Output |
|-----------------|---|-----------------|
| Instance | Routing instance for PIM snooping. | All levels |
| Learning-Domain | Learning domain for PIM snooping. | All levels |
| Tx J/P messages | Total number of transmitted join/prune packets. | All levels |
| RX J/P messages | Total number of received join/prune packets. | All levels |

Table 71: show pim snooping statistics Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-------------------------------------|---|-----------------|
| Rx J/P messages -- seen | Number of join/prune packets seen but not received on the upstream interface. | All levels |
| Rx J/P messages -- received | Number of join/prune packets received on the downstream interface. | All levels |
| Rx Hello messages | Total number of received hello packets. | All levels |
| Rx Version Unknown | Number of packets received with an unknown version number. | All levels |
| Rx Neighbor Unknown | Number of packets received from an unknown neighbor. | All levels |
| Rx Upstream Neighbor Unknown | Number of packets received with unknown upstream neighbor information. | All levels |
| Rx Bad Length | Number of packets received containing incorrect length information. | All levels |
| Rx J/P Busy Drop | Number of join/prune packets dropped while the router is busy. | All levels |
| Rx J/P Group Aggregate 0 | Number of join/prune packets received containing the aggregate group information. | All levels |
| Rx Malformed Packet | Number of malformed packets received. | All levels |
| Rx No PIM Interface | Number of packets received without the interface information. | All levels |
| Rx No Upstream Neighbor | Number of packets received without upstream neighbor information. | All levels |
| Rx Unknown Hello Option | Number of hello packets received with unknown options. | All levels |

Sample Output

**show pim snooping
statistics**

```
user@host> show pim snooping statistics
```

```
Instance: vpls1
```

```
Learning-Domain: vlan-id 10
```

```
Tx J/P messages 0
RX J/P messages 8
Rx J/P messages -- seen 0
Rx J/P messages -- received 8
Rx Hello messages 37
Rx Version Unknown 0
Rx Neighbor Unknown 0
Rx Upstream Neighbor Unknown 0
Rx Bad Length 0
Rx J/P Busy Drop 0
Rx J/P Group Aggregate 0
Rx Malformed Packet 0
Rx No PIM Interface 0
Rx No Upstream Neighbor 0
Rx Bad Length 0
Rx Neighbor Unknown 0
Rx Unknown Hello Option 0
Rx Malformed Packet 0
```

```
Learning-Domain: vlan-id 20
```

```
Tx J/P messages 0
RX J/P messages 2
Rx J/P messages -- seen 0
Rx J/P messages -- received 2
Rx Hello messages 39
Rx Version Unknown 0
Rx Neighbor Unknown 0
Rx Upstream Neighbor Unknown 0
Rx Bad Length 0
Rx J/P Busy Drop 0
Rx J/P Group Aggregate 0
Rx Malformed Packet 0
Rx No PIM Interface 0
Rx No Upstream Neighbor 0
Rx Bad Length 0
Rx Neighbor Unknown 0
Rx Unknown Hello Option 0
Rx Malformed Packet 0
```

**show pim snooping
statistics instance**

```
user@host> show pim snooping statistics instance vpls1
```

```
Instance: vpls1
```

```
Learning-Domain: vlan-id 10
```

```
Tx J/P messages 0
RX J/P messages 9
Rx J/P messages -- seen 0
Rx J/P messages -- received 9
Rx Hello messages 45
Rx Version Unknown 0
Rx Neighbor Unknown 0
Rx Upstream Neighbor Unknown 0
```

```

Rx Bad Length 0
Rx J/P Busy Drop 0
Rx J/P Group Aggregate 0
Rx Malformed Packet 0
Rx No PIM Interface 0
Rx No Upstream Neighbor 0
Rx Bad Length 0
Rx Neighbor Unknown 0
Rx Unknown Hello Option 0
Rx Malformed Packet 0

```

```
Learning-Domain: vlan-id 20
```

```

Tx J/P messages 0
RX J/P messages 3
Rx J/P messages -- seen 0
Rx J/P messages -- received 3
Rx Hello messages 47
Rx Version Unknown 0
Rx Neighbor Unknown 0
Rx Upstream Neighbor Unknown 0
Rx Bad Length 0
Rx J/P Busy Drop 0
Rx J/P Group Aggregate 0
Rx Malformed Packet 0
Rx No PIM Interface 0
Rx No Upstream Neighbor 0
Rx Bad Length 0
Rx Neighbor Unknown 0
Rx Unknown Hello Option 0
Rx Malformed Packet 0

```

show pim snooping statistics interface

```
user@host> show pim snooping statistics interface ge-1/3/1.20
```

```

Instance: vpls1
Learning-Domain: vlan-id 10
Learning-Domain: vlan-id 20

```

```
PIM Interface statistics for ge-1/3/1.20
```

```

Tx J/P messages 0
RX J/P messages 0
Rx J/P messages -- seen 0
Rx J/P messages -- received 0
Rx Hello messages 13
Rx Version Unknown 0
Rx Neighbor Unknown 0
Rx Upstream Neighbor Unknown 0
Rx Bad Length 0
Rx J/P Busy Drop 0
Rx J/P Group Aggregate 0
Rx Malformed Packet 0

```

show pim snooping statistics vlan-id

```
user@host> show pim snooping statistics vlan-id 10
```

```

Instance: vpls1
Learning-Domain: vlan-id 10

```

```

Tx J/P messages 0
RX J/P messages 11
Rx J/P messages -- seen 0
Rx J/P messages -- received 11

```

```
Rx Hello messages 64
Rx Version Unknown 0
Rx Neighbor Unknown 0
Rx Upstream Neighbor Unknown 0
Rx Bad Length 0
Rx J/P Busy Drop 0
Rx J/P Group Aggregate 0
Rx Malformed Packet 0
Rx No PIM Interface 0
Rx No Upstream Neighbor 0
Rx Bad Length 0
Rx Neighbor Unknown 0
```

show pim source

| | |
|---|--|
| Syntax | <pre>show pim source <brief detail> <inet inet6> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <source-prefix></pre> |
| Syntax (EX Series Switch and the QFX Series) | <pre>show pim source <brief detail> <inet inet6> <instance <i>instance-name</i>> <source-prefix></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> |
| Description | Display information about the Protocol Independent Multicast (PIM) source reverse path forwarding (RPF) state. |
| Options | <p>none—Display standard information about the PIM RPF state for all supported family addresses for all routing instances.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>inet inet6—(Optional) Display information for IPv4 or IPv6 family addresses, respectively.</p> <p>instance <i>instance-name</i>—(Optional) Display information about the RPF state for a specific PIM-enabled routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>source-prefix—(Optional) Display the state for source RPF states in the given range.</p> |
| Required Privilege Level | view |
| List of Sample Output | <p>show pim source on page 284</p> <p>show pim source brief on page 284</p> <p>show pim source detail on page 284</p> |
| Output Fields | Table 72 on page 284 describes the output fields for the show pim source command. Output fields are listed in the approximate order in which they appear. |

Table 72: show pim source Output Fields

| Field Name | Field Description |
|---------------------------|---|
| Instance | Name of the routing instance. |
| Source | Address of the source or reverse path. |
| Prefix/length | Prefix and prefix length for the route used to reach the RPF address. |
| Upstream interface | RPF interface toward the source address. |
| Upstream Neighbor | Address of the RPF neighbor used to reach the source address. |

Sample Output

show pim source

```

user@host> show pim source
Instance: PIM.master Family: INET

Source 10.255.14.144
  Prefix 10.255.14.144/32
  Upstream interface Local
  Upstream neighbor Local

Source 10.255.70.15
  Prefix 10.255.70.15/32
  Upstream interface so-1/0/0.0
  Upstream neighbor 10.111.10.2

Instance: PIM.master Family: INET6

```

show pim source brief

The output for the **show pim source brief** command is identical to that for the **show pim source** command. For sample output, see [show pim source on page 284](#).

show pim source detail

```

user@host> show pim source detail
Instance: PIM.master Family: INET

Source 10.255.14.144
  Prefix 10.255.14.144/32
  Upstream interface Local
  Upstream neighbor Local
  Active groups:228.0.0.0
    239.1.1.1
    239.1.1.1

Source 10.255.70.15
  Prefix 10.255.70.15/32
  Upstream interface so-1/0/0.0
  Upstream neighbor 10.111.10.2
  Active groups:239.1.1.1

Instance: PIM.master Family: INET6

```

show pim statistics

| | |
|---|---|
| Syntax | <pre>show pim statistics <inet inet6> <instance <i>instance-name</i>> <interface <i>interface-name</i>> <logical-system (all <i>logical-system-name</i>)></pre> |
| Syntax (EX Series Switch and the QFX Series) | <pre>show pim statistics <inet inet6> <instance <i>instance-name</i>> <interface <i>interface-name</i>></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional PIM added in Junos OS Release 12.1.</p> |
| Description | Display Protocol Independent Multicast (PIM) statistics. |
| Options | <p>none—Display PIM statistics.</p> <p>inet inet6—(Optional) Display IPv4 or IPv6 PIM statistics, respectively.</p> <p>instance <i>instance-name</i>—(Optional) Display statistics for a specific routing instance enabled by Protocol Independent Multicast (PIM).</p> <p>interface <i>interface-name</i>—(Optional) Display statistics about the specified interface.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • clear pim statistics on page 122 |
| List of Sample Output | <p>show pim statistics on page 293</p> <p>show pim statistics inet interface <interface-name> on page 294</p> <p>show pim statistics inet6 interface <interface-name> on page 294</p> <p>show pim statistics instance <instance-name> on page 295</p> <p>show pim statistics interface <interface-name> on page 297</p> |
| Output Fields | <p>Table 73 on page 286 describes the output fields for the show pim statistics command. Output fields are listed in the approximate order in which they appear.</p> |

Table 73: show pim statistics Output Fields

| Field Name | Field Description |
|-------------------------|---|
| Instance | <p>Name of the routing instance.</p> <p>This field only appears if you specify an interface, for example:</p> <ul style="list-style-type: none"> • inet interface <i>interface-name</i> • inet6 interface <i>interface-name</i> • interface <i>interface-name</i> |
| Family | <p>Output is for IPv4 or IPv6 PIM statistics. INET indicates IPv4 statistics, and INET6 indicates IPv6 statistics.</p> <p>This field only appears if you specify an interface, for example:</p> <ul style="list-style-type: none"> • inet interface <i>interface-name</i> • inet6 interface <i>interface-name</i> • interface <i>interface-name</i> |
| PIM statistics | PIM statistics for all interfaces or for the specified interface. |
| PIM message type | Message type for which statistics are displayed. |
| Received | Number of received statistics. |
| Sent | Number of messages sent of a certain type. |
| Rx errors | Number of received packets that contained errors. |
| V2 Hello | PIM version 2 hello packets. |
| V2 Register | PIM version 2 register packets. |
| V2 Register Stop | PIM version 2 register stop packets. |
| V2 Join Prune | PIM version 2 join and prune packets. |
| V2 Bootstrap | PIM version 2 bootstrap packets. |
| V2 Assert | PIM version 2 assert packets. |
| V2 Graft | PIM version 2 graft packets. |
| V2 Graft Ack | PIM version 2 graft acknowledgment packets. |
| V2 Candidate RP | PIM version 2 candidate RP packets. |

Table 73: show pim statistics Output Fields (*continued*)

| Field Name | Field Description |
|---|--|
| V2 State Refresh | PIM version 2 control messages related to PIM dense mode (PIM-DM) state refresh. State refresh is an extension to PIM-DM. It not supported in Junos OS. |
| V2 DF Election | PIM version 2 send and receive messages associated with bidirectional PIM designated forwarder election. |
| V1 Query | PIM version 1 query packets. |
| V1 Register | PIM version 1 register packets. |
| V1 Register Stop | PIM version 1 register stop packets. |
| V1 Join Prune | PIM version 1 join and prune packets. |
| V1 RP Reachability | PIM version 1 RP reachability packets. |
| V1 Assert | PIM version 1 assert packets. |
| V1 Graft | PIM version 1 graft packets. |
| V1 Graft Ack | PIM version 1 graft acknowledgment packets. |
| AutoRP Announce | Auto-RP announce packets. |
| AutoRP Mapping | Auto-RP mapping packets. |
| AutoRP Unknown type | Auto-RP packets with an unknown type. |
| Anycast Register | Auto-RP announce packets. |
| Anycast Register Stop | Auto-RP announce packets. |
| Global Statistics | Summary of PIM statistics for all interfaces. |
| Hello dropped on neighbor policy | Number of hello packets dropped because of a configured neighbor policy. |
| Unknown type | Number of PIM control packets received with an unknown type. |
| V1 Unknown type | Number of PIM version 1 control packets received with an unknown type. |
| Unknown Version | Number of PIM control packets received with an unknown version. The version is not version 1 or version 2. |

Table 73: show pim statistics Output Fields (*continued*)

| Field Name | Field Description |
|--|---|
| Neighbor unknown | Number of PIM control packets received (excluding PIM hello) without first receiving the hello packet. |
| Bad Length | Number of PIM control packets received for which the packet size does not match the PIM length field in the packet. |
| Bad Checksum | Number of PIM control packets received for which the calculated checksum does not match the checksum field in the packet. |
| Bad Receive If | Number of PIM control packets received on an interface that does not have PIM configured. |
| Rx Bad Data | Number of PIM control packets received that contain data for TCP Bad register packets. |
| Rx Intf disabled | Number of PIM control packets received on an interface that has PIM disabled. |
| Rx V1 Require V2 | Number of PIM version 1 control packets received on an interface configured for PIM version 2. |
| Rx V2 Require V1 | Number of PIM version 2 control packets received on an interface configured for PIM version 1. |
| Rx Register not RP | Number of PIM register packets received when the router is not the RP for the group. |
| Rx Register no route | Number of PIM register packets received when the RP does not have a unicast route back to the source. |
| Rx Register no decap if | Number of PIM register packets received when the RP does not have a de-encapsulation interface. |
| Null Register Timeout | Number of NULL register timeout packets. |
| RP Filtered Source | Number of PIM packets received when the router has a source address filter configured for the RP. |
| Rx Unknown Reg Stop | Number of register stop messages received with an unknown type. |
| Rx Join/Prune no state | Number of join and prune messages received for which the router has no state. |
| Rx Join/Prune on upstream if | Number of join and prune messages received on the interface used to reach the upstream router, toward the RP. |
| Rx Join/Prune for invalid group | Number of join or prune messages received for invalid multicast group addresses. |

Table 73: show pim statistics Output Fields (*continued*)

| Field Name | Field Description |
|---------------------------------------|---|
| Rx Join/Prune messages dropped | Number of join and prune messages received and dropped. |
| Rx sparse join for dense group | Number of PIM sparse mode join messages received for a group that is configured for dense mode. |
| Rx Graft/Graft Ack no state | Number of graft and graft acknowledgment messages received for which the router or switch has no state. |
| Rx Graft on upstream if | Number of graft messages received on the interface used to reach the upstream router, toward the RP. |
| Rx CRP not BSR | Number of BSR messages received in which the PIM message type is Candidate-RP-Advertisement, not Bootstrap. |
| Rx BSR when BSR | Number of BSR messages received in which the PIM message type is Bootstrap. |
| Rx BSR not RPF if | Number of BSR messages received on an interface that is not the RPF interface. |
| Rx unknown hello opt | Number of PIM hello packets received with options that Junos OS does not support. |
| Rx data no state | Number of PIM control packets received for which the router has no state for the data type. |
| Rx RP no state | Number of PIM control packets received for which the router has no state for the RP. |
| Rx aggregate | Number of PIM aggregate MDT packets received. |
| Rx malformed packet | Number of PIM control packets received with a malformed IP unicast or multicast address family. |
| No RP | Number of PIM control packets received with no RP address. |
| No register encaps if | Number of PIM register packets received when the first-hop router does not have an encapsulation interface. |
| No route upstream | Number of PIM control packets received when the router does not have a unicast route to the the interface used to reach the upstream router, toward the RP. |
| Nexthop Unusable | Number of PIM control packets with an unusable nexthop. A path can be unusable if the route is hidden or the link is down. |
| RP mismatch | Number of PIM control packets received for which the router has an RP mismatch. |

Table 73: show pim statistics Output Fields (*continued*)

| Field Name | Field Description |
|--|--|
| RP mode mismatch | RP mode (sparse or bidirectional) mismatches encountered when processing join and prune messages. |
| RPF neighbor unknown | Number of PIM control packets received for which the router has an unknown RPF neighbor for the source. |
| Rx Joins/Prunes filtered | The number of join and prune messages filtered because of configured route filters and source address filters. |
| Tx Joins/Prunes filtered | The number of join and prune messages filtered because of configured route filters and source address filters. |
| Embedded-RP invalid addr | Number of packets received with an invalid embedded RP address in PIM join messages and other types of messages sent between routing domains. |
| Embedded-RP limit exceed | Number of times the limit configured with the maximum-rps statement is exceeded. The maximum-rps statement limits the number of embedded RPs created in a specific routing instance. The range is from 1 through 500. The default is 100. |
| Embedded-RP added | <p>Number of packets in which the embedded RP for IPv6 is added.</p> <p>The following receive events trigger extraction of an IPv6 embedded RP address on the router:</p> <ul style="list-style-type: none"> • Multicast Listener Discovery (MLD) report for an embedded RP multicast group address • PIM join message with an embedded RP multicast group address • Static embedded RP multicast group address associated with an interface • Packets sent to an embedded RP multicast group address received on the DR <p>An embedded RP node discovered through these receive events is added if it does not already exist on the routing platform.</p> |
| Embedded-RP removed | Number of packets in which the embedded RP for IPv6 is removed. The embedded RP is removed whenever all PIM join states using this RP are removed or the configuration changes to remove the embedded RP feature. |
| Rx Register msgs filtering drop | Number of received register messages dropped because of a filter configured for PIM register messages. |
| Tx Register msgs filtering drop | Number of register messages dropped because of a filter configured for PIM register messages. |
| Rx Bidir Join/Prune on non-Bidir if | Error counter for join and prune messages received on non-bidirectional PIM interfaces. |

Table 73: show pim statistics Output Fields (*continued*)

| Field Name | Field Description |
|---|---|
| Rx Bidir Join/Prune on non-DF if | Error counter for join and prune messages received on non-designated forwarder interfaces. |
| V4 (S,G) Maximum | Maximum number of (S,G) IPv4 multicast routes accepted for the VPN routing and forwarding (VRF) routing instance. If this number is met, additional (S,G) entries are not accepted. |
| V4 (S,G) Accepted | Number of accepted (S,G) IPv4 multicast routes. |
| V4 (S,G) Threshold | Threshold at which a warning message is logged (percentage of the maximum number of (S,G) IPv4 multicast routes accepted by the device). |
| V4 (S,G) Log Interval | Time (in seconds) between consecutive log messages. |
| V6 (S,G) Maximum | Maximum number of (S,G) IPv6 multicast routes accepted for the VPN routing and forwarding (VRF) routing instance. If this number is met, additional (S,G) entries are not accepted. |
| V6 (S,G) Accepted | Number of accepted (S,G) IPv6 multicast routes. |
| V6 (S,G) Threshold | Threshold at which a warning message is logged (percentage of the maximum number of (S,G) IPv6 multicast routes accepted by the device). |
| V6 (S,G) Log Interval | Time (in seconds) between consecutive log messages. |
| V4 (grp-prefix, RP) Maximum | Maximum number of group-to-rendezvous point (RP) IPv4 multicast mappings accepted for the VRF routing instance. If this number is met, additional mappings are not accepted. |
| V4 (grp-prefix, RP) Accepted | Number of accepted group-to-RP IPv4 multicast mappings. |
| V4 (grp-prefix, RP) Threshold | Threshold at which a warning message is logged (percentage of the maximum number of group-to-RP IPv4 multicast mappings accepted by the device). |
| V4 (grp-prefix, RP) Log Interval | Time (in seconds) between consecutive log messages. |
| V6 (grp-prefix, RP) Maximum | Maximum number of group-to RP IPv6 multicast mappings accepted for the VRF routing instance. If this number is met, additional mappings are not accepted. |
| V6 (grp-prefix, RP) Accepted | Number of accepted group-to-RP IPv6 multicast mappings. |

Table 73: show pim statistics Output Fields (*continued*)

| Field Name | Field Description |
|---|---|
| V6 (grp-prefix, RP) Threshold | Threshold at which a warning message is logged (percentage of the maximum number of group-to-RP IPv6 multicast mappings accepted by the device). |
| V6 (grp-prefix, RP) Log Interval | Time (in seconds) between consecutive log messages. |
| V4 Register Maximum | Maximum number of IPv4 PIM registers accepted for the VRF routing instance. If this number is met, additional PIM registers are not accepted. You configure the register limits on the RP. |
| V4 Register Accepted | Number of accepted IPv4 PIM registers. |
| V4 Register Threshold | Threshold at which a warning message is logged (percentage of the maximum number of IPv4 PIM registers accepted by the device). |
| V4 Register Log Interval | Time (in seconds) between consecutive log messages. |
| V6 Register Maximum | Maximum number of IPv6 PIM registers accepted for the VRF routing instance. If this number is met, additional PIM registers are not accepted. You configure the register limits on the RP. |
| V6 Register Accepted | Number of accepted IPv6 PIM registers. |
| V6 Register Threshold | Threshold at which a warning message is logged (percentage of the maximum number of IPv6 PIM registers accepted by the device). |
| V6 Register Log Interval | Time (in seconds) between consecutive log messages. |

Sample Output

show pim statistics

```

user@host> show pim statistics
PIM Message type      Received      Sent  Rx errors
V2 Hello               15           32      0
V2 Register            0           362     0
V2 Register Stop      483           0      0
V2 Join Prune         18           518     0
V2 Bootstrap           0            0      0
V2 Assert              0            0      0
V2 Graft               0            0      0
V2 Graft Ack           0            0      0
V2 Candidate RP        0            0      0
V2 State Refresh       0            0      0
V2 DF Election         0            0      0
V1 Query               0            0      0
V1 Register            0            0      0
V1 Register Stop       0            0      0
V1 Join Prune          0            0      0
V1 RP Reachability     0            0      0
V1 Assert              0            0      0
V1 Graft               0            0      0
V1 Graft Ack           0            0      0
AutoRP Announce        0            0      0
AutoRP Mapping         0            0      0
AutoRP Unknown type    0            0      0
Anycast Register       0            0      0
Anycast Register Stop  0            0      0

Global Statistics

Hello dropped on neighbor policy  0
Unknown type                      0
V1 Unknown type                   0
Unknown Version                   0
Neighbor unknown                  0
Bad Length                        0
Bad Checksum                      0
Bad Receive If                    0
Rx Bad Data                       0
Rx Intf disabled                   0
Rx V1 Require V2                  0
Rx V2 Require V1                  0
Rx Register not RP                0
Rx Register no route              0
Rx Register no decap if           0
Null Register Timeout             0
RP Filtered Source                0
Rx Unknown Reg Stop               0
Rx Join/Prune no state            0
Rx Join/Prune on upstream if      0
Rx Join/Prune for invalid group   5
Rx Join/Prune messages dropped    0
Rx sparse join for dense group    0
Rx Graft/Graft Ack no state       0
Rx Graft on upstream if           0
Rx CRP not BSR                    0
Rx BSR when BSR                   0
Rx BSR not RPF if                 0
Rx unknown hello opt              0

```

| | |
|-------------------------------------|---|
| Rx data no state | 0 |
| Rx RP no state | 0 |
| Rx aggregate | 0 |
| Rx malformed packet | 0 |
| Rx illegal TTL | 0 |
| Rx illegal destination address | 0 |
| No RP | 0 |
| No register encap if | 0 |
| No route upstream | 0 |
| Nexthop Unusable | 0 |
| RP mismatch | 0 |
| RP mode mismatch | 0 |
| RPF neighbor unknown | 0 |
| Rx Joins/Prunes filtered | 0 |
| Tx Joins/Prunes filtered | 0 |
| Embedded-RP invalid addr | 0 |
| Embedded-RP limit exceed | 0 |
| Embedded-RP added | 0 |
| Embedded-RP removed | 0 |
| Rx Register msgs filtering drop | 0 |
| Tx Register msgs filtering drop | 0 |
| Rx Bidir Join/Prune on non-Bidir if | 0 |
| Rx Bidir Join/Prune on non-DF if | 0 |

Sample Output

```
show pim statistics
inet interface
<interface-name>
```

```
user@host> show pim statistics inet interface ge-0/3/0.0
Instance: PIM.master Family: INET
```

PIM Interface statistics for ge-0/3/0.0

| PIM Message type | Received | Sent | Rx errors |
|-----------------------|----------|------|-----------|
| V2 Hello | 0 | 4 | 0 |
| V2 Register | 0 | 0 | 0 |
| V2 Register Stop | 0 | 0 | 0 |
| V2 Join Prune | 0 | 0 | 0 |
| V2 Bootstrap | 0 | 0 | 0 |
| V2 Assert | 0 | 0 | 0 |
| V2 Graft | 0 | 0 | 0 |
| V2 Graft Ack | 0 | 0 | 0 |
| V2 Candidate RP | 0 | 0 | 0 |
| V1 Query | 0 | 0 | 0 |
| V1 Register | 0 | 0 | 0 |
| V1 Register Stop | 0 | 0 | 0 |
| V1 Join Prune | 0 | 0 | 0 |
| V1 RP Reachability | 0 | 0 | 0 |
| V1 Assert | 0 | 0 | 0 |
| V1 Graft | 0 | 0 | 0 |
| V1 Graft Ack | 0 | 0 | 0 |
| AutoRP Announce | 0 | 0 | 0 |
| AutoRP Mapping | 0 | 0 | 0 |
| AutoRP Unknown type | 0 | | |
| Anycast Register | 0 | 0 | 0 |
| Anycast Register Stop | 0 | 0 | 0 |

Sample Output

```
show pim statistics
inet6 interface
```

```
user@host> show pim statistics inet6 interface ge-0/3/0.0
Instance: PIM.master Family: INET6
```



```
<interface-name> PIM Interface statistics for ge-0/3/0.0

PIM Message type      Received      Sent  Rx errors
V2 Hello               0             4      0
V2 Register            0             0      0
V2 Register Stop       0             0      0
V2 Join Prune          0             0      0
V2 Bootstrap           0             0      0
V2 Assert              0             0      0
V2 Graft               0             0      0
V2 Graft Ack           0             0      0
V2 Candidate RP        0             0      0
Anycast Register       0             0      0
Anycast Register Stop  0             0      0
```

```
show pim statistics instance VPN-A
instance
<instance-name>

user@host> show pim statistics instance VPN-A
PIM Message type      Received      Sent  Rx errors
V2 Hello               31            37      0
V2 Register            0             0      0
V2 Register Stop       0             0      0
V2 Join Prune          0            16      0
V2 Bootstrap           0             0      0
V2 Assert              0             0      0
V2 Graft               0             0      0
V2 Graft Ack           0             0      0
V2 Candidate RP        0             0      0
V2 State Refresh       0             0      0
V2 DF Election         0             0      0
V1 Query              0             0      0
V1 Register            0             0      0
V1 Register Stop       0             0      0
V1 Join Prune          0             0      0
V1 RP Reachability     0             0      0
V1 Assert              0             0      0
V1 Graft               0             0      0
V1 Graft Ack           0             0      0
AutoRP Announce        0             0      0
AutoRP Mapping         0             0      0
AutoRP Unknown type    0             0      0
Anycast Register       0             0      0
Anycast Register Stop  0             0      0
```

Global Statistics

```
Hello dropped on neighbor policy      0
Unknown type                          0
V1 Unknown type                       0
Unknown Version                       0
Neighbor unknown                      0
Bad Length                           0
Bad Checksum                         0
Bad Receive If                       0
Rx Bad Data                          0
Rx Intf disabled                     0
Rx V1 Require V2                     0
Rx V2 Require V1                     0
Rx Register not RP                   0
Rx Register no route                 0
Rx Register no decap if              0
Null Register Timeout                 0
RP Filtered Source                    0
```

| | |
|-------------------------------------|-----|
| Rx Unknown Reg Stop | 0 |
| Rx Join/Prune no state | 0 |
| Rx Join/Prune on upstream if | 0 |
| Rx Join/Prune for invalid group | 0 |
| Rx Join/Prune messages dropped | 0 |
| Rx sparse join for dense group | 0 |
| Rx Graft/Graft Ack no state | 0 |
| Rx Graft on upstream if | 0 |
| Rx CRP not BSR | 0 |
| Rx BSR when BSR | 0 |
| Rx BSR not RPF if | 0 |
| Rx unknown hello opt | 0 |
| Rx data no state | 0 |
| Rx RP no state | 0 |
| Rx aggregate | 0 |
| Rx malformed packet | 0 |
| Rx illegal TTL | 0 |
| Rx illegal destination address | 0 |
| No RP | 0 |
| No register encap if | 0 |
| No route upstream | 28 |
| Nexthop Unusable | 0 |
| RP mismatch | 0 |
| RP mode mismatch | 0 |
| RPF neighbor unknown | 0 |
| Rx Joins/Prunes filtered | 0 |
| Tx Joins/Prunes filtered | 0 |
| Embedded-RP invalid addr | 0 |
| Embedded-RP limit exceed | 0 |
| Embedded-RP added | 0 |
| Embedded-RP removed | 0 |
| Rx Register msgs filtering drop | 0 |
| Tx Register msgs filtering drop | 0 |
| Rx Bidir Join/Prune on non-Bidir if | 0 |
| Rx Bidir Join/Prune on non-DF if | 0 |
| V4 (S,G) Maximum | 10 |
| V4 (S,G) Accepted | 9 |
| V4 (S,G) Threshold | 80 |
| V4 (S,G) Log Interval | 80 |
| V6 (S,G) Maximum | 8 |
| V6 (S,G) Accepted | 8 |
| V6 (S,G) Threshold | 50 |
| V6 (S,G) Log Interval | 100 |
| V4 (grp-prefix, RP) Maximum | 100 |
| V4 (grp-prefix, RP) Accepted | 5 |
| V4 (grp-prefix, RP) Threshold | 80 |
| V4 (grp-prefix, RP) Log Interval | 10 |
| V6 (grp-prefix, RP) Maximum | 20 |
| V6 (grp-prefix, RP) Accepted | 0 |
| V6 (grp-prefix, RP) Threshold | 90 |
| V6 (grp-prefix, RP) Log Interval | 20 |
| V4 Register Maximum | 100 |
| V4 Register Accepted | 10 |
| V4 Register Threshold | 80 |
| V4 Register Log Interval | 10 |
| V6 Register Maximum | 20 |
| V6 Register Accepted | 0 |
| V6 Register Threshold | 90 |
| V6 Register Log Interval | 20 |

Sample Output

```
show pim statistics
interface
<interface-name>
```

```
user@host> show pim statistics interface ge-0/3/0.0
Instance: PIM.master Family: INET
```

PIM Interface statistics for ge-0/3/0.0

| PIM Message type | Received | Sent | Rx errors |
|-----------------------|----------|------|-----------|
| V2 Hello | 0 | 3 | 0 |
| V2 Register | 0 | 0 | 0 |
| V2 Register Stop | 0 | 0 | 0 |
| V2 Join Prune | 0 | 0 | 0 |
| V2 Bootstrap | 0 | 0 | 0 |
| V2 Assert | 0 | 0 | 0 |
| V2 Graft | 0 | 0 | 0 |
| V2 Graft Ack | 0 | 0 | 0 |
| V2 Candidate RP | 0 | 0 | 0 |
| V1 Query | 0 | 0 | 0 |
| V1 Register | 0 | 0 | 0 |
| V1 Register Stop | 0 | 0 | 0 |
| V1 Join Prune | 0 | 0 | 0 |
| V1 RP Reachability | 0 | 0 | 0 |
| V1 Assert | 0 | 0 | 0 |
| V1 Graft | 0 | 0 | 0 |
| V1 Graft Ack | 0 | 0 | 0 |
| AutoRP Announce | 0 | 0 | 0 |
| AutoRP Mapping | 0 | 0 | 0 |
| AutoRP Unknown type | 0 | | |
| Anycast Register | 0 | 0 | 0 |
| Anycast Register Stop | 0 | 0 | 0 |

Instance: PIM.master Family: INET6

PIM Interface statistics for ge-0/3/0.0

| PIM Message type | Received | Sent | Rx errors |
|-----------------------|----------|------|-----------|
| V2 Hello | 0 | 3 | 0 |
| V2 Register | 0 | 0 | 0 |
| V2 Register Stop | 0 | 0 | 0 |
| V2 Join Prune | 0 | 0 | 0 |
| V2 Bootstrap | 0 | 0 | 0 |
| V2 Assert | 0 | 0 | 0 |
| V2 Graft | 0 | 0 | 0 |
| V2 Graft Ack | 0 | 0 | 0 |
| V2 Candidate RP | 0 | 0 | 0 |
| Anycast Register | 0 | 0 | 0 |
| Anycast Register Stop | 0 | 0 | 0 |

show sap listen

| | |
|---------------------------------|---|
| Syntax | show sap listen <brief detail> <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display the addresses that the router is listening to in order to receive multicast Session Announcement Protocol (SAP) session announcements. |
| Options | <p>none—Display standard information about the addresses that the router is listening to in order to receive multicast SAP session announcements.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show sap listen on page 298 show sap listen brief on page 298 show sap listen detail on page 298 |
| Output Fields | Table 74 on page 298 describes the output fields for the show sap listen command. Output fields are listed in the approximate order in which they appear. |

Table 74: show sap listen Output Fields

| Field Name | Field Description |
|----------------------|--|
| Group address | Address of the group that the local router is listening to for SAP messages. |
| Port | UDP port number used for SAP. |

Sample Output

```

show sap listen
user@host> show sap listen
Group address  Port
224.2.127.254  9875
239.255.255.255 9875

```

show sap listen brief The output for the **show sap listen brief** command is identical to that for the **show sap listen** command. For sample output, see [show sap listen on page 298](#).

show sap listen detail The output for the **show sap listen detail** command is identical to that for the **show sap listen** command. For sample output, see [show sap listen on page 298](#).

test msdp

| | |
|---------------------------------|--|
| Syntax | test msdp (dependent-peers <i>prefix</i> rpf-peer <i>originator</i>) <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 12.1 for the QFX Series. |
| Description | Find Multicast Source Discovery Protocol (MSDP) peers. |
| Options | <p>dependent-peers <i>prefix</i>—Find downstream dependent MSDP peers.</p> <p>rpf-peer <i>originator</i>—Find the MSDP reverse-path-forwarding (RPF) peer for the originator.</p> <p>instance <i>instance-name</i>—(Optional) Find MDSP peers for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | test msdp dependent-peers on page 299 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

```
test msdp
dependent-peers      user@host> test msdp dependent-peers 10.0.0.1/24
```


CHAPTER 6

IPv6 Operational Mode Commands

Table 75 on page 301 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot Internet Protocol version 6 (IPv6). Commands are listed in alphabetical order.

Table 75: IPv6 Operational Mode Commands

| Task | Command |
|---|--|
| Clear IPv6 neighbor cache information. | <code>clear ipv6 neighbors</code> |
| Clear IPv6 router advertisement counters. | <code>clear ipv6 router-advertisement</code> |
| Display neighbor discovery information. | <code>show ipv6 neighbors</code> |
| Display router advertisement information. | <code>show ipv6 router-advertisement</code> |



NOTE: For information about how to configure IPv6 parameters, see the Junos OS Routing Protocols Configuration Guide.

clear ipv6 neighbors

| | |
|---------------------------------|---|
| Syntax | <code>clear ipv6 neighbors</code> <code><all host <i>hostname</i>></code> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.3 for EX Series switches. Command introduced in Junos OS Release 12.2 for the QFX Series. |
| Description | Clear IPv6 neighbor cache information. |
| Options | none —Clear all IPv6 neighbor cache information. all —(Optional) Clear all IPv6 neighbor cache information. host <i>hostname</i> —(Optional) Clear the information for the specified IPv6 neighbors. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">• show ipv6 neighbors on page 304 |
| List of Sample Output | clear ipv6 neighbors on page 302 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

`clear ipv6 neighbors` `user@host> clear ipv6 neighbors`

clear ipv6 router-advertisement

| | |
|---------------------------------|---|
| Syntax | clear ipv6 router-advertisement <interface <i>interface</i> > <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Clear IPv6 router advertisement counters. |
| Options | <p>none—Clear IPv6 router advertisement counters for all interfaces.</p> <p>interface <i>interface</i>—(Optional) Clear IPv6 router advertisement counters for the specified interface.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • show ipv6 router-advertisement on page 306 |
| List of Sample Output | clear ipv6 router-advertisement on page 303 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

```
clear ipv6 router-advertisement      user@host> clear ipv6 router-advertisement
```

show ipv6 neighbors

| | |
|---------------------------------|--|
| Syntax | show ipv6 neighbors |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.3 for EX Series switches. Command introduced in Junos OS Release 12.2 for the QFX Series. |
| Description | Display information about the IPv6 neighbor cache. |
| Options | This command has no options. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">• clear ipv6 neighbors on page 302 |
| List of Sample Output | show ipv6 neighbors on page 305 |
| Output Fields | Table 76 on page 304 describes the output fields for the show ipv6 neighbors command. Output fields are listed in the approximate order in which they appear. |

Table 76: show ipv6 neighbors Output Fields

| Field Name | Field Description |
|-------------------|--|
| IPv6 Address | Name of the IPv6 interface. |
| Linklayer Address | Link-layer address. |
| State | State of the link: up , down , incomplete , reachable , stale , or unreachable . |
| Exp | Number of seconds until the entry expires. |
| Rtr | Whether the neighbor is a routing device: yes or no . |
| Secure | Whether this entry was created using the Secure Neighbor Discovery (SEND) protocol: yes or no . |
| Interface | Name of the interface. |

Sample Output

show ipv6 neighbors

user@host> show ipv6 neighbors

| IPv6 Address | Linklayer Address | State | Exp | Rtr | Secure |
|-----------------------------|-------------------|-------|-----|-----|--------|
| Interface | | | | | |
| 2001:db8:0:1:2a0:a514:0:24c | 00:05:85:8f:c8:bd | stale | 546 | yes | no |
| fe-1/2/0.1 | | | | | |
| fe80::2a0:a514:0:24c | 00:05:85:8f:c8:bd | stale | 258 | yes | no |
| fe-1/2/0.1 | | | | | |
| fe80::2a0:a514:0:64c | 00:05:85:8f:c8:bd | stale | 111 | yes | no |
| fe-1/2/1.5 | | | | | |
| fe80::2a0:a514:0:a4c | 00:05:85:8f:c8:bd | stale | 327 | yes | no |
| fe-1/2/2.9 | | | | | |

show ipv6 router-advertisement

| | |
|---------------------------------|--|
| Syntax | <pre>show ipv6 router-advertisement <conflicts> <interface <i>interface</i>> <logical-system (all <i>logical-system-name</i>)> <prefix <i>prefix/prefix length</i>></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 12.2 for the QFX Series.</p> |
| Description | Display information about IPv6 router advertisements, including statistics about messages sent and received on interfaces, and information received from advertisements from other routers. |
| Options | <p>none—Display all IPv6 router advertisement information for all interfaces.</p> <p>conflicts—(Optional) Display only the IPv6 router advertisement information that is conflicting.</p> <p>interface <i>interface</i>—(Optional) Display IPv6 router advertisement information for the specified interface.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>prefix <i>prefix/prefix length</i>—(Optional) Display IPv6 router advertisement information for the specified prefix.</p> |
| Additional Information | The display identifies conflicting information by enclosing the value the router is advertising in brackets. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • clear ipv6 router-advertisement on page 303 |
| List of Sample Output | show ipv6 router-advertisement on page 308 show ipv6 router-advertisement conflicts on page 308 show ipv6 router-advertisement prefix on page 308 |
| Output Fields | Table 77 on page 306 describes the output fields for the show ipv6 router-advertisement command. Output fields are listed in the approximate order in which they appear. |

Table 77: show ipv6 router-advertisement Output Fields

| Field Name | Field Description |
|---------------------|---|
| Interface | Name of the interface. |
| Advertisements sent | Number of router advertisements sent and the elapsed time since they were sent. |

Table 77: show ipv6 router-advertisement Output Fields (*continued*)

| Field Name | Field Description |
|--------------------------------|---|
| Solicits received | Number of solicitation messages received. |
| Advertisements received | Number of router advertisements received. |
| Advertisements from | Names of interfaces from which router advertisements have been received and the elapsed time since the last one was received. |
| Managed | Managed address configuration flag: 0 (stateless) or 1 (stateful). |
| Other configuration | Other stateful configuration flag: 0 (stateless) or 1 (stateful). |
| Reachable time | Time that a node identifies a neighbor as reachable after receiving a reachability confirmation, in milliseconds. |
| Default lifetime | Default lifetime, in seconds: from 0 seconds to 18.2 hours. A setting of 0 indicates that the router is not a default router. |
| Retransmit timer | Time between retransmitted Neighbor Solicitation messages, in milliseconds. |
| Current hop limit | Configured current hop limit. |
| Prefix | Name and length of the prefix. |
| Valid lifetime | How long the prefix remains valid for onlink determination. |
| Preferred lifetime | How long the prefix generated by stateless autoconfiguration remains preferred. |
| On link | Onlink flag: 0 (not onlink) or 1 (onlink). |
| Autonomous | Autonomous address configuration flag: 0 (not autonomous) or 1 (autonomous). |

Sample Output

**show ipv6
router-advertisement**

```
user@host> show ipv6 router-advertisement
Interface: fe-0/1/1.0
  Advertisements sent: 0
  Solicits received: 0
  Advertisements received: 0
Interface: fxp0.0
  Advertisements sent: 0
  Solicits received: 0
  Advertisements received: 1
  Advertisement from fe80::2d0:b7ff:fe1e:7b0e, heard 00:00:13 ago
  Managed: 0
  Other configuration: 0 [1]
    Reachable time: 0 ms
    Default lifetime: 1800 sec
    Retransmit timer: 0 ms
    Current hop limit: 64
```

**show ipv6
router-advertisement
conflicts**

```
user@host> show ipv6 router-advertisement conflicts
Interface: fxp0.0
  Advertisement from fe80::2d0:b7ff:fe1e:7b0e, heard 00:01:08 ago
  Other configuration: 0 [1]
```

**show ipv6
router-advertisement
prefix**

```
user@host> show ipv6 router-advertisement prefix 8040::/16
Interface: fe-0/1/3.0
  Advertisements sent: 3, last sent 00:04:11 ago
  Solicits received: 0
  Advertisements received: 3
  Advertisement from fe80::290:69ff:fe9a:5403, heard 00:00:05 ago
  Managed: 0
  Other configuration: 0
  Reachable time: 0 ms
  Default lifetime: 180 sec [1800 sec]
  Retransmit timer: 0 ms
  Current hop limit: 64
  Prefix: 8040:1::/64
    Valid lifetime: 2592000 sec
    Preferred lifetime: 604800 sec
    On link: 1
    Autonomous: 1
```

CHAPTER 7

IS-IS Operational Mode Commands

Table 78 on page 309 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot the Intermediate System-to-Intermediate System (IS-IS) protocol. Commands are listed in alphabetical order.

Table 78: IS-IS Operational Mode Commands

| Task | Command |
|--|---|
| Remove adjacencies. | <code>clear isis adjacency</code> |
| Remove database entries. | <code>clear isis database</code> |
| Reset IS-IS dynamic overload bit. | <code>clear isis overload</code> |
| Set IS-IS traffic statistics to zero. | <code>clear isis statistics</code> |
| Display adjacent routers. | <code>show isis adjacency</code> |
| Display authentication statistics. | <code>show isis authentication</code> |
| Display information about the level of backup coverage available for protected routes. | <code>show isis backup coverage</code> |
| Display information about MPLS LSPs designated as backup paths. | <code>show isis backup label-switched-path</code> |
| Display SPF calculations for backup paths. | <code>show isis backup spf results</code> |
| Display IS-IS context identifier information. | <code>show isis context-identifier</code> |
| Display database entries. | <code>show isis database</code> |
| Display hostname mapping. | <code>show isis hostname</code> |
| Display the status of interfaces on which IS-IS is running. | <code>show isis interface</code> |
| Display IS-IS overview information. | <code>show isis overview</code> |
| Display IS-IS routing table entries. | <code>show isis route</code> |

Table 78: IS-IS Operational Mode Commands (*continued*)

| Task | Command |
|-----------------------------------|-----------------------------------|
| Display SPF calculations. | <code>show isis spf</code> |
| Display IS-IS traffic statistics. | <code>show isis statistics</code> |



NOTE: For more IS-IS-related commands, such as `show route protocol`, `show route instance`, and `show route table`, see Protocol-Independent Routing Operational Mode Commands. For information about monitoring Bidirectional Forwarding Detection (BFD) sessions for IS-IS clients, see BFD Operational Mode Commands. For information about how to configure IS-IS, see the *Junos Routing Protocols Configuration Guide*.



NOTE: In IS-IS command output, the CLI displays the system ID numerically by default. To display the hostname instead, include the `static-host-mapping` statement at the `[edit system]` hierarchy level of the configuration.

clear isis adjacency

| | |
|---|---|
| Syntax | clear isis adjacency <instance <i>instance-name</i> > <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)> < <i>neighbor</i> > |
| Syntax (EX Series Switches and QFX Series) | clear isis adjacency <instance <i>instance-name</i> > <interface <i>interface-name</i> > < <i>neighbor</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series. |
| Description | Remove entries from the IS-IS adjacency database. |
| Options | <p>none—Remove all entries from the adjacency database.</p> <p>instance <i>instance-name</i>—(Optional) Clear all adjacencies for the specified routing instance only.</p> <p>interface <i>interface-name</i>—(Optional) Clear all adjacencies for the specified interface only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><i>neighbor</i>—(Optional) Clear adjacencies for the specified neighbor only.</p> |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none"> • show isis adjacency on page 319 |
| List of Sample Output | clear isis adjacency on page 312 |
| Output Fields | See show isis adjacency for an explanation of output fields. |

Sample Output

clear isis adjacency


The following sample output displays IS-IS adjacency database information before and after the **clear isis adjacency** command is entered:

```
user@host> show isis adjacency
IS-IS adjacency database:
Interface      System          L State          Hold (secs) SNPA
so-1/0/0.0     karakul         3 Up             26
so-1/1/3.0     1921.6800.5080 3 Up             23
so-5/0/0.0     1921.6800.5080 3 Up             19
```

```
user@host> clear isis adjacency karakul
```

```
user@host> show isis adjacency
IS-IS adjacency database:
Interface      System          L State          Hold (secs) SNPA
so-1/0/0.0     karakul         3 Initializing   26
so-1/1/3.0     1921.6800.5080 3 Up             24
so-5/0/0.0     1921.6800.5080 3 Up             21
```

clear isis database

| | |
|---|--|
| Syntax | clear isis database <entries> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> <purge> |
| Syntax (EX Series Switches and QFX Series) | clear isis database <entries> <instance <i>instance-name</i> > <purge> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. purge option introduced in Junos OS Release 9.0. Command introduced in Junos OS Release 12.1 for the QFX Series. |
| Description | Remove the entries from the IS-IS link-state database, which contains prefixes and topology information. You can also use purge with any of the options to initiate a network-wide purge of link-state PDUs rather than the local deletion of entries from the IS-IS link-state database. |
| <div>  <p>CAUTION: In a production network, the purge command option might cause short-term network-wide traffic disruptions.</p> </div> | |
| Options | <p>none—Remove all entries from the IS-IS link-state database for all routing instances.</p> <p>entries—(Optional) Name of the database entry.</p> <p>instance <i>instance-name</i>—(Optional) Clear all entries for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>purge—(Optional) Discard all entries in the IS-IS link-state database.</p> |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none"> • show isis database on page 335 |
| List of Sample Output | clear isis database on page 314 |
| Output Fields | See show isis database for an explanation of output fields. |

Sample Output

clear isis database

The following sample output displays IS-IS link-state database information before and after the **clear isis database** command is entered:

```
user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                Sequence Checksum Lifetime (secs)
crater.00-00          0x12   0x84dd           1139
    1 LSPs
IS-IS level 2 link-state database:
LSP ID                Sequence Checksum Lifetime (secs)
crater.00-00          0x19   0xe92c           1134
badlands.00-00        0x16   0x1454           985
carlsbad.00-00        0x33   0x220b          1015
ranier.00-00          0x2e   0xfc31          1007
1921.6800.5066.00-00  0x11   0x7313           566
1921.6800.5067.00-00  0x14   0xd9d4           939
    6 LSPs
```

```
user@host> clear isis database
```

```
user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                Sequence Checksum Lifetime (secs)

IS-IS level 2 link-state database:
LSP ID                Sequence Checksum Lifetime (secs)
```

clear isis overload

| | |
|---|---|
| Syntax | clear isis overload <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switches and QFX Series) | clear isis overload <instance <i>instance-name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series. |
| Description | <p>Reset the IS-IS dynamic overload bit. This command can appear to not work, continuing to display overload after execution. The bit is reset only if the root cause is corrected by configuration remotely or locally.</p> <p>When other routers detect that the overload bit is set, they do not use this routing device for transit traffic, but they do use it for packets destined to the overloaded routing device's directly connected networks and IP prefixes.</p> |
| Options | <p>none—Reset the IS-IS dynamic overload bit.</p> <p>instance <i>instance-name</i>—(Optional) Reset the IS-IS dynamic overload bit for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none"> • show isis database on page 335 |
| List of Sample Output | clear isis overload on page 316 |
| Output Fields | See show isis database for an explanation of output fields. |

Sample Output

clear isis overload

The following sample output displays IS-IS database information before and after the **clear isis overload** command is entered:

```
user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                Sequence Checksum Lifetime Attributes
pro3-c.00-00          0x4    0x10db    1185 L1 L2 Overload

  1 LSPs
IS-IS level 2 link-state database:
LSP ID                Sequence Checksum Lifetime Attributes
pro3-c.00-00          0x5    0x429f    1185 L1 L2 Overload

pro2-a.00-00          0x91e   0x2589     874 L1 L2
pro2-a.02-00          0x1     0xcbc     874 L1 L2
  3 LSPs

user@host> clear isis overload

user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                Sequence Checksum Lifetime Attributes
pro3-c.00-00          0xa    0x429e    1183 L1 L2
  1 LSPs

IS-IS level 2 link-state database:
LSP ID                Sequence Checksum Lifetime Attributes
pro3-c.00-00          0xc    0x9c39    1183 L1 L2
pro2-a.00-00          0x91e   0x2589     783 L1 L2
pro2-a.02-00          0x1     0xcbc     783 L1 L2
  3 LSPs
```

clear isis statistics

| | |
|---|---|
| Syntax | clear isis statistics <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switches and QFX Series) | clear isis statistics <instance <i>instance-name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series. |
| Description | Set statistics about IS-IS traffic to zero. |
| Options | <p>none—Set IS-IS traffic statistics to zero for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Set IS-IS traffic statistics to zero for the specified routing instance only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • show isis statistics on page 360 |
| List of Sample Output | clear isis statistics on page 318 |
| Output Fields | See show isis statistics for an explanation of output fields. |

Sample Output

clear isis statistics

The following sample output displays IS-IS statistics before and after the **clear isis statistics** command is entered:

```
user@host> show isis statistics
```

IS-IS statistics for merino:

| PDU type | Received | Processed | Drops | Sent | Rexmit |
|----------|----------|-----------|-------|--------|--------|
| LSP | 12793 | 12793 | 0 | 8666 | 719 |
| IIH | 116751 | 116751 | 0 | 118834 | 0 |
| CSNP | 203956 | 203956 | 0 | 204080 | 0 |
| PSNP | 7356 | 7350 | 6 | 8635 | 0 |
| Unknown | 0 | 0 | 0 | 0 | 0 |
| Totals | 340856 | 340850 | 6 | 340215 | 719 |

Total packets received: 340856 Sent: 340934

SNP queue length: 0 Drops: 0
LSP queue length: 0 Drops: 0

SPF runs: 1064
Fragments rebuilt: 1087
LSP regenerations: 436
Purges initiated: 0

```
user@host> clear isis statistics
```

```
user@host> show isis statistics
```

IS-IS statistics for merino:

| PDU type | Received | Processed | Drops | Sent | Rexmit |
|----------|----------|-----------|-------|------|--------|
| LSP | 0 | 0 | 0 | 0 | 0 |
| IIH | 3 | 3 | 0 | 3 | 0 |
| CSNP | 2 | 2 | 0 | 4 | 0 |
| PSNP | 0 | 0 | 0 | 0 | 0 |
| Unknown | 0 | 0 | 0 | 0 | 0 |
| Totals | 5 | 5 | 0 | 7 | 0 |

Total packets received: 5 Sent: 7

SNP queue length: 0 Drops: 0
LSP queue length: 0 Drops: 0

SPF runs: 0
Fragments rebuilt: 0
LSP regenerations: 0
Purges initiated: 0

show isis adjacency

| | | |
|---|--|--|
| Syntax | <pre>show isis adjacency <system-id> <brief detail extensive> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)></pre> | |
| Syntax (EX Series Switches and QFX Series) | <pre>show isis adjacency <system-id> <brief detail extensive> <instance <i>instance-name</i>></pre> | |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p> | |
| Description | Display information about IS-IS neighbors. | |
| Options | <p>none—Display standard information about IS-IS neighbors for all routing instances.</p> <p><i>system id</i>—(Optional) Display information about IS-IS neighbors for the specified intermediate system.</p> <p>brief detail extensive—(Optional) Display standard information about IS-IS neighbors with the specified level of output.</p> <p>instance <i>instance-name</i>—(Optional) Display information about IS-IS neighbors for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Display information about IS-IS neighbors for all logical systems or for a particular logical system.</p> | |
| Required Privilege Level | view | |
| Related Documentation | <ul style="list-style-type: none"> • clear isis adjacency on page 311 | |
| List of Sample Output | <p>show isis adjacency on page 322</p> <p>show isis adjacency brief on page 322</p> <p>show isis adjacency detail on page 322</p> <p>show isis adjacency extensive on page 322</p> | |
| Output Fields | <p>Table 79 on page 319 describes the output fields for the show isis adjacency command. Output fields are listed in the approximate order in which they appear.</p> | |

Table 79: show isis adjacency Output Fields

| Field Name | Field Description | Level of Output |
|------------|--|-----------------|
| Interface | Interface through which the neighbor is reachable. | All levels |

Table 79: show isis adjacency Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---|---|-------------------------|
| System | System identifier (sysid), displayed as a name, if possible. | brief |
| L or Level | Level: <ul style="list-style-type: none"> • 1—Level 1 only • 2—Level 2 only • 3—Level 1 and Level 2 An exclamation point (!) preceding the level number indicates that the adjacency is missing an IP address. | All levels |
| State | State of the adjacency: Up , Down , New , One-way , Initializing , or Rejected . | All levels |
| Hold (secs) | Remaining hold time of the adjacency. | brief |
| SNPA | Subnetwork point of attachment (MAC address of the next hop). | brief |
| Expires in | How long until the adjacency expires, in seconds. | detail |
| Priority | Priority to become the designated intermediate system. | detail extensive |
| Up/Down transitions | Count of adjacency status changes from Up to Down or from Down to Up . | detail |
| Last transition | Time of the last Up/Down transition. | detail |
| Circuit type | Bit mask of levels on this interface: 1=Level 1 router; 2=Level 2 router; 3=both Level 1 and Level 2 router. | detail |
| Speaks | Protocols supported by this neighbor. | detail extensive |
| MAC address | MAC address of the interface. | detail extensive |
| Topologies | Supported topologies. | detail extensive |
| Restart capable | Whether a neighbor is capable of graceful restart: Yes or No . | detail extensive |
| Adjacency advertisement: Advertise | This router has signaled to advertise this interface to its neighbors in their link-state PDUs. | detail extensive |
| Adjacency advertisement: Suppress | This neighbor has signaled not to advertise the interface in the router's outbound link-state PDUs. | detail extensive |
| IP addresses | IP address of this neighbor. | detail extensive |

Table 79: show isis adjacency Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|----------------|--|-----------------|
| Transition log | <p>List of recent transitions, including:</p> <ul style="list-style-type: none"> • When—Time at which an IS-IS adjacency transition occurred. • State—Current state of the IS-IS adjacency (up, down, or rejected). <ul style="list-style-type: none"> • Up—Adjacency is up and operational. • Down—Adjacency is down and not available. • Rejected—Adjacency has been rejected. • Event—Type of transition that occurred. <ul style="list-style-type: none"> • Seenself—Possible routing loop has been detected. • Interface down—IS-IS interface has gone down and is no longer available. • Error—Adjacency error. • Down reason—Reason that an IS-IS adjacency is down: <ul style="list-style-type: none"> • 3-Way Handshake Failed—Connection establishment failed. • Address Mismatch—Address mismatch caused link failure. • Aged Out—Link expired. • ISO Area Mismatch—IS-IS area mismatch caused link failure. • Bad Hello—Unacceptable hello message caused link failure. • BFD Session Down—Bidirectional failure detection caused link failure. • Interface Disabled—IS-IS interface is disabled. • Interface Down—IS-IS interface is unavailable. • Interface Level Disabled—IS-IS level is disabled. • Level Changed—IS-IS level has changed on the adjacency. • Level Mismatch—Levels on adjacency are not compatible. • MPLS LSP Down—Label-switched path (LSP) is unavailable. • MT Topology Changed—IS-IS topology has changed. • MT Topology Mismatch—IS-IS topology is mismatched. • Remote System ID Changed—Adjacency peer system ID changed. • Protocol Shutdown—IS-IS protocol is disabled. • CLI Command—Adjacency brought down by user. • Unknown—Unknown. | extensive |

Sample Output

show isis adjacency

```
user@host> show isis adjacency
Interface          System      L State      Hold (secs) SNPA
at-2/3/0.0         ranier      3 Up          23
```

show isis adjacency brief

The output for the **show isis adjacency brief** command is identical to that for the **show isis adjacency** command. For sample output, see [show isis adjacency on page 322](#).

show isis adjacency detail

```
user@host> show isis adjacency detail
ranier
  Interface: at-2/3/0.0, Level: 3, State: Up, Expires in 21 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 00:01:09 ago
  Circuit type: 3, Speaks: IP, IPv6
  Topologies: Unicast
  Restart capable: Yes
  IP addresses: 11.1.1.2
```

show isis adjacency extensive

```
user@host> show isis adjacency extensive
ranier
  Interface: at-2/3/0.0, Level: 3, State: Up, Expires in 22 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 00:01:16 ago
  Circuit type: 3, Speaks: IP, IPv6
  Topologies: Unicast
  Restart capable: Yes
  IP addresses: 11.1.1.2
  Transition log:
    When      State      Event      Down reason
    Wed Nov  8 21:24:25  Up        Seenself
```

show isis authentication

| | |
|---|--|
| Syntax | show isis authentication <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switches and QFX Series) | show isis authentication <instance <i>instance-name</i> > |
| Release Information | Command introduced in Junos OS Release 7.5. Command introduced in Junos OS Release 9.0 for EX Series switches. Support for hitless authentication key rollover introduced in Junos OS Release 11.2. Command introduced in Junos OS Release 12.1 for the QFX Series. |
| Description | Display information about IS-IS authentication. |
| Options | <p>none—Display information about IS-IS authentication.</p> <p>instance <i>instance-name</i>—(Optional) Display IS-IS authentication for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show isis authentication on page 324 show isis authentication (With Hitless Authentication Key Rollover Configured) on page 324 |
| Output Fields | Table 80 on page 323 describes the output fields for the show isis authentication command. Output fields are listed in the approximate order in which they appear. |

Table 80: show isis authentication Output Fields

| Field Name | Field Description |
|------------------|---|
| Interface | Interface name. |
| Level | IS-IS level. |
| IIH Auth | IS-IS Hello (IIH) packet authentication type. Displays the name of the active keychain if hitless authentication key rollover is configured. |
| CSN Auth | Complete sequence number authentication type. |
| PSN Auth | Partial sequence number authentication type. |

Table 80: show isis authentication Output Fields (*continued*)

| Field Name | Field Description |
|------------------------------|---|
| L1 LSP Authentication | Layer 1 link-state PDU authentication type. |
| L2 LSP Authentication | Layer 2 link-state PDU authentication type. |

Sample Output

show isis authentication

```

user@host> show isis authentication
Interface          Level IIH Auth  CSN Auth  PSN Auth
at-2/3/0.0         1      Simple    Simple    Simple
                   2      MD5       MD5       MD5

L1 LSP Authentication: Simple
L2 LSP Authentication: MD5

```

show isis authentication (With Hitless Authentication Key Rollover Configured)

```

user@host> show isis authentication
Interface          Level IIH Auth  CSN Auth  PSN Auth
so-0/1/3.0         2      hakrhello MD5       MD5

L2 LSP Authentication: MD5

```

show isis backup coverage

| | |
|---|---|
| Syntax | <pre>show isis backup coverage <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)></pre> |
| Syntax (EX Series Switches and QFX Series) | <pre>show isis backup coverage <instance <i>instance-name</i>></pre> |
| Release Information | <p>Command introduced in Junos OS Release 9.5.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p> |
| Description | Display information about the level of backup coverage available. |
| Options | <p>none—Display information about the level of backup coverage available for all the nodes and prefixes in the network.</p> <p>instance <i>instance-name</i>—(Optional) Display information about the level of backup coverage for a specific IS-IS routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> Example: Configuring Link and Node Protection for IS-IS Routes show isis backup label-switched-path on page 327 |
| List of Sample Output | show isis backup coverage on page 326 |
| Output Fields | <p>Table 81 on page 325 lists the output fields for the show isis backup coverage command. Output fields are listed in the approximate order in which they appear.</p> |

Table 81: show isis backup coverage Output Fields

| Field Name | Field Description |
|-----------------|--|
| Topology | Type of topology or address family: IPV4 Unicast or IPV6 Unicast . |
| Level | IS-IS level: <ul style="list-style-type: none"> 1—Level 1 2—Level 2 |
| Node | By topology, the percentage of all routes configured on the node that are protected through backup coverage. |

Table 81: show isis backup coverage Output Fields (*continued*)

| Field Name | Field Description |
|------------|--|
| IPv4 | Percentage of IPv4 unicast routes that are protected through backup coverage. |
| IPv6 | Percentage of IPv6 unicast routes that are protected through backup coverage. |
| CLNS | Percentage of Connectionless Network Service (CLNS) routes that are protected through backup coverage. |

Sample Output

**show isis backup
coverage**

```
user@host> show isis backup coverage
```

```
Backup Coverage:
```

| Topology | Level | Node | IPv4 | IPv6 | CLNS |
|--------------|-------|--------|--------|-------|-------|
| IPV4 Unicast | 2 | 28.57% | 22.22% | 0.00% | 0.00% |
| IPV6 Unicast | 2 | 0.00% | 0.00% | 0.00% | 0.00% |

show isis backup label-switched-path

| | |
|---|---|
| Syntax | show isis backup label-switched-path <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switches and QFX Series) | show isis backup label-switched-path |
| Release Information | Command introduced in Junos OS Release 9.5. Command introduced in Junos OS Release 9.5 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series. |
| Description | Display information about MPLS label-switched-paths (LSPs) designated as backup routes for IS-IS routes. |
| Options | none —Display information about MPLS LSPs designated as backup routes for IS-IS routes. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> Example: Configuring Link and Node Protection for IS-IS Routes show isis backup coverage on page 325 |
| List of Sample Output | show isis backup label-switched-path on page 328 |
| Output Fields | Table 82 on page 327 lists the output fields for the show isis backup label-switched-path command. Output fields are listed in the approximate order in which they appear. |

Table 82: show isis backup label-switched-path Output Fields

| Field Name | Field Description |
|-------------------------|---|
| Backup MPLS LSPs | List of MPLS LSPs designated as backup paths for IS-IS routes. |
| Egress | IP address of the egress routing device for the LSP. |
| Status | State of the LSP: <ul style="list-style-type: none"> Up—The router can detect RSVP hello messages from the neighbor. Down—The router has received one of the following indications: <ul style="list-style-type: none"> Communication failure from the neighbor. Communication from IGP that the neighbor is unavailable. Change in the sequence numbers in the RSVP hello messages sent by the neighbor. Deleted—LSP is no longer available as a backup path. |

Table 82: show isis backup label-switched-path Output Fields (*continued*)

| Field Name | Field Description |
|-------------|--|
| Last change | Time elapsed since the neighbor state changed either from up to down or from down to up. The format is <i>hh:mm:ss</i> . |
| TE-metric | Configured traffic engineering metric. |
| Metric | Configured metric. |

Sample Output

show isis backup
label-switched-path

```
user@host> show isis backup label-switched-path
Backup MPLS LSPs:
f-to-g, Egress: 192.168.1.4, Status: up, Last change: 06:12:03
TE-metric: 9, Metric: 0
```

show isis backup spf results

| | |
|------------------------------------|--|
| Syntax | <pre>show isis backup spf results <instance <i>instance-name</i>> <level (1 2)> <logical-system (all <i>logical-system-name</i>)> <no-coverage> <topology (ipv4-unicast ipv6-multicast ipv6-unicast unicast)></pre> |
| Syntax (EX Series Switches) | <pre>show isis backup spf results <instance <i>instance-name</i>> <level (1 2)> <no-coverage> <topology (ipv4-unicast unicast)></pre> |
| Release Information | Command introduced in Junos OS Release 9.5. |
| Description | Display information about IS-IS shortest-path-first (SPF) calculations for backup paths. |
| Options | <p>none—Display information about IS-IS SPF calculations for all backup paths for all destination nodes.</p> <p>instance <i>instance-name</i>—(Optional) Display SPF calculations for backup paths for the specified routing instance.</p> <p>level (1 2)—(Optional) Display SPF calculations for the backup paths for the specified IS-IS level.</p> <p>logical-system <i>logical-system-name</i>—(Optional) Display SPF calculations for the backup paths for all logical systems or on a particular logical system.</p> <p>no-coverage—(Optional) Display SPF calculations only for destinations that do not have backup coverage.</p> <p>topology (ipv4-multicast ipv6-multicast ipv6-unicast unicast)—(Optional) Display SPF calculations for backup paths for the specified topology only.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • Example: Configuring Link and Node Protection for IS-IS Routes • show isis backup coverage on page 325 |
| List of Sample Output | <p>show isis backup spf results on page 331</p> <p>show isis backup spf results no-coverage on page 332</p> |
| Output Fields | <p>Table 83 on page 330 lists the output fields for the show isis backup spf results command. Output fields are listed in the approximate order in which they appear.</p> |

Table 83: show isis backup spf results Output Fields

| Field Name | Field Description |
|-------------------------|--|
| <i>node-name</i> | Name of the destination node. |
| Address | Address of the destination node. |
| Primary next-hop | Interface and name of the node of the primary next hop to reach the destination. |
| Root | Name of the next-hop neighbor. |
| Metric | Metric to the node. |
| Eligible | Indicates that the next-hop neighbor has been designated as a backup path to the destination node. |
| Backup next-hop | Name of the interface of the backup next hop. |
| SNPA | Subnetwork point of attachment (MAC address of the next hop). |
| LSP | Name of the MPLS label-switched path (LSP) designated as a backup path. |
| Not eligible | Indicates that the next-hop neighbor cannot function as a backup path to the destination. |
| Reason | Describes why the next-hop neighbor is designated as Not eligible as a backup path. |

Sample Output

show isis backup spf results

user@host> **show isis backup spf results**

IS-IS level 1 SPF results:
0 nodes

IS-IS level 2 SPF results:
banff.00

```
Primary next-hop: so-6/0/0.0, IPV4, olympic
Primary next-hop: ae0.0, IPV4, camaro, SNPA: 0:90:69:f:67:f0
Primary next-hop: so-6/0/0.0, IPV6, olympic
Primary next-hop: ae0.0, IPV6, camaro, SNPA: 0:90:69:f:67:f0
Root: camaro, Root Metric: 10, Metric: 10
  Not eligible, Reason: Primary next-hop multipath
Root: olympic, Root Metric: 10, Metric: 10
  Not eligible, Reason: Primary next-hop multipath
Root: glacier, Root Metric: 10, Metric: 25
  Not eligible, Reason: Primary next-hop multipath
```

crater.00

```
Primary next-hop: so-6/0/0.0, IPV4, olympic
Primary next-hop: so-6/0/0.0, IPV6, olympic
Root: olympic, Root Metric: 10, Metric: 10
  Not eligible, Reason: Primary next-hop link fate sharing
Root: glacier, Root Metric: 10, Metric: 15
  Eligible, Backup next-hop: as0.0, IPV4, glacier
  Eligible, Backup next-hop: as0.0, IPV6, glacier
Root: camaro, Root Metric: 10, Metric: 20
  Not eligible, Reason: Interface is already covered
```

olympic.00

```
Primary next-hop: so-6/0/0.0, IPV4, olympic
Primary next-hop: so-6/0/0.0, IPV6, olympic
Root: olympic, Root Metric: 10, Metric: 0
  Not eligible, Reason: Primary next-hop link fate sharing
Root: camaro, Root Metric: 10, Metric: 20
  track-item: olympic.00-00
  track-item: kobuk.00-00
  Not eligible, Reason: Path loops
Root: glacier, Root Metric: 10, Metric: 20
  track-item: olympic.00-00
  track-item: kobuk.00-00
  Not eligible, Reason: Path loops
```

camaro.00

```
Primary next-hop: ae0.0, IPV4, camaro, SNPA: 0:90:69:f:67:f0
Primary next-hop: ae0.0, IPV6, camaro, SNPA: 0:90:69:f:67:f0
Root: camaro, Root Metric: 10, Metric: 0
  Not eligible, Reason: Primary next-hop link fate sharing
Root: glacier, Root Metric: 10, Metric: 20
  track-item: camaro.00-00
  track-item: kobuk.00-00
  Not eligible, Reason: Path loops
Root: olympic, Root Metric: 10, Metric: 20
  track-item: camaro.00-00
  track-item: kobuk.00-00
  Not eligible, Reason: Path loops
```

glacier.00

```
Primary next-hop: as0.0, IPV4, glacier
Primary next-hop: as0.0, IPV6, glacier
Root: glacier, Root Metric: 10, Metric: 0
  Not eligible, Reason: Primary next-hop link fate sharing
```

```

Root: camaro, Root Metric: 10, Metric: 20
  track-item: glacier.00-00
  track-item: kobuk.00-00
  Not eligible, Reason: Path loops
Root: olympic, Root Metric: 10, Metric: 20
  track-item: glacier.00-00
  track-item: kobuk.00-00
  Not eligible, Reason: Path loops
5 nodes

```

show isis backup spf results no-coverage

user@host> show isis backup spf results no-coverage

```

IS-IS level 1 SPF results:
  0 nodes

IS-IS level 2 SPF results:
olympic.00
  Primary next-hop: so-6/0/0.0, IPV4, olympic
  Primary next-hop: so-6/0/0.0, IPV6, olympic
  Root: olympic, Root Metric: 10, Metric: 0
    Not eligible, Reason: Primary next-hop link fate sharing
  Root: camaro, Root Metric: 10, Metric: 20
    track-item: olympic.00-00
    track-item: kobuk.00-00
    Not eligible, Reason: Path loops
  Root: glacier, Root Metric: 10, Metric: 20
    track-item: olympic.00-00
    track-item: kobuk.00-00
    Not eligible, Reason: Path loops
camaro.00
  Primary next-hop: ae0.0, IPV4, camaro, SNPA: 0:90:69:f:67:f0
  Primary next-hop: ae0.0, IPV6, camaro, SNPA: 0:90:69:f:67:f0
  Root: camaro, Root Metric: 10, Metric: 0
    Not eligible, Reason: Primary next-hop link fate sharing
  Root: glacier, Root Metric: 10, Metric: 20
    track-item: camaro.00-00
    track-item: kobuk.00-00
    Not eligible, Reason: Path loops
  Root: olympic, Root Metric: 10, Metric: 20
    track-item: camaro.00-00
    track-item: kobuk.00-00
    Not eligible, Reason: Path loops
glacier.00
  Primary next-hop: as0.0, IPV4, glacier
  Primary next-hop: as0.0, IPV6, glacier
  Root: glacier, Root Metric: 10, Metric: 0
    Not eligible, Reason: Primary next-hop link fate sharing
  Root: camaro, Root Metric: 10, Metric: 20
    track-item: glacier.00-00
    track-item: kobuk.00-00
    Not eligible, Reason: Path loops
  Root: olympic, Root Metric: 10, Metric: 20
    track-item: glacier.00-00
    track-item: kobuk.00-00
    Not eligible, Reason: Path loops
3 nodes

```

show isis context-identifier

| | |
|---------------------------------|--|
| Syntax | show isis context-identifier <brief detail extensive> <identifier name> <instance instance-name> <logical-system (all logical-system-name)> |
| Release Information | Command introduced in Junos OS Release 10.4. |
| Description | Display IS-IS context identifier information. |
| Options | <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>identifier name—(Optional) Display information about the specified context identifier.</p> <p>instance instance-name—(Optional) Display entries for the specified routing instance.</p> <p>logical-system (all logical-system-name)—(Optional) Display the context identifier information for all logical systems or for a particular logical system.</p> |
| Required Privilege Level | View |
| Output Fields | Table 84 on page 333 lists the output fields for the show isis context-identifier command. Output fields are listed in the approximate order in which they appear. |

Table 84: show isis context-identifier Output Fields

| Field Name | Field Description | Level of Output |
|----------------|--|-----------------|
| Context | IPv4 address that defines a protection pair. The context is manually configured on both primary and protector PEs. | detail |
| Owner | Protocol that requires the context. | detail |
| Role | Role of the PE, which is either primary or protector. | detail |
| Primary | Name of the primary PE. | detail |
| Metric | Advertised interior gateway protocol (IGP) metric. | detail |

Sample Output

```
user@host> show isis context-identifier detail
```

```
IS-IS context database:
```

| Context | Owner | Role | Primary | Metric |
|---|-------|---------|---------|--------|
| 2.2.4.3 | MPLS | Primary | pro3-e | 1 |
| Advertiser pro3-e, Router ID 10.255.245.198, Metric 1, Level 1 | | | | |
| Advertiser pro3-e, Router ID 10.255.245.198, Metric 1, Level 2 | | | | |
| Advertiser pro3-c, Router ID 10.255.245.196, Metric 11, Level 2 | | | | |

show isis database

| | |
|---|---|
| Syntax | <pre>show isis database <system-id> <brief detail extensive> <instance <i>instance-name</i>> <level (1 2)> <logical-system (all <i>logical-system-name</i>)></pre> |
| Syntax (EX Series Switches and QFX Series) | <pre>show isis database <system-id> <brief detail extensive> <level (1 2)> <instance <i>instance-name</i>></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p> |
| Description | Display the entries in the IS-IS link-state database, which contains data about PDU packets. |
| Options | <p>none—Display standard information about IS-IS link-state database entries for all routing instances.</p> <p><i>system id</i>—(Optional) Display IS-IS link-state database entries for the specified intermediate system.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>instance <i>instance-name</i>—(Optional) Display IS-IS link-state database entries for the specified routing instance.</p> <p>level (1 2)—(Optional) Display IS-IS link-state database entries for the specified IS-IS level.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Display standard information about IS-IS link-state database entries for all logical systems or for a particular logical system.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • clear isis database on page 313 |
| List of Sample Output | <p>show isis database on page 338</p> <p>show isis database brief on page 338</p> <p>show isis database detail on page 338</p> <p>show isis database extensive on page 339</p> |

Output Fields Table 85 on page 336 describes the output fields for the **show isis database** command. Output fields are listed in the approximate order in which they appear. Fields that contain internal IS-IS information useful only in troubleshooting obscure problems are not described in the table. For more details about these fields, contact your customer support representative.

Table 85: show isis database Output Fields

| Field Name | Field Description | Level of Output |
|------------------------|--|-------------------------|
| Interface name | Name of the interface on which the link-state PDU has been received; always IS-IS for this command. | All levels |
| level | Level of intermediate system: <ul style="list-style-type: none"> 1—Intermediate system routes within an area; when the destination is outside an area, it routes toward a Level 2 system. 2—Intermediate system routes between areas and toward other ASs. | All levels |
| LSP ID | Link-state PDU identifier. | All levels |
| Sequence | Sequence number of the link-state PDU. | All levels |
| Checksum | Checksum value of the link-state PDU. | All levels |
| Lifetime (secs) | Remaining lifetime of the link-state PDU, in seconds. | All levels |
| Attributes | Attributes of the specified database: L1 , L2 , Overload , or Attached (L1 only). | none brief |
| # LSPs | Total number of link-state PDUs in the specified link-state database. | none brief |
| IP prefix | Prefix advertised by this link-state PDU. | detail extensive |
| IS neighbor | IS-IS neighbor of the advertising system. | detail extensive |
| ES neighbor | (J Series routers only) An ES-IS neighbor of the advertising system. | detail extensive |
| IP prefix | IPv4 prefix advertised by this link-state PDU. | detail extensive |
| V6 prefix | IPv6 prefix advertised by this link-state PDU. | detail extensive |
| Metric | Metric of the prefix or neighbor. | detail extensive |
| Header | <ul style="list-style-type: none"> LSP ID—Link state PDU identifier of the header. Length—Header length. Allocated Length—Amount of length available for the header. Router ID—Address of the local routing device. Remaining Lifetime—Remaining lifetime of the link-state PDU, in seconds. | extensive |

Table 85: show isis database Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------|--|------------------|
| Packet | <ul style="list-style-type: none"> • LSP ID—The identifier for the link-state PDU. • Length—Packet length. • Lifetime—Remaining lifetime, in seconds. • Checksum—The checksum of the link-state PDU. • Sequence—The sequence number of the link-state PDU. Every time the link-state PDU is updated, this number increments. • Attributes—Packet attributes. • NLPID—Network layer protocol identifier. • Fixed length—Specifies the set length for the packet. | extensive |
| TLVs | <ul style="list-style-type: none"> • Area Address—Area addresses that the routing device can reach. • Speaks—Supported routing protocols. • IP router id—ID of the routing device (usually the IP address). • IP address—IPv4 address. • Hostname—Assigned name of the routing device. • IP prefix—IP prefix of the routing device. • Metric—IS-IS metric that measures the cost of the adjacency between the originating routing device and the advertised routing device. • IP extended prefix—Extended IP prefix of the routing device. • IS neighbor—Directly attached neighbor's name and metric. • IS extended neighbor—Directly attached neighbor's name, metric, and IP address. | extensive |

Sample Output

show isis database

```
user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                Sequence Checksum Lifetime Attributes
kobuk.00-00           0x3    0x3167    1057 L1 L2
camaro.00-00          0x5    0x770e    1091 L1 L2
ranier.00-00          0x4    0xaa95    1091 L1 L2
glacier.00-00         0x4    0x206f    1089 L1 L2
glacier.02-00         0x1    0xd141    1089 L1 L2
badlands.00-00       0x3    0x87a2    1093 L1 L2
    6 LSPs

IS-IS level 2 link-state database:
LSP ID                Sequence Checksum Lifetime Attributes
kobuk.00-00           0x6    0x8d6b    1096 L1 L2
camaro.00-00          0x9    0x877b    1101 L1 L2
ranier.00-00          0x8    0x855d    1103 L1 L2
glacier.00-00         0x7    0xf892    1098 L1 L2
glacier.02-00         0x1    0xd141    1089 L1 L2
badlands.00-00       0x6    0x562    1105 L1 L2
    6 LSPs
```

show isis database brief

The output for the **show isis database brief** command is identical to that for the **show isis database** command. For sample output, see [show isis database on page 338](#).

show isis database detail

```
user@host> show isis database logical-system CE3 sisira.00-00 detail

IS-IS level 1 link-state database:

sisira.00-00 Sequence: 0x11, Checksum: 0x10fc, Lifetime: 975 secs
  IS neighbor: hemantha-CE3.02                Metric: 10
  ES neighbor: 0015.0015.0015                  Metric: 10 Down
  ES neighbor: 0025.0025.0025                  Metric: 10 Down
  ES neighbor: 0030.0030.0030                  Metric: 10 Down
  ES neighbor: 0040.0040.0040                  Metric: 10 Down
  ES neighbor: sisira                          Metric: 0
  IP prefix: 1.0.0.0/24                        Metric: 10 External Down
  IP prefix: 3.0.0.0/24                        Metric: 10 External Down
  IP prefix: 4.0.0.0/24                        Metric: 10 External Down
  IP prefix: 5.0.0.0/24                        Metric: 10 Internal Up
  IP prefix: 15.15.15.15/32                    Metric: 10 External Down
  IP prefix: 25.25.25.25/32                    Metric: 10 External Down
  IP prefix: 30.30.30.30/32                    Metric: 10 External Down
  IP prefix: 40.40.40.40/32                    Metric: 10 External Down
  IP prefix: 60.60.60.60/32                    Metric: 0 Internal Up

IS-IS level 2 link-state database:

sisira.00-00 Sequence: 0x13, Checksum: 0x69ac, Lifetime: 993 secs
  IS neighbor: hemantha-CE3.02                Metric: 10
  IP prefix: 1.0.0.0/24                        Metric: 10 External Down
  IP prefix: 3.0.0.0/24                        Metric: 10 External Down
  IP prefix: 4.0.0.0/24                        Metric: 10 External Down
  IP prefix: 5.0.0.0/24                        Metric: 10 Internal Up
  IP prefix: 15.15.15.15/32                    Metric: 10 External Down
  IP prefix: 25.25.25.25/32                    Metric: 10 External Down
  IP prefix: 30.30.30.30/32                    Metric: 10 External Down
```

```

IP prefix: 40.40.40.40/32          Metric:      10 External Down
IP prefix: 50.50.50.50/32          Metric:      10 Internal Up
IP prefix: 60.60.60.60/32          Metric:       0 Internal Up
ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0015.0015.0015/152
                                          Metric:      10 External Down
ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0025.0025.0025/152
                                          Metric:      10 External Down
ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0030.0030.0030/152
                                          Metric:      10 External Down
ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0040.0040.0040/152
                                          Metric:      10 External Down
ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0060.0060.0060/152
                                          Metric:       0 Internal Up

```

show isis database extensive

```
user@host> show isis database logical-system CE3 sisira.00-00 extensive
```

IS-IS level 1 link-state database:

```

sisira.00-00 Sequence: 0x11, Checksum: 0x10fc, Lifetime: 970 secs
  IS neighbor: hemantha-CE3.02          Metric:      10
    Two-way fragment: hemantha-CE3.02-00, Two-way first fragment:
hemantha-CE3.02-00
  ES neighbor: 0015.0015.0015          Metric:      10 Down
  ES neighbor: 0025.0025.0025          Metric:      10 Down
  ES neighbor: 0030.0030.0030          Metric:      10 Down
  ES neighbor: 0040.0040.0040          Metric:      10 Down
  ES neighbor: sisira                  Metric:       0
  IP prefix: 1.0.0.0/24                Metric:      10 External Down
  IP prefix: 3.0.0.0/24                Metric:      10 External Down
  IP prefix: 4.0.0.0/24                Metric:      10 External Down
  IP prefix: 5.0.0.0/24                Metric:      10 Internal Up
  IP prefix: 15.15.15.15/32            Metric:      10 External Down
  IP prefix: 25.25.25.25/32            Metric:      10 External Down
  IP prefix: 30.30.30.30/32            Metric:      10 External Down
  IP prefix: 40.40.40.40/32            Metric:      10 External Down
  IP prefix: 60.60.60.60/32            Metric:       0 Internal Up

```

```

Header: LSP ID: sisira.00-00, Length: 336 bytes
  Allocated length: 336 bytes, Router ID: 0.0.0.0
  Remaining lifetime: 970 secs, Level: 1, Interface: 333
  Estimated free bytes: 144, Actual free bytes: 0
  Aging timer expires in: 970 secs
  Protocols: IP, IPv6, CLNS

```

```

Packet: LSP ID: sisira.00-00, Length: 336 bytes, Lifetime : 1198 secs
  Checksum: 0x10fc, Sequence: 0x11, Attributes: 0xb L1 L2 Attached
  NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
  Packet type: 18, Packet version: 1, Max area: 0

```

TLVs:

```

  Area address: 60.0006.80ff.f800.0000.0108.0001 (13)
  Speaks: IP
  Speaks: IPV6
  Speaks: CLNP
  Hostname: sisira
  ES neighbor TLV: Internal, Metric: default 0, Up
    ES: sisira
  IS neighbor: hemantha-CE3.02, Internal, Metric: default 10
  IS extended neighbor: hemantha-CE3.02, Metric: default 10
  ES neighbor TLV: External, Metric: default 10, Down
    ES: 0040.0040.0040

```

```

ES neighbor TLV: External, Metric: default 10, Down
  ES: 0025.0025.0025
ES neighbor TLV: External, Metric: default 10, Down
  ES: 0015.0015.0015
ES neighbor TLV: External, Metric: default 10, Down
  ES: 0030.0030.0030
IP external prefix: 3.0.0.0/24, Internal, Metric: default 10, Down
IP external prefix: 40.40.40.40/32, Internal, Metric: default 10, Down
IP external prefix: 4.0.0.0/24, Internal, Metric: default 10, Down
IP external prefix: 25.25.25.25/32, Internal, Metric: default 10, Down
IP external prefix: 15.15.15.15/32, Internal, Metric: default 10, Down
IP external prefix: 1.0.0.0/24, Internal, Metric: default 10, Down
IP external prefix: 30.30.30.30/32, Internal, Metric: default 10, Down
IP extended prefix: 3.0.0.0/24 metric 10 down
IP extended prefix: 40.40.40.40/32 metric 10 down
IP extended prefix: 4.0.0.0/24 metric 10 down
IP extended prefix: 25.25.25.25/32 metric 10 down
IP extended prefix: 15.15.15.15/32 metric 10 down
IP extended prefix: 1.0.0.0/24 metric 10 down
IP extended prefix: 30.30.30.30/32 metric 10 down
IP prefix: 60.60.60.60/32, Internal, Metric: default 0, Up
IP prefix: 5.0.0.0/24, Internal, Metric: default 10, Up
IP extended prefix: 60.60.60.60/32 metric 0 up
IP extended prefix: 5.0.0.0/24 metric 10 up
No queued transmissions

```

IS-IS level 2 link-state database:

```

sisira.00-00 Sequence: 0x13, Checksum: 0x69ac, Lifetime: 988 secs
  IS neighbor: hemantha-CE3.02 Metric: 10
    Two-way fragment: hemantha-CE3.02-00, Two-way first fragment:
hemantha-CE3.02-00
  IP prefix: 1.0.0.0/24 Metric: 10 External Down
  IP prefix: 3.0.0.0/24 Metric: 10 External Down
  IP prefix: 4.0.0.0/24 Metric: 10 External Down
  IP prefix: 5.0.0.0/24 Metric: 10 Internal Up
  IP prefix: 15.15.15.15/32 Metric: 10 External Down
  IP prefix: 25.25.25.25/32 Metric: 10 External Down
  IP prefix: 30.30.30.30/32 Metric: 10 External Down
  IP prefix: 40.40.40.40/32 Metric: 10 External Down
  IP prefix: 50.50.50.50/32 Metric: 10 Internal Up
  IP prefix: 60.60.60.60/32 Metric: 0 Internal Up
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0015.0015.0015/152
    Metric: 10 External Down
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0025.0025.0025/152
    Metric: 10 External Down
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0030.0030.0030/152
    Metric: 10 External Down
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0040.0040.0040/152
    Metric: 10 External Down
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0060.0060.0060/152
    Metric: 0 Internal Up

Header: LSP ID: sisira.00-00, Length: 427 bytes
  Allocated length: 427 bytes, Router ID: 0.0.0.0
  Remaining lifetime: 988 secs, Level: 2, Interface: 333
  Estimated free bytes: 130, Actual free bytes: 0
  Aging timer expires in: 988 secs
  Protocols: IP, IPv6, CLNS

Packet: LSP ID: sisira.00-00, Length: 427 bytes, Lifetime : 1198 secs

```

Checksum: 0x69ac, Sequence: 0x13, Attributes: 0x3 L1 L2
 NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
 Packet type: 20, Packet version: 1, Max area: 0

TLVs:

Area address: 60.0006.80ff.f800.0000.0108.0001 (13)
 Speaks: IP
 Speaks: IPV6
 Speaks: CLNP
 Hostname: sisira
 IS neighbor: hemantha-CE3.02, Internal, Metric: default 10
 IS extended neighbor: hemantha-CE3.02, Metric: default 10
 IP external prefix: 3.0.0.0/24, Internal, Metric: default 10, Down
 IP external prefix: 40.40.40.40/32, Internal, Metric: default 10, Down
 IP external prefix: 4.0.0.0/24, Internal, Metric: default 10, Down
 IP external prefix: 25.25.25.25/32, Internal, Metric: default 10, Down
 IP external prefix: 15.15.15.15/32, Internal, Metric: default 10, Down
 IP external prefix: 1.0.0.0/24, Internal, Metric: default 10, Down
 IP external prefix: 30.30.30.30/32, Internal, Metric: default 10, Down
 IP extended prefix: 3.0.0.0/24 metric 10 down
 IP extended prefix: 40.40.40.40/32 metric 10 down
 IP extended prefix: 4.0.0.0/24 metric 10 down
 IP extended prefix: 25.25.25.25/32 metric 10 down
 IP extended prefix: 15.15.15.15/32 metric 10 down
 IP extended prefix: 1.0.0.0/24 metric 10 down
 IP extended prefix: 30.30.30.30/32 metric 10 down
 ISO prefix-neighbor TLV: Internal, Metric: default 0, Up
 Prefix : 60.0006.80ff.f800.0000.0108.0001.0060.0060.0060/152
 ISO prefix-neighbor TLV: External, Metric: default 10, Down
 Prefix : 60.0006.80ff.f800.0000.0108.0001.0040.0040.0040/152
 ISO prefix-neighbor TLV: External, Metric: default 10, Down
 Prefix : 60.0006.80ff.f800.0000.0108.0001.0025.0025.0025/152
 ISO prefix-neighbor TLV: External, Metric: default 10, Down
 Prefix : 60.0006.80ff.f800.0000.0108.0001.0015.0015.0015/152
 ISO prefix-neighbor TLV: External, Metric: default 10, Down
 Prefix : 60.0006.80ff.f800.0000.0108.0001.0030.0030.0030/152
 IP prefix: 60.60.60.60/32, Internal, Metric: default 0, Up
 IP prefix: 5.0.0.0/24, Internal, Metric: default 10, Up
 IP prefix: 50.50.50.50/32, Internal, Metric: default 10, Up
 IP extended prefix: 60.60.60.60/32 metric 0 up
 IP extended prefix: 5.0.0.0/24 metric 10 up
 IP extended prefix: 50.50.50.50/32 metric 10 up

No queued transmissions

show isis hostname

| | |
|---|---|
| Syntax | show isis hostname <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switches and QFX Series) | show isis hostname |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series. |
| Description | Display IS-IS hostname database information. |
| Options | <p>none—Display IS-IS hostname database information.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show isis hostname on page 342 |
| Output Fields | Table 86 on page 342 describes the output fields for the show isis hostname command. Output fields are listed in the approximate order in which they appear. |

Table 86: show isis hostname Output Fields

| Field Name | Field Description |
|------------------|---|
| System Id | System identifier mapped to the hostname. |
| Hostname | Hostname mapped to the system identifier. |
| Type | <p>Type of mapping between system identifier and hostname.</p> <ul style="list-style-type: none"> Dynamic—Hostname mapping determined as described in RFC 2763, <i>Dynamic Hostname Exchange Mechanism for IS-IS</i>. Static—Hostname mapping configured by user. |

Sample Output

```

user@host> show isis hostname
IS-IS hostname database:
System Id      Hostname      Type
1921.6800.4201 isis1         Dynamic
1921.6800.4202 isis2         Static
1921.6800.4203 isis3         Dynamic

```


show isis interface


| | |
|---|--|
| Syntax | <pre>show isis interface <brief detail extensive> <interface-name> <logical-system (all logical-system-name)></pre> |
| Syntax (EX Series Switches and QFX Series) | <pre>show isis interface <brief detail extensive> <interface-name></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p> |
| Description | <p>Display status information about IS-IS-enabled interfaces.</p> |
| | <div>  <p>NOTE: If the configured metric for an IS-IS level is above 63, and the <code>wide-metrics-only</code> statement is not configured, the <code>show isis interface detail</code> command and the <code>show isis interface extensive</code> command display 63 as the metric value for that level. Configure the <code>wide-metrics-only</code> statement to generate metric values greater than 63 on a per IS-IS level basis.</p> <p>The <code>show isis interface</code> command displays the configured metric value for an IS-IS level irrespective of whether is configured or not.</p> </div> |
| Options | <p>none—Display standard information about all IS-IS-enabled interfaces.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>interface-name—(Optional) Display information about the specified interface only.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> Example: Enabling Wide IS-IS Metrics for Traffic Engineering |
| List of Sample Output | <p>show isis interface on page 346</p> <p>show isis interface brief on page 346</p> <p>show isis interface detail on page 346</p> <p>show isis interface extensive on page 346</p> <p>show isis interface extensive (With LDP) on page 346</p> |
| Output Fields | <p>Table 87 on page 344 describes the output fields for the <code>show isis interface</code> command. Output fields are listed in the approximate order in which they appear.</p> |

Table 87: show isis interface Output Fields

| Field Name | Field Description | Level of Output |
|---|--|-------------------|
| <i>interface-name</i> | Name of the interface. | detail |
| Designated router | Routing device selected by other routers that is responsible for sending link-state advertisements that describe the network. Used only on broadcast networks. | detail |
| Index | Interface index assigned by the Junos OS kernel. | detail |
| State | Internal implementation information. | detail |
| Circuit id | Circuit identifier. | detail |
| Circuit type | Circuit type: <ul style="list-style-type: none"> • 1—Level 1 only • 2—Level 2 only • 3—Level 1 and Level 2 | detail |
| LSP interval | Interval between link-state PDUs sent from the interface. | detail |
| CSNP interval | Interval between complete sequence number PDUs sent from the interface. | detail extensive |
| Sysid | System identifier. | detail |
| Interface | Interface through which the adjacency is made. | none brief |
| L or Level | Level: <ul style="list-style-type: none"> • 1—Level 1 only • 2—Level 2 only • 3—Level 1 and Level 2 | All levels |
| CirID | Circuit identifier. | none brief |
| Level 1 DR | Level 1 designated intermediate system. | none brief |
| Level 2 DR | Level 2 designated intermediate system. | none brief |
| L1/L2 Metric | Interface's metric for Level 1 and Level 2. If there is no information, the metric is 0. | none brief |
| Adjacency advertisement: Advertise | This routing device has signaled to advertise this interface to its neighbors in their label-switched paths (LSPs). | detail extensive |
| Adjacency advertisement: Suppress | This neighbor has signaled not to advertise this interface in the routing device's outbound LSPs. | detail extensive |
| Adjacencies | Number of adjacencies established on this interface. | detail |

Table 87: show isis interface Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|----------------------------------|---|-------------------------|
| Priority | Priority value for this interface. | detail |
| Metric | Metric value for this interface. | detail |
| Hello(s) / Hello Interval | Interface's hello interval. | detail extensive |
| Hold(s) / Hold Time | Interface's hold time. | detail extensive |
| Designated Router | Router responsible for sending network link-state advertisements, which describe all the routers attached to the network. | detail |
| Hello padding | Type of hello padding: <ul style="list-style-type: none"> • Adaptive—On point-to-point connections, the hello packets are padded from the initial detection of a new neighbor until the neighbor verifies the adjacency as Up in the adjacency state TLV. If the neighbor does not support the adjacency state TLV, then padding continues. On LAN connections, padding starts from the initial detection of a new neighbor until there is at least one active adjacency on the interface. • Loose—(Default) The hello packet is padded from the initial detection of a new neighbor until the adjacency transitions to the Up state. • Strict—Padding is performed on all interface types and for all adjacency states, and is continuous. | extensive |
| LDP sync state | Current LDP synchronization state: in sync , in holddown , or not supported . | extensive |
| reason | Reason for being in the LDP sync state. | extensive |
| config holdtime | Configured value of the hold timer. | extensive |
| remaining | If the state is not in sync and the hold time is not infinity, then this field displays the remaining hold time in seconds. | extensive |

Sample Output

show isis interface

```
user@host> show isis interface
IS-IS interface database:
Interface          L CirID Level 1 DR      Level 2 DR      L1/L2 Metric
at-2/3/0.0         3   0x1 Point to Point    Point to Point    10/10
lo0.0              0   0x1 Passive          Passive           0/0
```

show isis interface brief

The output for the **show isis interface brief** command is identical to that for the **show isis interface** command. For sample output, see [show isis interface on page 346](#).

show isis interface detail

```
user@host> show isis interface detail
IS-IS interface database:
at-2/3/0.0
  Index: 66, State: 0x6, Circuit id: 0x1, Circuit type: 3
  LSP interval: 100 ms, CSNP interval: 5 s
  Level Adjacencies Priority Metric Hello (s) Hold (s) Designated Router
    1                1       64     10    9.000     27
    2                1       64     10    9.000     27
lo0.0
  Index: 64, State: 0x6, Circuit id: 0x1, Circuit type: 0
  LSP interval: 100 ms, CSNP interval: disabled
  Level Adjacencies Priority Metric Hello (s) Hold (s) Designated Router
    1                0       64      0  Passive
    2                0       64      0  Passive
```

show isis interface extensive

```
user@host> show isis interface extensive
IS-IS interface database:
at-2/3/0.0
  Index: 66, State: 0x6, Circuit id: 0x1, Circuit type: 3
  LSP interval: 100 ms, CSNP interval: 5 s, Loose Hello padding
  Level 1
    Adjacencies: 1, Priority: 64, Metric: 10
    Hello Interval: 9.000 s, Hold Time: 27 s
  Level 2
    Adjacencies: 1, Priority: 64, Metric: 10
    Hello Interval: 9.000 s, Hold Time: 27 s
lo0.0
  Index: 64, State: 0x6, Circuit id: 0x1, Circuit type: 0
  LSP interval: 100 ms, CSNP interval: disabled, Loose Hello padding
  Level 1
    Adjacencies: 0, Priority: 64, Metric: 0
    Passive
  Level 2
    Adjacencies: 0, Priority: 64, Metric: 0
    Passive
```

show isis interface extensive (With LDP)

```
user@host> show isis interface extensive
IS-IS interface database:
so-1/1/2.0
  Index: 114, State: 0x6, Circuit id: 0x1, Circuit type: 2
  LSP interval: 100 ms, CSNP interval: 20 s, Loose Hello padding
  Adjacency advertisement: Advertise
  LDP sync state: in sync, for: 00:01:28, reason: LDP up during config
  config holdtime: 20 seconds
```

Level 2

Adjacencies: 1, Priority: 64, Metric: 11
Hello Interval: 9.000 s, Hold Time: 27 s
IPV4 MulticastMetric: 10
IPV6 UnicastMetric: 10

show isis overview

| | |
|---|---|
| Syntax | show isis overview <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switches and QFX Series) | show isis overview <instance <i>instance-name</i> > |
| Release Information | Command introduced in Junos OS Release 8.5. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series. |
| Description | Display IS-IS overview information. |
| Options | none —Display standard overview information about IS-IS for all routing instances. instance <i>instance-name</i> —(Optional) Display overview information for the specified routing instance. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | view |
| List of Sample Output | show isis overview on page 350 |
| Output Fields | Table 88 on page 348 lists the output fields for the show isis overview command. Output fields are listed in the approximate order in which they appear. |

Table 88: show isis overview Output Fields

| Field Name | Field Description |
|-------------------------|--|
| Instance | IS-IS routing instance. |
| Router ID | Router ID of the routing device. |
| Adjacency holddown | Adjacency holddown capability: enabled or disabled . |
| Maximum Areas | Maximum number of IS-IS areas advertised by the routing device. |
| LSP life time | Lifetime of the link-state PDU, in seconds. |
| Attached bit evaluation | Attached bit capability: enabled or disabled . |
| SPF delay | Delay before performing consecutive shortest-path-first (SPF) calculations. |
| SPF holddown | Delay before performing additional SPF calculations after the maximum number of consecutive SPF calculations is reached. |

Table 88: show isis overview Output Fields (*continued*)

| Field Name | Field Description |
|--------------------------------|---|
| SPF rapid runs | Maximum number of SPF calculations that can be performed in succession before the holddown timer begins. |
| Overload bit at startup is set | Overload bit capability is enabled. |
| Overload high metrics | Overload high metrics capability: enabled or disabled . |
| Overload timeout | Time period after which overload is reset and the time that remains before the timer is set to expire. |
| Traffic engineering | Traffic engineering capability: enabled or disabled . |
| Restart | Graceful restart capability: enabled or disabled . |
| Restart duration | Time period for complete reacquisition of IS-IS neighbors. |
| Helper mode | Graceful restart helper capability: enabled or disabled . |
| Level | IS-IS level: <ul style="list-style-type: none"> • 1—Level 1 information • 2—Level 2 information |
| IPv4 is enabled | IP Protocol version 4 capability is enabled. |
| IPv6 is enabled | IP Protocol version 6 capability is enabled. |
| CLNS is enabled | (J Series routers only) OSI CLNP capability is enabled. |
| Internal route preference | Preference value of internal routes. |
| External route preference | Preference value of external routes. |
| Wide area metrics are enabled | Wide area metrics capability is enabled. |
| Narrow metrics are enabled | Narrow metrics capability is enabled. |

Sample Output

show isis overview

```
user@host> show isis overview
Instance: master
  Router ID: 192.168.1.220
  Adjacency holddown: enabled
  Maximum Areas: 3
  LSP life time: 65535
  Attached bit evaluation: enabled
  SPF delay: 200 msec, SPF holddown: 5000 msec, SPF rapid runs: 3
  Overload bit at startup is set
    Overload high metrics: disabled
    Overload timeout: 300 sec, expires in 295 seconds
  IPv4 is enabled, IPv6 is enabled
  Traffic engineering: enabled
  Restart: Enabled
    Restart duration: 210 sec
    Helper mode: Enabled
Level 1
  Internal route preference: 15
  External route preference: 160
  Wide metrics are enabled, Narrow metrics are enabled
Level 2
  Internal route preference: 18
  External route preference: 165
  Wide metrics are enabled
```


show isis route

| | |
|---|--|
| Syntax | <pre>show isis route <destination> <inet inet6> <instance instance-name> <logical-system (all logical-system-name)> <topology (ipv4-multicast ipv6-multicast ipv6-unicast unicast)></pre> |
| Syntax (EX Series Switches and QFX Series) | <pre>show isis route <destination> <inet inet6> <instance instance-name> <topology (ipv4-multicast ipv6-multicast ipv6-unicast unicast)></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p> |
| Description | Display the routes in the IS-IS routing table. |
| Options | <p>none—Display all routes in the IS-IS routing table for all supported address families for all routing instances.</p> <p>destination—(Optional) Destination address for the route.</p> <p>inet inet6—(Optional) Display inet (IPv4) or inet6 (IPv6) routes, respectively.</p> <p>instance instance-name—(Optional) Display routes for the specified routing instance only.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>topology (ipv4-multicast ipv6-multicast ipv6-unicast unicast)—(Optional) Display routes for the specified topology only, or use unicast to display information, if available, for both IPv4 and IPv6 unicast topologies.</p> |
| Required Privilege Level | view |
| List of Sample Output | <p>show isis route logical-system on page 353</p> <p>show isis route (CLNS) on page 353</p> <p>show isis route on page 353</p> |
| Output Fields | <p>Table 89 on page 351 describes the output fields for the show isis route command. Output fields are listed in the approximate order in which they appear.</p> |

Table 89: show isis route Output Fields

| Field Name | Field Description |
|-----------------|---|
| Current version | Number of the current version of the IS-IS routing table. |

Table 89: show isis route Output Fields (*continued*)

| Field Name | Field Description |
|-------------------|--|
| L1 | Version of Level 1 SPF that was run. |
| L2 | Version of Level 2 SPF that was run. |
| Prefix | Destination of the route. |
| L | IS-IS level: <ul style="list-style-type: none">• 1—Level 1 only• 2—Level 2 only• 3—Level 1 and Level 2 |
| Version | Version of SPF that generated the route. |
| Metric | Metric value associated with the route. |
| Type | Metric type: int (internal) or ext (external). |
| Interface | Interface to the next hop. |
| Via | System identifier of the next hop, displayed as a name if possible. |
| ISO Routes | ISO routing table entries. |
| snpa | MAC address. |

Sample Output

show isis route logical-system

```
user@host> show isis route logical-system ls1
IS-IS routing table                      Current version: L1: 8 L2: 11
Prefix                                L Version Metric Type Interface  Via
10.9.7.0/30                          2      11     20 int  gr-0/2/0.0  h
10.9.201.1/32                        2      11     60 int  gr-0/2/0.0  h
IPv6 Unicast IS-IS routing table        Current version: L1: 9 L2: 11
Prefix                                L Version Metric Type Interface  Via
8009:3::a09:3200/126                 2      11     20 int  gr-0/2/0.0  h
```

show isis route (CLNS)

```
user@host> show isis route
IS-IS routing table                      Current version: L1: 10 L2: 8
IPv4/IPv6 Routes
Prefix                                L Version Metric Type Interface  Via
0.0.0.0/0                            1      10      10 int  fe-0/0/1.0  ISIS.0
ISO Routes
Prefix L Version Metric Type Interface  Via snpa
0/0
1      10      10 int  fe-0/0/1.0  isis.0 0:12:0:34:0:56
47.0005.80ff.f800.0000.0108.0001/104
1      10      0 int
47.0005.80ff.f800.0000.0108.0001.1921.6800.4001/152
1      10      10 int  fe-0/0/1.0  isis.0 0:12:0:34:0:56
47.0005.80ff.f800.0000.0108.0001.1921.6800.4002/152
1      10      20 int  fe-0/0/1.0  isis.0 0:12:0:34:0:56
47.0005.80ff.f800.0000.0108.0002/104
1      10      0 int
47.0005.80ff.f800.0000.0108.0002.1921.6800.4001/152
1      10      10 int  fe-0/0/1.0  isis.0 0:12:0:34:0:56
```

show isis route

```
user@host> show isis route

IS-IS routing table                      Current version: L1: 4 L2: 13
IPv4/IPv6 Routes
-----
Prefix                                L Version Metric Type Interface  NH Via
10.255.71.52/32                      2      13      10 int  ae0.0          IPV4 camaro
10.255.71.238/32                     2      13      20 int  so-6/0/0.0    IPV4 olympic
                                         as0.0          IPV4 glacier
10.255.71.239/32                     2      13      20 int  so-6/0/0.0    IPV4 olympic
                                         ae0.0          IPV4 camaro
10.255.71.242/32                     2      13      10 int  as0.0          IPV4 glacier
10.255.71.243/32                     2      13      10 int  so-6/0/0.0    IPV4 olympic
12.13.0.0/30                        2      13      20 int  so-6/0/0.0    IPV4 olympic
12.15.0.0/30                        2      13      20 int  so-6/0/0.0    IPV4 olympic
13.15.0.0/30                        2      13      30 int  ae0.0          IPV4 camaro
                                         so-6/0/0.0    IPV4 olympic
```

| | | | | | | |
|-------------------------|---|----|----|-----|------------|--------------|
| | | | | | as0.0 | IPv4 glacier |
| 13.16.0.0/30 | 2 | 13 | 25 | int | as0.0 | IPv4 glacier |
| 14.15.0.0/30 | 2 | 13 | 20 | int | ae0.0 | IPv4 camaro |
| 192.2.1.0/30 | 2 | 13 | 30 | int | so-6/0/0.0 | IPv4 olympic |
| | | | | | as0.0 | IPv4 glacier |
| 1eee::/64 | 2 | 13 | 30 | int | so-6/0/0.0 | IPv6 olympic |
| | | | | | as0.0 | IPv6 glacier |
| abcd::10:255:71:52/128 | 2 | 13 | 10 | int | ae0.0 | IPv6 camaro |
| abcd::10:255:71:238/128 | 2 | 13 | 20 | int | so-6/0/0.0 | IPv6 olympic |
| | | | | | as0.0 | IPv6 glacier |
| abcd::10:255:71:239/128 | 2 | 13 | 20 | int | so-6/0/0.0 | IPv6 olympic |
| | | | | | ae0.0 | IPv6 camaro |
| abcd::10:255:71:242/128 | 2 | 13 | 10 | int | as0.0 | IPv6 glacier |
| abcd::10:255:71:243/128 | 2 | 13 | 10 | int | so-6/0/0.0 | IPv6 olympic |

show isis spf

| | |
|------------------------------------|---|
| Syntax | show isis spf (brief log results) <instance <i>instance-name</i> > <level (1 2)> <logical-system (all <i>logical-system-name</i>)> <topology (ipv4-multicast ipv6-multicast ipv6-unicast unicast)> |
| Syntax (EX Series Switches) | show isis spf (brief log results) <instance <i>instance-name</i> > <level (1 2)> <topology (ipv4-multicast ipv6-multicast ipv6-unicast unicast)> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. |
| Description | Display information about IS-IS shortest-path-first (SPF) calculations. |
| Options | <p>brief—Display an overview of SPF calculations.</p> <p>instance <i>instance instance-name</i>—(Optional) Display SPF calculations for the specified routing instance.</p> <p>level (1 2)—(Optional) Display SPF calculations for the specified IS-IS level.</p> <p>log—Display the log of SPF calculations.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>results—Display the results of SPF calculations.</p> <p>topology (ipv4-multicast ipv6-multicast ipv6-unicast unicast)—(Optional) Display SPF calculations for the specified topology only.</p> |
| Required Privilege Level | view |
| List of Sample Output | show isis spf log on page 357 show isis spf results logical-system on page 357 show isis spf results (CLNS) on page 359 |
| Output Fields | Table 90 on page 355 describes the output fields for the show isis spf command. Output fields are listed in the approximate order in which they appear. |

Table 90: show isis spf Output Fields

| Field Name | Field Description |
|------------|----------------------|
| Node | System ID of a node. |
| Metric | Metric to the node. |

Table 90: show isis spf Output Fields (*continued*)

| Field Name | Field Description |
|----------------|---|
| Interface | Interface of the next hop. |
| Via | System ID of the next hop. |
| SNPA | Subnetwork point of attachment (MAC address of the next hop). |
| Start time | (log option only) Time that the SPF computation started. |
| Elapsed (secs) | (log option only) Length of time, in seconds, required to complete the SPF computation. |
| Count | (log option only) Number of times the SPF was triggered. |
| Reason | (log option only) Reason that the SPF computation was completed. |

Sample Output

show isis spf log

```

user@host> show isis spf log logical-system ls1
IS-IS level 1 SPF log:
Start time           Elapsed (secs) Count Reason
Fri Oct 31 12:41:18   0.000069      1 Reconfig
Fri Oct 31 12:41:18   0.000107      3 Updated LSP fix.00-00
Fri Oct 31 12:41:18   0.000050      3 Address change on so-1/2/2.0
Fri Oct 31 12:41:23   0.000033      1 Updated LSP fix.00-00
Fri Oct 31 12:41:28   0.000178      5 New adjacency scat on ge-1/1/0.0
Fri Oct 31 12:41:59   0.000060      1 Updated LSP fix.00-00
Fri Oct 31 12:42:30   0.000161      2 Multi area attachment change
Fri Oct 31 12:56:58   0.000198      1 Periodic SPF
Fri Oct 31 13:10:29   0.000209      1 Periodic SPF
IS-IS level 2 SPF log:

Start time           Elapsed (secs) Count Reason
Fri Oct 31 12:41:18   0.000035      1 Reconfig
Fri Oct 31 12:41:18   0.000047      2 Updated LSP fix.00-00
Fri Oct 31 12:41:18   0.000043      5 Address change on gr-0/2/0.0
Fri Oct 31 12:41:23   0.000022      1 Updated LSP fix.00-00
Fri Oct 31 12:41:59   0.000144      3 New adjacency h on gr-0/2/0.0
Fri Oct 31 12:42:30   0.000257      3 New LSP skag.00-00
Fri Oct 31 12:54:37   0.000195      1 Periodic SPF
Fri Oct 31 12:55:50   0.000178      1 Updated LSP fix.00-00
Fri Oct 31 12:55:55   0.000174      1 Updated LSP h.00-00
Fri Oct 31 12:55:58   0.000176      1 Updated LSP skag.00-00
Fri Oct 31 13:08:14   0.000198      1 Periodic SPF
IPV6 Unicast IS-IS level 1 SPF log:

Start time           Elapsed (secs) Count Reason
Fri Oct 31 12:41:18   0.000028      1 Reconfig
Fri Oct 31 12:41:18   0.000043      3 Updated LSP fix.00-00
Fri Oct 31 12:41:18   0.000112      4 Updated LSP fix.00-00
Fri Oct 31 12:41:23   0.000059      1 Updated LSP fix.00-00
Fri Oct 31 12:41:25   0.000041      1 Updated LSP fix.00-00
Fri Oct 31 12:41:28   0.000103      5 New adjacency scat on ge-1/1/0.0
Fri Oct 31 12:41:59   0.000040      1 Updated LSP fix.00-00
Fri Oct 31 12:42:30   0.000118      2 Multi area attachment change
Fri Oct 31 12:56:08   0.000289      1 Periodic SPF
Fri Oct 31 13:11:07   0.000214      1 Periodic SPF
IPV6 Unicast IS-IS level 2 SPF log:

Start time           Elapsed (secs) Count Reason
Fri Oct 31 12:41:18   0.000027      1 Reconfig
Fri Oct 31 12:41:18   0.000039      2 Updated LSP fix.00-00
Fri Oct 31 12:41:18   0.000049      6 Updated LSP fix.00-00
Fri Oct 31 12:41:23   0.000025      1 Updated LSP fix.00-00
Fri Oct 31 12:41:25   0.000023      1 Updated LSP fix.00-00
Fri Oct 31 12:41:59   0.000087      3 New adjacency h on gr-0/2/0.0
Fri Oct 31 12:42:30   0.000123      3 New LSP skag.00-00
Fri Oct 31 12:55:50   0.000121      1 Updated LSP fix.00-00
Fri Oct 31 12:55:55   0.000121      1 Updated LSP h.00-00
Fri Oct 31 12:55:58   0.000121      1 Updated LSP skag.00-00
Fri Oct 31 13:09:46   0.000201      1 Periodic SPF
...

```

show isis spf results

```

user@host> show isis spf results logical-system ls1
IS-IS level 1 SPF results:

```

logical-system

| Node | Metric | Interface | Via | SNPA |
|---------|--------|---------------|------|------------------|
| scat.00 | 10 | ge-1/1/0.0 | scat | 0:90:69:a6:48:9d |
| | 20 | 10.9.1.0/30 | | |
| fix.02 | 10 | | | |
| fix.00 | 0 | | | |
| | 10 | 10.9.1.0/30 | | |
| | 10 | 10.9.5.0/30 | | |
| | 10 | 10.9.6.0/30 | | |
| | 20 | 10.9.7.0/30 | | |
| | 60 | 10.9.201.1/32 | | |

3 nodes

IS-IS level 2 SPF results:

| Node | Metric | Interface | Via | SNPA |
|---------|--------|---------------|-----|------|
| skag.00 | 20 | gr-0/2/0.0 | h | |
| | 30 | 10.9.7.0/30 | | |
| skag.02 | 20 | gr-0/2/0.0 | h | |
| h.00 | 10 | gr-0/2/0.0 | h | |
| | 20 | 10.9.6.0/30 | | |
| | 20 | 10.9.7.0/30 | | |
| | 60 | 10.9.201.1/32 | | |
| fix.00 | 0 | | | |
| | 10 | 10.9.1.0/30 | | |
| | 10 | 10.9.5.0/30 | | |
| | 10 | 10.9.6.0/30 | | |

4 nodes

IPv6 Unicast IS-IS level 1 SPF results:

| Node | Metric | Interface | Via | SNPA |
|---------|--------|----------------------|------|------------------|
| scat.00 | 10 | ge-1/1/0.0 | scat | 0:90:69:a6:48:9d |
| | | ge-1/1/0.0 | scat | 0:90:69:a6:48:9d |
| | 20 | 8009:1::a09:1400/126 | | |
| fix.02 | 10 | | | |
| fix.00 | 0 | | | |
| | 10 | 8009:1::a09:1400/126 | | |
| | 10 | 8009:2::a09:1e00/126 | | |
| | 20 | 8009:3::a09:3200/126 | | |
| | 10 | 8009:4::a09:2800/126 | | |

3 nodes

IPv6 Unicast IS-IS level 2 SPF results:

| Node | Metric | Interface | Via | SNPA |
|---------|--------|----------------------|-----|------|
| skag.00 | 20 | gr-0/2/0.0 | h | |
| | | gr-0/2/0.0 | h | |
| | 30 | 8009:3::a09:3200/126 | | |
| skag.02 | 20 | gr-0/2/0.0 | h | |
| | | gr-0/2/0.0 | h | |
| h.00 | 10 | gr-0/2/0.0 | h | |
| | | gr-0/2/0.0 | h | |
| | 20 | 8009:3::a09:3200/126 | | |
| | 20 | 8009:4::a09:2800/126 | | |
| fix.00 | 0 | | | |
| | 10 | 8009:1::a09:1400/126 | | |
| | 10 | 8009:2::a09:1e00/126 | | |
| | 10 | 8009:4::a09:2800/126 | | |

4 nodes

Multicast IS-IS level 1 SPF results:

| Node | Metric | Interface | Via | SNPA |
|---------|--------|------------|------|------------------|
| scat.00 | 10 | ge-1/1/0.0 | scat | 0:90:69:a6:48:9d |
| fix.02 | 10 | | | |


```
fix.00      0
  3 nodes
```

Multicast IS-IS level 2 SPF results:

| Node | Metric | Interface | Via | SNPA |
|---------|--------|------------|-----|------|
| skag.00 | 20 | gr-0/2/0.0 | h | |
| skag.02 | 20 | gr-0/2/0.0 | h | |
| h.00 | 10 | gr-0/2/0.0 | h | |
| fix.00 | 0 | | | |
| 4 nodes | | | | |
| ... | | | | |

show isis spf results (CLNS)

```
user@host> show isis spf results
```

IS-IS level 1 SPF results:

| Node | Metric | Interface | Via | SNPA |
|------------|--------|------------------|-----------|----------------|
| skag.00 10 | | fe-0/0/1.0 | toothache | 0:12:0:34:0:56 |
| | | fe-0/0/1.0 | toothache | 0:12:0:34:0:56 |
| | 20 | 192.168.37.64/29 | | |
| | 10 | 1921.6800.4001 | | |
| | 20 | 1921.6800.4002 | | |
| pro1-a.02 | 10 | | | |
| pro1-a.00 | 0 | | | |
| | 0 | 10.255.245.1/32 | | |
| | 10 | 192.168.37.64/29 | | |
| | 0 | 1921.6800.4211 | | |
| 3 nodes | | | | |

IS-IS level 2 SPF results:

| Node | Metric | Interface | Via | SNPA |
|------------|--------|--------------------------------------|-----------|----------------|
| skag.00 10 | | fe-0/0/1.0 | toothache | 0:12:0:34:0:56 |
| | | fe-0/0/1.0 | toothache | 0:12:0:34:0:56 |
| | 20 | 10.255.245.1/32 | | |
| | 20 | 192.168.37.64/29 | | |
| | 20 | 47.0005.80ff.f800.0000.0109.0010/104 | | |
| pro1-a.02 | 10 | | | |
| pro1-a.00 | 0 | | | |
| | 0 | 10.255.245.1/32 | | |
| | 10 | 192.168.37.64/29 | | |
| 3 nodes | | | | |

show isis statistics

| | |
|---|--|
| Syntax | show isis statistics <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switches and QFX Series) | show isis statistics <instance <i>instance-name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series. |
| Description | Display statistics about IS-IS traffic. |
| Options | none —Display IS-IS traffic statistics for all routing instances. instance <i>instance-name</i> —(Optional) Display statistics for the specified routing instance. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">• clear isis statistics on page 317 |
| List of Sample Output | show isis statistics on page 362 |
| Output Fields | Table 91 on page 361 describes the output fields for the show isis statistics command. Output fields are listed in the approximate order in which they appear. |

Table 91: show isis statistics Output Fields

| Field Name | Field Description |
|-----------------------------|---|
| PDU type | <p>PDU type:</p> <ul style="list-style-type: none"> • CSNP—Complete sequence number PDUs contain a complete list of all link-state PDUs in the IS-IS database. CSNPs are sent periodically on all links, and the receiving systems use the information in the CSNP to update and synchronize their link-state PDU databases. The designated router multicasts CSNPs on broadcast links in place of sending explicit acknowledgments for each link-state PDU. • IIH—IS-IS hello packets are broadcast to discover the identity of neighboring IS-IS systems and to determine whether the neighbors are Level 1 or Level 2 intermediate systems. • LSP—Link-state PDUs contain information about the state of adjacencies to neighboring IS-IS systems. Link-state PDUs are flooded periodically throughout an area. • PSNP—Partial sequence number PDUs are sent multicast by a receiver when it detects that it is missing a link-state PDU (when its link-state PDU database is out of date). The receiver sends a PSNP to the system that transmitted the CSNP, effectively requesting that the missing link-state PDU be transmitted. That routing device, in turn, forwards the missing link-state PDU to the requesting routing device. • Unknown—The PDU type is unknown. |
| Received | Number of PDUs received since IS-IS started or since the statistics were set to zero. |
| Processed | Number of PDUs received less the number dropped. |
| Drops | Number of PDUs dropped. |
| Sent | Number of PDUs transmitted since IS-IS started or since the statistics were set to zero. |
| Rexmit | Number of PDUs retransmitted since IS-IS started or since the statistics were set to zero. |
| Total packets received/sent | Total number of PDUs received and transmitted since IS-IS started or since the statistics were set to zero. |
| SNP queue length | Number of CSPN and PSNP packets currently waiting in the queue for processing. This value is almost always 0. |
| LSP queue length | Number of link-state PDUs waiting in the queue for processing. This value is almost always 0. |
| SPF runs | Number of shortest-path-first (SPF) calculations that have been performed. If this number is incrementing rapidly, it indicates that the network is unstable. |
| Fragments rebuilt | Number of link-state PDU fragments that the local system has computed. |
| LSP regenerations | Number of link-state PDUs that have been regenerated. A link-state PDU is regenerated when it is nearing the end of its lifetime and it has not changed. |
| Purges initiated | Number of purges that the system initiated. A purge is initiated if the software decides that a link-state PDU must be removed from the network. |

Sample Output

show isis statistics

user@host> **show isis statistics**

IS-IS statistics for merino:

| PDU type | Received | Processed | Drops | Sent | Rexmit |
|----------|----------|-----------|-------|--------|--------|
| LSP | 12227 | 12227 | 0 | 8184 | 683 |
| IIH | 113808 | 113808 | 0 | 115817 | 0 |
| CSNP | 198868 | 198868 | 0 | 198934 | 0 |
| PSNP | 6985 | 6979 | 6 | 8274 | 0 |
| Unknown | 0 | 0 | 0 | 0 | 0 |
| Totals | 331888 | 331882 | 6 | 331209 | 683 |

Total packets received: 331888 Sent: 331892

| | | | |
|-------------------|---|--------|---|
| SNP queue length: | 0 | Drops: | 0 |
| LSP queue length: | 0 | Drops: | 0 |

| | |
|--------------------|------|
| SPF runs: | 1014 |
| Fragments rebuilt: | 1038 |
| LSP regenerations: | 425 |
| Purges initiated: | 0 |

CHAPTER 8

LLDP Operational Mode Commands

[Table 92 on page 363](#) summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot the Link Layer Discovery Protocol (LLDP) protocol. Commands are listed in alphabetical order.

Table 92: LLDP Operational Mode Commands

| Task | Command |
|--|---|
| Clear LLDP neighbor information. | <code>clear lldp neighbor</code> |
| Clear LLDP statistics. | <code>clear lldp statistics</code> |
| Display basic LLDP information. | <code>show lldp</code> |
| Display LLDP local information. | <code>show lldp local-information</code> |
| Display LLDP neighbor information. | <code>show lldp neighbors</code> |
| Display LLDP remote global statistics. | <code>show lldp remote-global-statistics</code> |
| Display LLDP statistics. | <code>show lldp statistics</code> |

clear lldp neighbor

| | |
|---------------------------------|--|
| Syntax | clear lldp neighbor <interface <i>interface-name</i>> |
| Release Information | Command introduced in Junos OS Release 9.6. |
| Description | <p>On MX Series and T Series routers, clear information regarding all Link Layer Discovery Protocol (LLDP) neighbors or LLDP neighbors of the specified interface.</p> <p>For information about interface names, see Interface Naming Overview. For information about interface names for TX Matrix routers, see TX Matrix Router Chassis and Interface Names. For information about FPC numbering on TX Matrix routers, see Routing Matrix with a TX Matrix Router FPC Numbering.</p> |
| Options | interface <i>interface-name</i> —(Optional) Clear the LLDP neighbors on the specified interface. |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• clear lldp statistics on page 365 |
| List of Sample Output | clear lldp statistics on page 364 |
| Output Fields | When you enter this command, you are provided no feedback on the status of your request. You can enter the show lldp neighbors command before and after clearing the LLDP neighbors to verify the clear operation. |

Sample Output

```
clear lldp statistics      user@host> clear lldp statistics
                           user@host> clear lldp statistics interface ge-0/2/0
```

clear lldp statistics

| | |
|---------------------------------|--|
| Syntax | <code>clear lldpp neighbor</code> <code><interface <i>interface-name</i>></code> |
| Release Information | Command introduced in Junos OS Release 9.6. |
| Description | <p>On MX Series and T Series routers, clear all Link Layer Discovery Protocols (LLDP) statistics or LLDP statistics associated with the specified interface.</p> <p>For information about interface names, see Interface Naming Overview. For information about interface names for TX Matrix routers, see TX Matrix Router Chassis and Interface Names. For information about FPC numbering on TX Matrix routers, see Routing Matrix with a TX Matrix Router FPC Numbering.</p> |
| Options | <code>interface <i>interface-name</i></code> —(Optional) Clear LLDP statistics on the specified interface. |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none"> • clear lldp neighbor on page 364 |
| List of Sample Output | clear lldp neighbor on page 365 |
| Output Fields | When you enter this command, you are provided no feedback on the status of your request. You can enter the show lldp statistics command before and after clearing the LLDP statistics to verify the clear operation. |

Sample Output

```
clear lldp neighbor      user@host> clear lldp neighbors
                        user@host> clear lldp neighbors interface ge-0/2/2
```

show lldp

| | |
|---------------------------------|--|
| Syntax | <code>show lldp</code> <code><detail></code> |
| Release Information | Command introduced in Junos OS Release 9.6. |
| Description | On MX Series and T Series routers, display information about the Link Layer Discovery Protocol (LLDP). |
| Options | detail —(Optional) Display the detailed output level. |
| Required Privilege Level | view |
| List of Sample Output | show lldp on page 368 show lldp detail on page 368 |
| Output Fields | Table 93 on page 366 describes the output fields for the show lldp command. Output fields are listed in the approximate order in which they appear. |

Table 93: show lldp Output Fields

| Field Name | Field Description |
|---------------------------|--|
| LLDP | Status of LLDP: Enabled or Disabled . |
| Advertisement interval | Value of the advertisement interval parameter. |
| Transmit delay | Value of the transmit delay parameter. |
| Hold timer | Value of the hold timer parameter. |
| Notification interval | Value of the notification interval parameter. |
| Config Trap Interval | Value of the configuration trap parameter. |
| Connection Hold timer | Value of the connection hold timer parameter. |
| Interface | <p>List of LLDP interfaces, showing status (Enabled or Disabled) and Neighbor count (detail only).</p> <p>For information about interface names, see Interface Naming Overview. For information about interface names for TX Matrix routers, see TX Matrix Router Chassis and Interface Names. For information about FPC numbering on TX Matrix routers, see Routing Matrix with a TX Matrix Router FPC Numbering.</p> |
| LLDP basic TLVs supported | List of basic LLDP TLVs supported by this device (detail only). |

Table 93: show lldp Output Fields (*continued*)

| Field Name | Field Description |
|--------------------------------|--|
| LLDP 802 TLVs supported | List of IEEE 802.1 LLDP TLVs supported by this device (detail only). |

Sample Output

show lldp

```
user@host> show lldp
LLDP : Enabled
Advertisement interval : 30 Second(s)
Transmit delay : 2 Second(s)
Hold timer : 4 Second(s)
Notification interval : 30 Second(s)
Config Trap Interval : 300 Second(s)
Connection Hold timer : 60 Second(s)
```

| Interface | LLDP |
|-----------|---------|
| ge-0/0/0 | Enabled |
| ge-0/0/1 | Enabled |
| ge-0/0/4 | Enabled |

Sample Output

show lldp detail

```
user@host> show lldp detail
LLDP : Enabled
Advertisement interval : 30 Second(s)
Transmit delay : 2 Second(s)
Hold timer : 4 Second(s)
Notification interval : 30 Second(s)
Config Trap Interval : 300 Second(s)
Connection Hold timer : 60 Second(s)
```

| Interface | LLDP | Neighbor count |
|-----------|---------|----------------|
| ge-0/0/0 | Enabled | 0 |
| ge-0/0/1 | Enabled | 0 |
| ge-0/0/4 | Enabled | 0 |

LLDP basic TLVs supported:

Chassis identifier, Port identifier, Port description, System name, System description, System capabilities, Management address.

LLDP 802 TLVs supported:

Link aggregation, Maximum frame size, MAC/PHY Configuration/Status, Port VLAN ID, Port VLAN name.

show lldp local-information

| | |
|---------------------------------|--|
| Syntax | show lldp local-information |
| Release Information | Command introduced in Junos OS Release 9.6. |
| Description | On MX Series and T Series routers, display local Link Layer Discovery Protocol (LLDP) information. |
| Options | This command has no options. |
| Required Privilege Level | view |
| List of Sample Output | show lldp local-information on page 370 |
| Output Fields | Table 94 on page 369 describes the output fields for the show lldp local-information command. Output fields are listed in the approximate order in which they appear. |

Table 94: show lldp local-information Output Fields

| Field Name | Field Description |
|---------------------------------------|---|
| LLDP Local Information details | Information that follows pertains to the local system. |
| Chassis ID | List of chassis identifiers for local information. |
| System name | Local system name reported by LLDP. |
| System descr | Local system description reported by LLDP. |
| System Capabilities | Capabilities (such as Bridge or Router) that are Supported or Enabled by system on the interface. |
| Management Information | Listed by Interface Name , Address Subtype (such as ipv4), Address (such as 192.168.168.229), Interface Number , and Interface Numbering Subtype . |
| Interface Name | List of local interfaces. For information about interface names, see Interface Naming Overview. For information about interface names for TX Matrix routers, see TX Matrix Router Chassis and Interface Names. For information about FPC numbering on TX Matrix routers, see Routing Matrix with a TX Matrix Router FPC Numbering. |
| Interface ID | List of local interface identifiers. |
| Interface Description | List of local interface descriptions. |
| Status | List of interface conditions: UP or DOWN . |

Sample Output

show lldp
local-information

user@host> **show lldp local-information**

LLDP Local Information details

Chassis ID : 00:90:69:0a:77:c0

System name : sw-mx-u

System descr : Juniper Networks, Inc. MX 960, Version 9.4I0.1, Build date
2008-09-04 14:51:50 UTC

System Capabilities

Supported : Bridge Router

Enabled : Bridge Router

Management Information

Interface Name : fxp0

Address Subtype : IPv4(1)

Address : 192.168.168.229

Interface Number : 1

Interface Numbering Subtype : ifIndex(2)

| Interface Name | Interface ID | Interface Description | Status |
|----------------|--------------|-----------------------|--------|
| ge-0/1/0 | 18 | Avaya Port | UP |
| ge-0/1/1 | 27 | - | DOWN |
| ge-0/1/2 | 13 | Port for Hub | UP |

show lldp neighbors

| | |
|---------------------------------|--|
| Syntax | <code>show lldp neighbors</code> <code><interface <i>interface-name</i>></code> |
| Release Information | Command introduced in Junos OS Release 9.6. |
| Description | On MX Series and T Series routers, display information about LLDP neighbors. For information about interface names, see Interface Naming Overview. For information about interface names for TX Matrix routers, see TX Matrix Router Chassis and Interface Names. For information about FPC numbering on TX Matrix routers, see Routing Matrix with a TX Matrix Router FPC Numbering. |
| Options | <code>interface <i>interface-name</i></code> —(Optional) Display the neighbor information about a particular physical interface. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> clear lldp neighbor on page 364 |
| List of Sample Output | show lldp neighbors on page 373 show lldp neighbors interface ge-0/0/4 on page 373 |
| Output Fields | Table 95 on page 371 describes the output fields for the show lldp neighbors command. Output fields are listed in the approximate order in which they appear. |

Table 95: show lldp neighbors Output Fields

| Field Name | Field Description |
|---------------------------------|--|
| LLDP Remote Devices Information | Information about remote devices. |
| LocalInterface | List of local interfaces for which neighbor information is available. |
| ChassisId | List of chassis identifiers for neighbors. |
| PortInfo | List of port information gathered from neighbors. This could be the port identifier or port description. |
| SysName | List of system names gathered from neighbors. |
| LLDP Neighbor Information | Information about both local and neighbor systems on the interface (appears when the <code>interface</code> option is used). |
| Local Information | Information about local systems on the interface (appears when the <code>interface</code> option is used). |

Table 95: show lldp neighbors Output Fields (*continued*)

| Field Name | Field Description |
|-----------------------------|---|
| Neighbor Information | Information about both local and neighbor system on the interface (appears when the interface option is used). |
| Index | Local interface index (appears when the interface option is used). |
| Time Mark | Date and timestamp of information (appears when the interface option is used). |
| Time To Live | Number of seconds for which this information is valid (appears when the interface option is used). |
| Local Interface | Name of the local physical interface (appears when the interface option is used). |
| Local Port ID | Local port identifier (appears when the interface option is used). |
| Neighbor Information | Information about neighbor systems on the interface (appears when the interface option is used). |
| Chassis type | Type of chassis identifier supplied, such as MAC address (appears when the interface option is used). |
| Chassis ID | Chassis identifier of type listed (appears when the interface option is used). |
| Port type | Type of port identifier supplied, such as local (appears when the interface option is used). |
| Port ID | Port identifier of type listed (appears when the interface option is used). |
| Port description | Port description (appears when the interface option is used). |
| System name | Name supplied by the system on the interface (appears when the interface option is used). |
| System Description | Description supplied by the system on the interface (appears when the interface option is used). |
| System Capabilities | Capabilities (such as bridge or router) that are Supported or Enabled by the system on the interface (appears when the interface option is used). |
| Management address | Details of the management address: Address Type (such as ipv4), Address (such as 10.204.34.35), Interface Number , Interface Subtype , and Organization Identifier (OID) (appears when the interface option is used). |
| Organization Info | One or more entries listing remote information by Organizationally Unique Identifier (OUI), Subtype , Index , and Info (appears when the interface option is used). |

Sample Output

show lldp neighbors

```
user@host> show lldp neighbors
LLDP Remote Devices Information
```

| LocalInterface | ChassisId | PortInfo | SysName |
|----------------|-------------------|-------------------|-----------|
| ge-0/0/0 | 10.209.192.12 | 00 19 bb 20 de 80 | AVA4C357D |
| ge-0/0/1 | 10.209.192.12 | 00 19 bb 20 de 80 | AVA4C357D |
| ge-0/0/1 | 10.209.192.13 | 00 19 bb 20 de 81 | AVA4C357E |
| ge-0/0/3 | 00 19 bb 20 de 79 | 5 | apg-hp1 |
| ge-0/0/3 | 00 19 bb 20 de 80 | 3 | apg-hp1 |
| ge-0/0/4 | 00 19 bb 20 de 79 | 5 | apg-hp1 |
| ge-0/0/4 | 00 19 bb 20 de 80 | 3 | apg-hp1 |
| ge-0/0/5 | 00 19 bb 20 de 81 | ge-0/0/3 | MX480-1 |
| ge-0/0/6 | 00 19 bb 20 de 82 | ge-0/0/4 | MX960-2 |

Sample Output

**show lldp neighbors
interface ge-0/0/4**

```
user@host> show lldp neighbors interface ge-0/0/4
LLDP Neighbor Information:
Local Information:
  Index 6 Time Mark Wed Jun 20 07:34:11 2007 Time To Live 120 seconds
  Local Interface   : ge-0/0/4
  Local Port ID     : 4

Neighbor Information:
  Chassis type      : Mac address
  Chassis ID        : 00 19 bb 20 de 80
  Port type         : local
  Port ID           : 3
  Port description  : 3
  System name       : apg-hp1

System Description : ProCurve J9049A Switch 2900-24G, revision
                    T.11.X1, ROM K....

System Capabilities
  Supported : bridge, router
  Enabled   : bridge

Management address
  Address Type : ipv4
  Address      : 10.204.34.35
  Interface Number : 1
  Interface Subtype : ifIndex(2)
  OID          : 1.3.6.1.2.1.31.1.1.1.1.1

Organization Info
  OUI : 0.18.15
  Subtype : 1
  Index : 1
  Info : 00A0000000

Organization Info
  OUI : 0.18.15
  Subtype : 3
  Index : 2
  Info : 0100000000
```

Organization Info

OUI : 0.18.15
Subtype : 4
Index : 3
Info : 05EA

show lldp remote-global-statistics

| | |
|---------------------------------|---|
| Syntax | show lldp remote-global-statistics |
| Release Information | Command introduced in Junos OS Release 9.6. |
| Description | On MX Series and T Series routers, display remote Link Layer Discovery Protocol (LLDP) global statistics. |
| Options | This command has no options. |
| Required Privilege Level | view |
| List of Sample Output | show lldp remote-global-statistics on page 376 |
| Output Fields | Table 96 on page 375 describes the output fields for the show lldp remote-global-statistics command. Output fields are listed in the approximate order in which they appear. |

Table 96: show lldp remote-global-statistics Output Fields

| Field Name | Field Description |
|-------------------------------------|---|
| LLDP Remote Database Table Counters | Information about remote database table counters. |
| LastchangeTime | Time elapsed between LLDP agent startup and the last change to the remote database table information. |
| Inserts | Number of insertions made in the remote database table. |
| Deletes | Number of deletions made in the remote database table. |
| Drops | Number of LLDP frames dropped from the remote database table because of errors. |
| Ageouts | Number of remote database table entries that have aged out of the table. |

Sample Output

```
show lldp remote-global-statistics
user@host> show lldp remote-global-statistics
LLDP Remote Database Table Counters
LastchangeTime      Inserts    Deletes    Drops    Ageouts
00:00:76 (76 sec)   192        0           0         0
```

show lldp statistics

| | |
|---------------------------------|---|
| Syntax | show lldp statistics <interface <i>interface-name</i> > |
| Release Information | Command introduced in Junos OS Release 9.6. |
| Description | On MX Series and T Series routers, display information about Link Layer Discovery Protocol (LLDP) statistics. |
| Options | interface <i>interface-name</i> —(Optional) Display the statistics about a particular physical interface. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • clear lldp statistics on page 365 |
| List of Sample Output | show lldp statistics on page 378 show lldp statistics interface ge-0/1/1 on page 378 |
| Output Fields | Table 97 on page 377 describes the output fields for the show lldp statistics command. Output fields are listed in the approximate order in which they appear. |

Table 97: show lldp statistics Output Fields

| Field Name | Field Description |
|---------------------|---|
| Interface | Interface name. For information about interface names, see Interface Naming Overview. For information about interface names for TX Matrix routers, see TX Matrix Router Chassis and Interface Names. For information about FPC numbering on TX Matrix routers, see Routing Matrix with a TX Matrix Router FPC Numbering. |
| Received | Number of LLDP frames received on this interface. |
| Transmitted | Number of LLDP frames sent on this interface. |
| Unknown-TLVs | Number of LLDP frames with unsupported content received on this interface. |
| With-Errors | Number of LLDP frames with errors received on this interface. |
| Discarded | Number of LLDP frames received on this interface that were discarded because of problems. |

Sample Output

show lldp statistics

```
user@host> show lldp statistics
Interface Received Transmitted Unknown-TLVs With-Errors Discarded
-----
ge-0/1/1  544      540          0           0           0
ge-0/1/2  540      500          0           0           0
ge-0/1/3  544      540          0           0           0
ge-0/1/4  544      540          0           0           0
ge-0/1/5  544      540          0           0           0
ge-0/1/6  544      540          0           0           0
ge-0/1/7   0         0           0           0           0
```

Sample Output

**show lldp statistics
interface ge-0/1/1**

```
user@host> show lldp statistics interface ge-0/1/1
Interface Received Transmitted Unknown-TLVs With-Errors Discarded
-----
ge-0/1/1  544      540          0           0           0
```

CHAPTER 9

MVRP Operational Mode Commands

Table 98 on page 379 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot the Multiple VLAN Registration Protocol (MVRP). Commands are listed in alphabetical order.

Table 98: MVRP Operational Mode Commands

| Task | Command |
|---|---|
| Display Multiple VLAN Registration Protocol (MVRP) configuration information. | <code>show mvrp</code> |
| Display Multiple VLAN Registration Protocol (MVRP) applicant state information. | <code>show mvrp applicant-state</code> |
| Display all Virtual LANs (VLANs) that have been created dynamically using Multiple VLAN Registration Protocol (MVRP) on the router. | <code>show mvrp dynamic-vlan-memberships</code> |
| Display Multiple VLAN Registration Protocol (MVRP) interface-specific information. | <code>show mvrp interface</code> |
| Display Multiple VLAN Registration Protocol (MVRP) registration state information. | <code>show mvrp registration-state</code> |
| Display Multiple VLAN Registration Protocol (MVRP) statistics in the form of Multiple Registration Protocol data unit (MRPDU) messages. | <code>show mvrp statistics</code> |

show mvrp

| | |
|---------------------------------|--|
| Syntax | show mvrp |
| Release Information | Command introduced in Junos OS Release 10.1. |
| Description | For MX Series routers, display Multiple VLAN Registration Protocol (MVRP) configuration information. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • show mvrp applicant-state on page 382 • show mvrp dynamic-vlan-memberships on page 384 • show mvrp interface on page 385 • show mvrp registration-state on page 386 • show mvrp statistics on page 388 |
| List of Sample Output | show mvrp on page 381 |
| Output Fields | Table 99 on page 380 lists the output fields for the show mvrp command. Output fields are listed in the approximate order in which they appear. |

Table 99: show mvrp Output Fields

| Field Name | Field Description |
|----------------------------|--|
| MVRP dynamic VLAN creation | Displays whether global MVRP dynamic Virtual LAN (VLAN) creation is Enabled or Disabled . |
| MVRP BPDU MAC address | Displays the multicast media access control (MAC) address for MVRP. If configured, the provider MVRP multicast MAC address is used; otherwise, the customer MVRP multicast MAC address is used. |
| MVRP timers (ms) | Displays MVRP timer information: <ul style="list-style-type: none"> • Interface—The interface on which MVRP is configured. • Join—The maximum number of milliseconds the interfaces must wait before sending VLAN advertisements. • Leave—The number of milliseconds an interface must wait after receiving a Leave message to remove the interface from the VLAN specified in the message. • LeaveAll—The interval at which LeaveAll messages are sent on interfaces. LeaveAll messages maintain current MVRP VLAN membership information in the network. |

Sample Output

`show mvrp`

```
user@host> show mvrp
MVRP configuration for routing instance 'default-switch'
MVRP dynamic VLAN creation : Enabled
MVRP BPDU MAC address      : Customer bridge group (01-80-C2-00-00-21)
MVRP timers (ms)
  Interface      Join   Leave  LeaveAll
  ge-11/2/8      200   800    10000
  ge-11/0/9      200   800    10000
  ge-11/3/0      200   800    10000
```

show mvrp applicant-state

| | |
|---------------------------------|--|
| Syntax | show mvrp applicant-state |
| Release Information | Command introduced in Junos OS Release 10.1. |
| Description | For MX Series routers, display Multiple VLAN Registration Protocol (MVRP) applicant state information. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • show mvrp on page 380 • show mvrp interface on page 385 • show mvrp registration-state on page 386 • show mvrp statistics on page 388 |
| List of Sample Output | show mvrp applicant-state on page 383 |
| Output Fields | Table 100 on page 382 lists the output fields for the show mvrp applicant-state command. Output fields are listed in the approximate order in which they appear. |

Table 100: show mvrp applicant-state Output Fields

| Field Name | Field Description |
|------------------|---|
| VLAN Id | Displays the Virtual LAN (VLAN) ID number. |
| Interface | Displays the interface number associated with the VLAN ID. |
| State | Displays one of the following MVRP registrar states: <ul style="list-style-type: none"> • VO—Very anxious observer. • VP—Very anxious passive. • VA—Very anxious new. • AN—Anxious new. • AA—Anxious active. • QA—Quiet active. • LA—Leaving active. • AO—Anxious observer. • QO—Quiet observer. • LO—Leaving observer. • AP—Anxious passive. • QA—Quiet passive. |

Sample Output

`show mvrp
applicant-state`

```
user@host> show mvrp applicant-state
MVRP applicant state for routing instance 'default-switch'
(V0) Very anxious observer, (VP) Very anxious passive, (VA) Very anxious new,
(AN) Anxious new, (AA) Anxious active, (QA) Quiet active, (LA) Leaving active,
(AO) Anxious observer, (QO) Quiet observer, (LO) Leaving observer,
(AP) Anxious passive, (QP) Quiet passive

VLAN Id      Interface      State
-----
100          ge-11/3/0      Declaring (QA)
200          ge-11/3/0      Declaring (QA)
300          ge-11/3/0      Declaring (QA)
```

show mvrp dynamic-vlan-memberships

| | |
|---------------------------------|---|
| Syntax | show mvrp dynamic-vlan-memberships |
| Release Information | Command introduced in Junos OS Release 10.1. |
| Description | For MX Series routers, display all Virtual LANs (VLANs) that have been created dynamically using Multiple VLAN Registration Protocol (MVRP) on the router. |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none"> • show mvrp on page 380 • show mvrp applicant-state on page 382 • show mvrp interface on page 385 • show mvrp registration-state on page 386 • show mvrp statistics on page 388 |
| List of Sample Output | show mvrp dynamic-vlan-memberships on page 384 |
| Output Fields | Table 101 on page 384 lists the output fields for the show mvrp dynamic-vlan-memberships command. Output fields are listed in the approximate order in which they appear. |

Table 101: show mvrp dynamic-vlan-memberships Output Fields

| Field Name | Field Description |
|------------|---|
| VLAN Id | The VLAN ID of the dynamically created VLAN. |
| Interfaces | The interface or interfaces that are bound to the dynamically created VLAN. |

Sample Output

```

show mvrp dynamic-vlan-memberships
user@host> show mvrp dynamic-vlan-memberships
MVRP dynamic vlans for routing instance 'default-switch'
(s) static vlan, (f) fixed registration

VLAN Id      Interfaces
 100 (s)     ge-11/3/0
 200 (s)     ge-11/3/0
 300 (s)
```

show mvrp interface

| | |
|---------------------------------|--|
| Syntax | show mvrp interface |
| Release Information | Command introduced in Junos OS Release 10.1. |
| Description | For MX Series routers, display Multiple VLAN Registration Protocol (MVRP) interface-specific information. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • show mvrp on page 380 • show mvrp applicant-state on page 382 • show mvrp dynamic-vlan-memberships on page 384 • show mvrp registration-state on page 386 • show mvrp statistics on page 388 |
| List of Sample Output | show mvrp interface on page 385 |
| Output Fields | Table 102 on page 385 lists the output fields for the show mvrp interface command. Output fields are listed in the approximate order in which they appear. |

Table 102: show mvrp interface Output Fields

| Field Name | Field Description |
|--------------------------|--|
| Interface | Interface on which MVRP is configured. |
| Status | Status of the MVRP: Enabled or Disabled . |
| Registration Mode | Registration for the interface: Fixed , Forbidden , or Normal . |
| Applicant Mode | Applicant mode. |

Sample Output

```

show mvrp interface      user@host> show mvrp interface
                          MVRP interface information for routing instance 'default-switch'

Interface      Status      Registration      Applicant
                Mode        Mode
ge-11/2/8      Enabled     Normal           Normal
ge-11/0/9      Enabled     Normal           Normal
ge-11/3/0      Enabled     Normal           Normal

```

show mvrp registration-state

| | |
|---------------------------------|--|
| Syntax | show mvrp registration-state |
| Release Information | Command introduced in Junos OS Release 10.1. |
| Description | For MX Series routers, display Multiple VLAN Registration Protocol (MVRP) registration state information. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • show mvrp on page 380 • show mvrp dynamic-vlan-memberships on page 384 • show mvrp interface on page 385 • show mvrp statistics on page 388 |
| List of Sample Output | show mvrp registration-state on page 387 |
| Output Fields | Table 103 on page 386 lists the output fields for the show mvrp registration-state command. Output fields are listed in the approximate order in which they appear. |

Table 103: show mvrp registration-state Output Fields

| Field Name | Field Description |
|------------------------|---|
| VLAN Id | Displays the Virtual LAN (VLAN) ID number. |
| Interface | Displays the interface number associated with the VLAN ID. |
| Registrar State | Displays whether the registrar state is Registered or Empty . |
| Forced State | Displays whether the forced state is Registered or Empty . |
| Managed State | Displays one of the following states: <ul style="list-style-type: none"> • Fixed—VLANs always stay in a registered state and are declared as such on all other forwarding ports. • Normal—VLANs participate in the MVRP protocol and honor incoming join requests normally. • Forbidden—VLANs ignore the incoming join requests and always stay in an unregistered state. |
| STP State | Displays whether the Spanning Tree Protocol (STP) is Blocking or Forwarding . |

Sample Output

`show mvrp
registration-state`

```
user@host> show mvrp registration-state
MVRP registration state for routing instance 'default-switch'
```

| VLAN Id | Interface | Registrar State | Forced State | Managed State | STP State |
|---------|-----------|--------------------|-----------------|------------------|--------------|
| 100 | ge-11/2/8 | Empty | Registered | Fixed | Forwarding |
| | ge-11/0/9 | Empty | Empty | Normal | Forwarding |
| | ge-11/3/0 | Registered | Registered | Normal | Forwarding |
| 101 | ge-11/2/8 | Empty | Registered | Fixed | Forwarding |
| | ge-11/0/9 | Empty | Empty | Normal | Forwarding |
| | ge-11/3/0 | Registered | Registered | Normal | Forwarding |

show mvrp statistics

| | |
|---------------------------------|---|
| Syntax | show mvrp statistics |
| Release Information | Command introduced in Junos OS Release 10.1. |
| Description | For MX Series routers, display Multiple VLAN Registration Protocol (MVRP) statistics in the form of Multiple Registration Protocol data unit (MRPDU) messages. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">• show mvrp on page 380• show mvrp applicant-state on page 382• show mvrp dynamic-vlan-memberships on page 384• show mvrp interface on page 385• show mvrp registration-state on page 386 |
| List of Sample Output | show mvrp statistics on page 389 |
| Output Fields | Table 104 on page 388 lists the output fields for the show mvrp statistics command. Output fields are listed in the approximate order in which they appear. |

Table 104: show mvrp statistics Output Fields

| Field Name | Field Description |
|----------------------------------|--|
| interface name | Interface for which MVRP statistics are displayed. |
| VLAN IDs registered | Number of Virtual LAN (VLAN) IDs registered. |
| Sent MVRP PDUs | Number of MRPDU messages transmitted from the router. |
| Received MVRP PDUs without error | Number of MRPDU messages received on the router. |
| Received MVRP PDUs with error | Number of invalid MRPDU messages received on the router. |

Sample Output

```
show mvrp statistics      user@host> show mvrp statistics
                          MVRP statistics for routing instance 'default-switch'

Interface name           : ge-11/2/8
VLAN IDs registered      : 0
Sent MVRP PDUs           : 1467
Received MVRP PDUs without error: 0
Received MVRP PDUs with error  : 0

Interface name           : ge-11/0/9
VLAN IDs registered      : 0
Sent MVRP PDUs           : 1418
Received MVRP PDUs without error: 702
Received MVRP PDUs with error  : 0

Interface name           : ge-11/3/0
VLAN IDs registered      : 2
Sent MVRP PDUs           : 1524
Received MVRP PDUs without error: 1366
Received MVRP PDUs with error  : 0
```


CHAPTER 10

OSPF Operational Mode Commands

Table 105 on page 391 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot the Open Shortest Path First (OSPF) protocol. Commands are listed in alphabetical order.

Table 105: OSPF Operational Mode Commands

| Task | Command |
|---|--|
| Clear link-state database entries. | <code>clear (ospf ospf3) database</code> |
| Clear OSPF input and output statistics. | <code>clear (ospf ospf3) io-statistics</code> |
| Tear down neighbor connections. | <code>clear (ospf ospf3) neighbor</code> |
| Clear the OSPF overload bit. | <code>clear (ospf ospf3) overload</code> |
| Clear OSPF statistics. | <code>clear (ospf ospf3) statistics</code> |
| Display information about the level of backup coverage available for OSPF nodes and prefixes. | <code>show (ospf ospf3) backup coverage</code> |
| Display information about MPLS label-switched-paths (LSPs) designated as backup routes for OSPF routes. | <code>show (ospf ospf3) backup lsp</code> |
| Display the neighbors through which direct next hops for the backup paths are available. | <code>show (ospf ospf3) backup neighbor</code> |
| Display information about OSPF shortest-path-first calculations for backup paths. | <code>show (ospf ospf3) backup spf</code> |
| Display context identifier information processed and advertised by OSPF for egress protection. | <code>show ospf context-identifier</code> |
| Display link-state database entries for OSPFv2. | <code>show ospf database</code> |
| Display link-state database entries for OSPFv3. | <code>show ospf3 database</code> |
| Display OSPF interface status. | <code>show (ospf ospf3) interface</code> |

Table 105: OSPF Operational Mode Commands (*continued*)

| Task | Command |
|---|--|
| Display OSPF input and output statistics. | <code>show (ospf ospf3) io-statistics</code> |
| Display the SPF log. | <code>show (ospf ospf3) log</code> |
| Display adjacent routers. | <code>show (ospf ospf3) neighbor</code> |
| Display overview statistics. | <code>show (ospf ospf3) overview</code> |
| Display OSPF routing table entries. | <code>show (ospf ospf3) route</code> |
| Display OSPF statistics. | <code>show (ospf ospf3) statistics</code> |



NOTE: For more OSPF-related commands, such as `show route protocol`, `show route instance`, and `show route table`, see Protocol-Independent Routing Operational Mode Commands. For information about monitoring Bidirectional Forwarding Detection (BFD) sessions for OSPF clients, see BFD Operational Mode Commands. For information about how to configure OSPF, see the *Junos Routing Protocols Configuration Guide*.

clear (ospf | ospf3) database

| | |
|---|--|
| Syntax | <pre>clear (ospf ospf3) database <advertising-router (<i>router-id</i> self)> <area <i>area-id</i>> <asbrsummary> <external> <instance <i>instance-name</i>> <inter-area-prefix> <inter-area-router> <intra-area-prefix> <link-local> <logical-system (all <i>logical-system-name</i>)> <lsa-id <i>lsa-id</i>> <netsummary> <network> <nssa> <opaque-area> <purge> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)> <router></pre> |
| Syntax (EX Series Switch and QFX Series) | <pre>clear (ospf ospf3) database <advertising-router (<i>router-id</i> self)> <area <i>area-id</i>> <asbrsummary> <external> <instance <i>instance-name</i>> <inter-area-prefix> <inter-area-router> <intra-area-prefix> <link-local> <lsa-id <i>lsa-id</i>> <netsummary> <network> <nssa> <opaque-area> <purge> <router></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>advertising-router <i>router-id</i>, area <i>area-id</i>, asbrsummary, external, inter-area-prefix, inter-area-router, intra-area-prefix, link-local, lsa-id <i>lsa-id</i>, netsummary, network, nssa, opaque-area, and router options added in Junos OS Release 8.3. You must use the purge command with these options.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>realm option added in Junos OS Release 9.2.</p> <p>advertising-router (<i>router-id</i> self) option added in Junos OS Release 9.5.</p> <p>advertising-router (<i>router-id</i> self) option introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> |

Description With the master Routing Engine, delete entries in the Open Shortest Path First (OSPF) link-state advertisement (LSA) database. With the backup Routing Engine, delete the OSPF LSA database and sync the new database with the master Routing Engine. You can also use the **purge** command with any of the options to discard rather than delete the specified LSA entries.



CAUTION: This command is useful only for testing. Use it with care, because it causes significant network disruption.

Options **none**—Delete all LSAs other than the system's own LSAs, which are regenerated. To resynchronize the database, the system destroys all adjacent neighbors that are in the state **EXSTART** or higher. The neighbors are then reacquired and the databases are synchronized.

advertising-router (*router-id* | **self**)—(Optional) Discard entries for the LSA entries advertised by the specified routing device or by this routing device.

area *area-id*—(Optional) Discard entries for the LSAs in the specified area.

asbrsummary—(Optional) Discard summary AS boundary router LSA entries.

external—(Optional) Discard external LSAs.

instance *instance-name*—(Optional) Delete or discard entries for the specified routing instance only.

inter-area-prefix—(OSPFv3 only) (Optional) Discard interarea prefix LSAs.

inter-area-router—(OSPFv3 only) (Optional) Discard interarea router LSAs.

intra-area-prefix—(OSPFv3 only) (Optional) Discard intra-area prefix LSAs.

logical-system (**all** | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

link-local—(Optional) Delete link-local LSAs.

lsa-id *lsa-id*—(Optional) Discard the LSA entries with the specified LSA identifier.

netsummary—(Optional) Discard summary network LSAs.

network—(Optional) Discard network LSAs.

nssa—(Optional) Discard not-so-stubby area (NSSA) LSAs.

opaque-area—(Optional) Discard opaque area-scope LSAs.

realm (**ipv4-multicast** | **ipv4-unicast** | **ipv6-multicast**)—(OSPFv3 only) (Optional) Delete the entries for the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

router—(Optional) Discard router LSAs.

purge—(Optional) Discard all entries in the link-state advertisement database. All link-state advertisements are set to **MAXAGE** and are flooded. The database is repopulated when the originators of the link-state advertisements receive the **MAXAGE** link-state advertisements and reissue them.

Required Privilege Level

clear

Related Documentation

- [show ospf database on page 420](#)
- [show ospf3 database on page 429](#)

List of Sample Output [clear ospf database on page 395](#)

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

Sample Output

`clear ospf database` user@host> **clear ospf database**

clear (ospf | ospf3) database-protection

| | |
|---------------------------------|---|
| Syntax | clear (ospf ospf3) database-protection <instance <i>instance-name</i> > |
| Release Information | Command introduced in Junos OS Release 10.2. Command introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | Clear the Open Shortest Path First (OSPF) link-state database from its isolated state. Reset the ignore count, ignore timer, and reset timer, and resume normal operations. |
| Options | instance <i>instance-name</i> —(Optional) Clear the OSPF link-state database for the specified routing instance only. |
| Required Privilege Level | clear |
| Output Fields | This command produces no output. |

Sample Output

| | |
|-----------------------------------|---|
| clear ospf database-protection | user@host> clear ospf database-protection |
|-----------------------------------|---|

clear (ospf | ospf3) io-statistics

| | |
|---|---|
| Syntax | clear (ospf ospf3) io-statistics <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switch and QFX Series) | clear (ospf ospf3) io-statistics |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | Clear Open Shortest Path First (OSPF) input and output statistics. |
| Options | none —Clear OSPF input and output statistics. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | clear |
| List of Sample Output | clear ospf io-statistics on page 397 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear ospf io-statistics user@host> clear ospf io-statistics

clear (ospf | ospf3) neighbor

| | |
|---|--|
| Syntax | clear (ospf ospf3) neighbor <area <i>area-id</i> > <instance <i>instance-name</i> > <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)> <neighbor> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)> |
| Syntax (EX Series Switch and QFX Series) | clear (ospf ospf3) neighbor <area <i>area-id</i> > <instance <i>instance-name</i> > <interface <i>interface-name</i> > <neighbor> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. realm option introduced in Junos OS Release 9.2. Command introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | Tear down Open Shortest Path First (OSPF) neighbor connections. |
| Options | none —Tear down OSPF connections with all neighbors for all routing instances. area <i>area-id</i> —(Optional) Tear down neighbor connections for the specified area only. instance <i>instance-name</i> —(Optional) Tear down neighbor connections for the specified routing instance only. interface <i>interface-name</i> —(Optional) Tear down neighbor connections for the specified interface only. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. neighbor —(Optional) Clear the state of the specified neighbor only. realm (ipv4-multicast ipv4-unicast ipv6-multicast) —(Optional) (OSPFv3 only) Clear the state of the specified OSPFv3 realm, or address family. Use the realm option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default. |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• show (ospf ospf3) neighbor on page 452 |
| List of Sample Output | clear ospf neighbor on page 399 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

```
clear ospf neighbor      user@host> clear ospf neighbor
```

clear (ospf | ospf3) overload

| | |
|------------------------------------|---|
| Syntax | clear (ospf ospf3) overload <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switches) | clear (ospf ospf3) overload <instance <i>instance-name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | Clear the Open Shortest Path First (OSPF) overload bit and rebuild link-state advertisements (LSAs). |
| Options | none —Clear the overload bit and rebuild LSAs for all routing instances. instance <i>instance-name</i> —(Optional) Clear the overload bit and rebuild LSAs for the specified routing instance only. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | clear |
| List of Sample Output | clear ospf overload on page 400 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear ospf overload user@host> clear ospf overload

clear (ospf | ospf3) statistics

| | |
|---|---|
| Syntax | clear (ospf ospf3) statistics <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)> |
| Syntax (EX Series Switch and QFX Series) | clear (ospf ospf3) statistics <instance <i>instance-name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. realm option introduced in Junos OS Release 9.2. Command introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | Clear Open Shortest Path First (OSPF) statistics. |
| Options | <p>none—Clear OSPF statistics.</p> <p>instance <i>instance-name</i>—(Optional) Clear statistics for the specified routing instance only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>realm (ipv4-multicast ipv4-unicast ipv6-multicast)—(Optional) (OSPFv3 only) Clear statistics for the specified OSPFv3 realm, or address family. Use the realm option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.</p> |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none"> • show (ospf ospf3) statistics on page 470 |
| List of Sample Output | clear ospf statistics on page 402 |
| Output Fields | See show (ospf ospf3) statistics for an explanation of output fields. |

Sample Output

clear ospf statistics

The following sample output displays OSPF statistics before and after the **clear ospf statistics** command is entered:

user@host> show ospf statistics

| Packet type | Total | | Last 5 seconds | |
|-------------|-------|----------|----------------|----------|
| | Sent | Received | Sent | Received |
| Hello | 3254 | 2268 | 3 | 1 |
| DbD | 41 | 46 | 0 | 0 |
| LSReq | 8 | 7 | 0 | 0 |
| LSUpdate | 212 | 154 | 0 | 0 |
| LSAck | 65 | 98 | 0 | 0 |

| | | | | |
|--------------------------|---|--------------------|---|---|
| DBDs retransmitted | : | 3, last 5 seconds | : | 0 |
| LSAs flooded | : | 12, last 5 seconds | : | 0 |
| LSAs flooded high-prio | : | 0, last 5 seconds | : | 0 |
| LSAs retransmitted | : | 0, last 5 seconds | : | 0 |
| LSAs transmitted to nbr: | : | 3, last 5 seconds | : | 0 |
| LSAs requested | : | 5, last 5 seconds | : | 0 |
| LSAs acknowledged | : | 19, last 5 seconds | : | 0 |

| | | |
|----------------------|---|---|
| Flood queue depth | : | 0 |
| Total rexmit entries | : | 0 |
| db summaries | : | 0 |
| lsreq entries | : | 0 |

Receive errors:

626 subnet mismatches

user@host> clear ospf statistics

user@host> show ospf statistics

| Packet type | Total | | Last 5 seconds | |
|-------------|-------|----------|----------------|----------|
| | Sent | Received | Sent | Received |
| Hello | 3 | 1 | 3 | 1 |
| DbD | 0 | 0 | 0 | 0 |
| LSReq | 0 | 0 | 0 | 0 |
| LSUpdate | 0 | 0 | 0 | 0 |
| LSAck | 0 | 0 | 0 | 0 |

| | | | | |
|--------------------------|---|-------------------|---|---|
| DBDs retransmitted | : | 0, last 5 seconds | : | 0 |
| LSAs flooded | : | 0, last 5 seconds | : | 0 |
| LSAs flooded high-prio | : | 0, last 5 seconds | : | 0 |
| LSAs retransmitted | : | 0, last 5 seconds | : | 0 |
| LSAs transmitted to nbr: | : | 0, last 5 seconds | : | 0 |
| LSAs requested | : | 0, last 5 seconds | : | 0 |
| LSAs acknowledged | : | 0, last 5 seconds | : | 0 |

| | | |
|----------------------|---|---|
| Flood queue depth | : | 0 |
| Total rexmit entries | : | 0 |
| db summaries | : | 0 |
| lsreq entries | : | 0 |

Receive errors:

None

show (ospf | ospf3) backup coverage

| | |
|---------------------------------|---|
| Syntax | <pre>show (ospf ospf3) backup coverage <instance <i>instance-name</i>> < logical-system (all <i>logical-system-name</i>)> <realm (ipv4-unicast ipv6-unicast)> <topology <i>topology-name</i>></pre> |
| Syntax (QFX Series) | <pre>show (ospf ospf3) backup coverage <instance <i>instance-name</i>> <topology <i>topology-name</i>></pre> |
| Release Information | <p>Command introduced in Junos OS Release 10.0.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> |
| Description | Display information about the level of backup coverage available for all the nodes and prefixes in the network. |
| Options | <p>none—Display information about the level backup coverage for all OSPF routing instances in all logical systems.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Display information about the level of backup coverage for all logical systems or for a specific logical system.</p> <p>instance <i>instance-name</i>—(Optional) Display information about the level of backup coverage for a specific OSPF routing instance.</p> <p>realm (ipv4-unicast ipv6-unicast)—(Optional) (OSPFv3 only) Display information about the level of backup coverage for the specific OSPFv3 realm, or address family.</p> <p>topology (default <i>topology-name</i>)—(Optional) (OSPFv2 only) Display information about the level of backup coverage for the specific OSPF topology.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> show (ospf ospf3) backup lsp on page 406 |
| List of Sample Output | <p>show ospf backup coverage on page 405</p> <p>show ospf3 backup coverage on page 405</p> |
| Output Fields | <p>Table 106 on page 403 lists the output fields for the show (ospf ospf3) backup coverage command. Output fields are listed in the approximate order in which they appear.</p> |

Table 106: show (ospf | ospf3) backup coverage Output Fields

| Field Name | Field Description |
|---------------|---|
| Node Coverage | Information about backup coverage for each OSPF node. |
| Area | Area number. Area 0.0.0.0 is the backbone. |

Table 106: show (ospf | ospf3) backup coverage Output Fields (*continued*)

| Field Name | Field Description |
|------------------------|---|
| Covered Nodes | Number of nodes for which backup coverage is available. |
| Total Nodes | Total number of OSPF nodes. |
| Route Coverage | Information about backup coverage for each type of OSPF route. |
| Path Type | Type of OSPF path: Intra , Inter , Ext1 , Ext2 , and All . |
| Covered Routes | For each path type, the number of routes for which backup coverage is available. |
| Total Routes | For each path type, the total number of configured routes. |
| Percent Covered | For all nodes and for each path type, the percentage for which backup coverage is available. |

Sample Output

show ospf backup coverage

user@host> **show ospf backup coverage**

Topology default coverage:

Node Coverage:

| Area | Covered Nodes | Total Nodes | Percent Covered |
|---------|------------------|----------------|--------------------|
| 0.0.0.0 | 4 | 5 | 80.00% |

Route Coverage:

| Path Type | Covered Routes | Total Routes | Percent Covered |
|-----------|-------------------|-----------------|--------------------|
| Intra | 8 | 14 | 57.14% |
| Inter | 0 | 0 | 100.00% |
| Ext1 | 0 | 0 | 100.00% |
| Ext2 | 1 | 1 | 100.00% |
| All | 9 | 15 | 60.00% |

show ospf3 backup coverage

user @host > **show ospf3 backup coverage**

show ospf3 backup coverage

Node Coverage:

| Area | Covered Nodes | Total Nodes | Percent Covered |
|---------|------------------|----------------|--------------------|
| 0.0.0.0 | 4 | 5 | 80.00% |

Route Coverage:

| Path Type | Covered Routes | Total Routes | Percent Covered |
|-----------|-------------------|-----------------|--------------------|
| Intra | 4 | 6 | 66.67% |
| Inter | 0 | 0 | 100.00% |
| Ext1 | 0 | 0 | 100.00% |
| Ext2 | 1 | 1 | 100.00% |
| All | 5 | 7 | 71.43% |

show (ospf | ospf3) backup lsp

Syntax `show (ospf | ospf3) backup lsp`
`<logical-system (all | logical-system-name)>`
`<realm (ipv4-unicast | ipv6-unicast)>`

Release Information Command introduced in Junos OS Release 10.0.

Description Display information about MPLS label-switched-paths (LSPs) designated as backup routes for OSPF routes.



NOTE: MPLS LSPs can be used as backup routes only for routes in the default OSPFv2 topology and not for any configured topology. Additionally, MPLS LSPs cannot be used as backup routes for nondefault instances either for OSPFv2 or OSPFv3.

Options `none`—Display information all MPLS LSPs designated as backup routes.

`logical-system (all | logical-system-name)`—(Optional) Display information about MPLS LSPs designated as backup routes for all logical systems or a specific logical system.

`realm (ipv4-unicast | ipv6-unicast)`—(Optional) (OSPFv3 only) Display information about MPLS LSPs designated as backup routes for a specific realm, or address family.

Required Privilege Level view

Related Documentation

- [show \(ospf | ospf3\) backup coverage on page 403](#)

List of Sample Output [show ospf backup lsp on page 407](#)
[show ospf3 backup lsp on page 407](#)

Output Fields [Table 107 on page 406](#) lists the output fields for the `show (ospf | ospf3) backup lsp` command. Output fields are listed in the approximate order in which they appear.

Table 107: show (ospf | ospf3) backup lsp Output Fields

| Field Name | Field Description |
|----------------------|--|
| <i>MPLS LSP name</i> | Name of each MPLS LSP designated as a backup path. |
| Egress | IP address of the egress router for the LSP. |

Table 107: show (ospf | ospf3) backup lsp Output Fields (*continued*)

| Field Name | Field Description |
|-------------|--|
| Status | <p>State of the LSP:</p> <ul style="list-style-type: none"> • Up—The router can detect RSVP hello messages from the neighbor. • Down—The router has received one of the following indications: <ul style="list-style-type: none"> • Communication failure from the neighbor. • Communication from IGP that the neighbor is unavailable. • Change in the sequence numbers in the RSVP hello messages sent by the neighbor. • Deleted—The LSP is no longer available as a backup path. |
| Last change | Time elapsed since the neighbor state changed either from up or down or from down to up . The format is <i>hh:mm:ss</i> . |
| TE-metric | Configured traffic engineering metric. |
| Metric | Configured metric. |

Sample Output

show ospf backup lsp

```
user@host> show ospf backup lsp
tobanff
  Egress: 10.255.71.239, Status: up, Last change: 00:00:23
  TE-metric: 0, Metric: 0
```

Sample Output

show ospf3 backup lsp

```
user@host> show ospf3 backup lsp
tobanff
  Egress: 10.255.71.239, Status: up, Last change: 00:00:45
  TE-metric: 0, Metric: 0
```

show (ospf | ospf3) backup neighbor

| | |
|---------------------------------|--|
| Syntax | <pre>show (ospf ospf3) backup neighbor <area <i>area-id</i>> <instance (default <i>instance-name</i>)> <logical-system (default ipv4-multicast <i>logical-system-name</i>)> <topology (default ipv4-multicast <i>topology-name</i>)></pre> |
| Syntax (QFX Series) | <pre>show (ospf ospf3) backup neighbor <area <i>area-id</i>> <instance <i>instance-name</i>> <topology (default ipv4-multicast <i>topology-name</i>)></pre> |
| Release Information | <p>Command introduced in Junos OS Release 10.0.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> |
| Description | Display the neighbors through which direct next hops for the backup paths are available. |
| Options | <p>none—Display all neighbors that have direct next hops for backup paths.</p> <p>area <i>area-id</i>—(Optional) Display the area information.</p> <p>instance (default <i>instance-name</i>)—(Optional) Display information about the default routing instance or a particular routing instance.</p> <p>logical-system (default ipv4-multicast <i>logical-system-name</i>)—(Optional) Display information about the default logical system, IPv4 multicast logical system, or a particular logical system.</p> <p>topology (default ipv4-multicast <i>topology-name</i>)—(OSPFv2 only) (Optional) Display information about the default topology, IPv4 multicast topology, or a particular topology.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> show (ospf ospf3) backup spf on page 410 |
| List of Sample Output | show ospf backup neighbor on page 409 |
| Output Fields | Table 108 on page 408 lists the output fields for the show (ospf ospf3) backup neighbor command. Output fields are listed in the approximate order in which they appear. |

Table 108: show (ospf | ospf3) backup neighbor Output Fields

| Field Name | Field Description | Level of Output |
|-------------------------|---|-----------------|
| Neighbor to Self Metric | Metric from the backup neighbor to the OSPF node. | All levels |
| Self to Neighbor Metric | Metric from the OSPF node to the backup neighbor. | All levels |

Table 108: show (ospf |ospf3) backup neighbor Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-----------------|---|-----------------|
| Direct next-hop | Interface and address of the direct next hop. | All levels |

Sample Output

**show ospf backup
neighbor**

```
user@host> show ospf backup neighbor
Topology default backup neighbors:
```

```
Area 0.0.0.5 backup neighbors:
```

```
10.0.0.5
```

```
Neighbor to Self Metric: 5
```

```
Self to Neighbor Metric: 5
```

```
Direct next-hop: ge-4/0/0.111 via 10.0.175.5
```

```
10.0.0.6
```

```
Neighbor to Self Metric: 5
```

```
Self to Neighbor Metric: 5
```

```
Direct next-hop: ge-4/1/0.110 via 10.0.176.6
```

show (ospf | ospf3) backup spf

| | |
|---------------------------------|--|
| Syntax | <pre>show (ospf ospf3) backup spf <brief detail> <area <i>area-id</i>> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <no-coverage> <node-id> <realm (ipv4-unicast ipv6-unicast)> <topology (default ipv4-multicast <i>topology-name</i>)></pre> |
| Syntax (QFX Series) | <pre>show (ospf ospf3) backup spf <brief detail> <area <i>area-id</i>> <instance <i>instance-name</i>> <no-coverage> <node-id> <topology (default ipv4-multicast <i>topology-name</i>)></pre> |
| Release Information | <p>Command introduced in JUNOS Release 10.0.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> |
| Description | Display information about OSPF shortest-path-first calculations for backup paths. |
| Options | <p>none—Display information about OSPF shortest-path-first (SPF) calculations for all backup paths for all destination nodes.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>area <i>area-id</i>—(Optional) Display the area information.</p> <p>instance <i>instance-name</i>—(Optional) Display information about the routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Display information about all logical systems or a specific logical system.</p> <p>no-coverage—(Optional) Display information if there is no backup coverage.</p> <p>node-id—(Optional) Display information about the node specified.</p> <p>realm (ipv4-unicast ipv6-unicast)—(Optional) Display information about the ipv4 or ipv6 realm.</p> <p>topology (default ipv4-multicast <i>topology-name</i>)—(Optional) (OSPFv2 only) Display information about the default topology, IPv4 multicast topology, or a specific topology.</p> |
| Required Privilege Level | view |
| List of Sample Output | <p>show ospf backup spf on page 412</p> <p>show ospf backup spf detail on page 412</p> <p>show ospf3 backup spf on page 414</p> |

Output Fields Table 109 on page 411 lists the output fields for the **show (ospf |ospf3) backup spf** command. Output fields are listed in the approximate order in which they appear.

Table 109: show (ospf |ospf3) backup spf Output Fields

| Field Name | Field Description | Level of Output |
|------------------------------------|--|-----------------|
| Area <i>area-id</i> results | Area for which the results are displayed. Area 0.0.0.0 is the backbone area. | All levels |
| <i>address</i> | Address of the node for which the results are displayed. | All levels |
| Self to Destination Metric | Metric from the node to the destination. | All levels |
| Parent Node | Address of the parent node. | All levels |
| Primary next-hop | Address of the next hop. | All levels |
| Backup Neighbor | Address of the backup neighbor or LSP endpoint and the following information: <ul style="list-style-type: none"> • Neighbor to Destination Metric • Neighbor to Self Metric • Self to Neighbor Metric • Status (Eligible, Not Eligible, Not Evaluated) and the reason for the status. <p>NOTE: If the backup neighbor is an LSP endpoint, it is indicated as such after the neighbor address.</p> | All levels |

Sample Output

```
show ospf backup spf      user@host> show ospf backup spf
                           Topology default results:

                           Area 0.0.0.0 results:

                           pro16-d-lo0.xxx.yyyy.net
                             Self to Destination Metric: 1
                             Parent Node: pro16-b-lo0.xxx.yyyy.net
                             Primary next-hop: at-1/0/1.0
                             Backup Neighbor: pro16-c-lo0.xxx.yyyy.net (LSP endpoint)
                               Neighbor to Destination Metric: 4, Neighbor to Self Metric: 3
                               Self to Neighbor Metric: 3
                               Not eligible, Reason: Path loops
                             Backup Neighbor: pro16-d-lo0.xxx.yyyy.net
                               Neighbor to Destination Metric: 0, Neighbor to Self Metric: 1
                               Self to Neighbor Metric: 1
                               Not eligible, Reason: Primary next-hop link fate sharing
                             ...

show ospf backup spf      user@host> show ospf backup spf detail
detail                    Topology default results:

                           Area 0.0.0.0 results:

                           11.14.10.2
                             Self to Destination Metric: 1
                             Parent Node: 10.255.70.103
                             Primary next-hop: ae0.0
                             Backup Neighbor: 10.255.71.243
                               Neighbor to Destination Metric: 2, Neighbor to Self Metric: 1
                               Self to Neighbor Metric: 1
                               Not eligible, Reason: Path loops
                             Backup Neighbor: 10.255.71.52
                               Neighbor to Destination Metric: 15, Neighbor to Self Metric: 15
                               Self to Neighbor Metric: 1
                               Not eligible, Reason: Primary next-hop link fate sharing
                             Backup Neighbor: 10.255.71.242
                               Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
                               Self to Neighbor Metric: 1
                               Not eligible, Reason: Path loops

                           10.255.71.52
                             Self to Destination Metric: 1
                             Parent Node: 11.14.10.2
                             Primary next-hop: ae0.0 via 11.14.10.2
                             Backup Neighbor: 10.255.71.52
                               Neighbor to Destination Metric: 0, Neighbor to Self Metric: 15
                               Self to Neighbor Metric: 1
                               Not eligible, Reason: Primary next-hop link fate sharing
                             Backup Neighbor: 10.255.71.243
                               Neighbor to Destination Metric: 2, Neighbor to Self Metric: 1
                               Self to Neighbor Metric: 1
                               Not eligible, Reason: Path loops
                             Backup Neighbor: 10.255.71.242
                               Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
                               Self to Neighbor Metric: 1
```

Not eligible, Reason: Path loops

10.255.71.242

Self to Destination Metric: 1

Parent Node: 10.255.70.103

Primary next-hop: as0.0

Backup Neighbor: 10.255.71.242

Neighbor to Destination Metric: 0, Neighbor to Self Metric: 15

Self to Neighbor Metric: 1

Not eligible, Reason: Primary next-hop link fate sharing

Backup Neighbor: 10.255.71.243

Neighbor to Destination Metric: 2, Neighbor to Self Metric: 1

Self to Neighbor Metric: 1

Not eligible, Reason: Path loops

Backup Neighbor: 10.255.71.52

Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15

Self to Neighbor Metric: 1

Not eligible, Reason: Path loops

10.255.71.243

Self to Destination Metric: 1

Parent Node: 10.255.70.103

Primary next-hop: so-6/0/0.0

Backup Neighbor: 10.255.71.243

Neighbor to Destination Metric: 0, Neighbor to Self Metric: 1

Self to Neighbor Metric: 1

Not eligible, Reason: Primary next-hop link fate sharing

Backup Neighbor: 10.255.71.52

Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15

Self to Neighbor Metric: 1

Not eligible, Reason: Path loops

Backup Neighbor: 10.255.71.242

Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15

Self to Neighbor Metric: 1

Not eligible, Reason: Path loops

12.15.0.1

Self to Destination Metric: 2

Parent Node: 10.255.71.243

Primary next-hop: so-6/0/0.0

Backup next-hop: ae0.0 via 11.14.10.2

Backup Neighbor: 10.255.71.243

Neighbor to Destination Metric: 1, Neighbor to Self Metric: 1

Self to Neighbor Metric: 1

Not eligible, Reason: Primary next-hop link fate sharing

Backup Neighbor: 10.255.71.52

Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15

Self to Neighbor Metric: 1

Eligible, Reason: Contributes backup next-hop

Backup Neighbor: 10.255.71.242

Neighbor to Destination Metric: 17, Neighbor to Self Metric: 15

Self to Neighbor Metric: 1

Not evaluated, Reason: Interface is already covered

10.255.71.238

Self to Destination Metric: 2

Parent Node: 10.255.71.243

Primary next-hop: so-6/0/0.0

Backup next-hop: as0.0

Backup Neighbor: 10.255.71.243

Neighbor to Destination Metric: 1, Neighbor to Self Metric: 1

```

    Self to Neighbor Metric: 1
    Not eligible, Reason: Primary next-hop link fate sharing
Backup Neighbor: 10.255.71.242
    Neighbor to Destination Metric: 15, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Eligible, Reason: Contributes backup next-hop
Backup Neighbor: 10.255.71.52
    Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Not evaluated, Reason: Interface is already covered

10.255.71.239
    Self to Destination Metric: 2
    Parent Node: 12.15.0.1
    Primary next-hop: so-6/0/0.0
    Backup next-hop: ae0.0 via 11.14.10.2
Backup Neighbor: 10.255.71.243
    Neighbor to Destination Metric: 1, Neighbor to Self Metric: 1
    Self to Neighbor Metric: 1
    Not eligible, Reason: Primary next-hop link fate sharing
Backup Neighbor: 10.255.71.52
    Neighbor to Destination Metric: 15, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Eligible, Reason: Contributes backup next-hop
Backup Neighbor: 10.255.71.242
    Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Not evaluated, Reason: Interface is already covered

14.15.0.2
    Self to Destination Metric: 3
    Parent Node: 10.255.71.239
    Primary next-hop: so-6/0/0.0
    Backup next-hop: ae0.0 via 11.14.10.2
Backup Neighbor: 10.255.71.243
    Neighbor to Destination Metric: 2, Neighbor to Self Metric: 1
    Self to Neighbor Metric: 1
    Not eligible, Reason: Primary next-hop link fate sharing
Backup Neighbor: 10.255.71.52
    Neighbor to Destination Metric: 15, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Eligible, Reason: Contributes backup next-hop
Backup Neighbor: 10.255.71.242
    Neighbor to Destination Metric: 17, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Not evaluated, Reason: Interface is already covered

```

show ospf3 backup spf user@host> **show ospf3 backup spf**
Area 0.0.0.0 results:

```

10.255.71.52;0.0.0.5
    Self to Destination Metric: 1
    Parent Node: 10.255.70.103
    Primary next-hop: ae0.0
Backup Neighbor: 10.255.71.243
    Neighbor to Destination Metric: 2, Neighbor to Self Metric: 1
    Self to Neighbor Metric: 1
    Not eligible, Reason: Path loops
Backup Neighbor: 10.255.71.52
    Neighbor to Destination Metric: 15, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1

```



```

    Not eligible, Reason: Primary next-hop link fate sharing
Backup Neighbor: 10.255.71.242
    Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Not eligible, Reason: Path loops

10.255.71.52
    Self to Destination Metric: 1
    Parent Node: 10.255.71.52;0.0.0.5
    Primary next-hop: ae0.0 via fe80::290:69ff:fe0f:67f0
Backup Neighbor: 10.255.71.52
    Neighbor to Destination Metric: 0, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Not eligible, Reason: Primary next-hop link fate sharing
Backup Neighbor: 10.255.71.243
    Neighbor to Destination Metric: 2, Neighbor to Self Metric: 1
    Self to Neighbor Metric: 1
    Not eligible, Reason: Path loops
Backup Neighbor: 10.255.71.242
    Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Not eligible, Reason: Path loops

10.255.71.242
    Self to Destination Metric: 1
    Parent Node: 10.255.70.103
    Primary next-hop: as0.0
Backup Neighbor: 10.255.71.242
    Neighbor to Destination Metric: 0, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Not eligible, Reason: Primary next-hop link fate sharing
Backup Neighbor: 10.255.71.243
    Neighbor to Destination Metric: 2, Neighbor to Self Metric: 1
    Self to Neighbor Metric: 1
    Not eligible, Reason: Path loops
Backup Neighbor: 10.255.71.52
    Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Not eligible, Reason: Path loops

10.255.71.243
    Self to Destination Metric: 1
    Parent Node: 10.255.70.103
    Primary next-hop: so-6/0/0.0
Backup Neighbor: 10.255.71.243
    Neighbor to Destination Metric: 0, Neighbor to Self Metric: 1
    Self to Neighbor Metric: 1
    Not eligible, Reason: Primary next-hop link fate sharing
Backup Neighbor: 10.255.71.52
    Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Not eligible, Reason: Path loops
Backup Neighbor: 10.255.71.242
    Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Not eligible, Reason: Path loops

10.255.71.243;0.0.0.2
    Self to Destination Metric: 2
    Parent Node: 10.255.71.243
    Primary next-hop: so-6/0/0.0

```

```
Backup next-hop: ae0.0 via fe80::290:69ff:fe0f:67f0
Backup Neighbor: 10.255.71.243
  Neighbor to Destination Metric: 1, Neighbor to Self Metric: 1
  Self to Neighbor Metric: 1
  Not eligible, Reason: Primary next-hop link fate sharing
Backup Neighbor: 10.255.71.52
  Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
  Self to Neighbor Metric: 1
  Eligible, Reason: Contributes backup next-hop
Backup Neighbor: 10.255.71.242
  Neighbor to Destination Metric: 17, Neighbor to Self Metric: 15
  Self to Neighbor Metric: 1
  Not evaluated, Reason: Interface is already covered

10.255.71.238
  Self to Destination Metric: 2
  Parent Node: 10.255.71.243
  Primary next-hop: so-6/0/0.0
  Backup next-hop: as0.0
Backup Neighbor: 10.255.71.243
  Neighbor to Destination Metric: 1, Neighbor to Self Metric: 1
  Self to Neighbor Metric: 1
  Not eligible, Reason: Primary next-hop link fate sharing
Backup Neighbor: 10.255.71.242
  Neighbor to Destination Metric: 15, Neighbor to Self Metric: 15
  Self to Neighbor Metric: 1
  Eligible, Reason: Contributes backup next-hop
Backup Neighbor: 10.255.71.52
  Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
  Self to Neighbor Metric: 1
  Not evaluated, Reason: Interface is already covered

10.255.71.239
  Self to Destination Metric: 2
  Parent Node: 10.255.71.243;0.0.0.2
  Primary next-hop: so-6/0/0.0
  Backup next-hop: ae0.0 via fe80::290:69ff:fe0f:67f0
Backup Neighbor: 10.255.71.243
  Neighbor to Destination Metric: 1, Neighbor to Self Metric: 1
  Self to Neighbor Metric: 1
  Not eligible, Reason: Primary next-hop link fate sharing
Backup Neighbor: 10.255.71.52
  Neighbor to Destination Metric: 15, Neighbor to Self Metric: 15
  Self to Neighbor Metric: 1
  Eligible, Reason: Contributes backup next-hop
Backup Neighbor: 10.255.71.242
  Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
  Self to Neighbor Metric: 1
  Not evaluated, Reason: Interface is already covered

10.255.71.239;0.0.0.4
  Self to Destination Metric: 3
  Parent Node: 10.255.71.239
  Primary next-hop: so-6/0/0.0
  Backup next-hop: ae0.0 via fe80::290:69ff:fe0f:67f0
Backup Neighbor: 10.255.71.243
  Neighbor to Destination Metric: 2, Neighbor to Self Metric: 1
  Self to Neighbor Metric: 1
  Not eligible, Reason: Primary next-hop link fate sharing
Backup Neighbor: 10.255.71.52
  Neighbor to Destination Metric: 15, Neighbor to Self Metric: 15
```

```
Self to Neighbor Metric: 1
Eligible, Reason: Contributes backup next-hop
Backup Neighbor: 10.255.71.242
Neighbor to Destination Metric: 17, Neighbor to Self Metric: 15
Self to Neighbor Metric: 1
Not evaluated, Reason: Interface is already covered
```

show ospf context-identifier

| | |
|---|---|
| Syntax | <code>show ospf context-identifier</code> <code><brief detail></code> <code><area <i>area-id</i>></code> <code><context-id></code> <code><instance <i>instance-name</i>></code> <code><logical-system (all <i>logical-system-name</i>)></code> |
| Syntax (EX Series Switches and QFX Series) | <code>show ospf context-identifier</code> <code><brief detail></code> <code><area <i>area-id</i>></code> <code><context-id></code> <code><instance <i>instance-name</i>></code> |
| Release Information | Command introduced in Junos OS Release 10.4. Command introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | Display the context identifier information processed and advertised by Open Shortest Path First (OSPF) for egress protection. |
| Options | none —Display information about all context identifiers. brief detail —(Optional) Display the specified level of output. area <i>area-id</i> —(Optional) Display information about the context identifier for the specified area. context-id —(Optional) Display information about the specified context identifier. instance <i>instance-name</i> —(Optional) Display information about the context identifier for the specified routing instance. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">• egress-protection (Layer 2 circuit) in the Junos OS VPNs Configuration Guide• egress-protection (MPLS) in the Junos OS VPNs Configuration Guide |
| List of Sample Output | show ospf context-identifier on page 419 show ospf context-identifier detail on page 419 |
| Output Fields | Table 110 on page 419 lists the output fields for the show ospf context-identifier command. Output fields are listed in the approximate order in which they appear. |

Table 110: show ospf context-identifier Output Fields

| Field Name | Field Description | Level of Output |
|-----------------------------|--|-----------------|
| Context | IPv4 address that defines a protection pair. The context is manually configured on both primary and protector provider edge (PE) devices. | All levels |
| Status | State of the path: active or inactive . | All levels |
| Metric | Advertised OSPF metric. | All levels |
| Area | OSPF area number. | All levels |
| Other Advertisements | Other advertisements received by the OSPF node: <ul style="list-style-type: none"> • Advertising router—Address of the device that sent the advertisement. • Type—Type of OSPF path: inter-area and stub. • Metric—Advertised OSPF metric. • None—No additional advertisements were received by the OSPF node. | detail |

Sample Output

show ospf context-identifier

```
user@host> show ospf context-identifier
Context-id: 2.2.4.3
Status: active, Metric: 65534, PE role: protector, Area: 0.0.0.0
```

show ospf context-identifier detail

```
user@host> show ospf context-identifier detail
Context-id: 88.24.13.1
Status: inactive, Metric: 0, PE role: protector, Area: 0.0.0.13
Other Advertisements:
Advertising router: 8.8.8.103
Type: stub link
Metric: 65534
```

show ospf database

| | |
|---|--|
| Syntax | <pre>show ospf database <brief detail extensive summary> <advertising-router (<i>address</i> self)> <area <i>area-id</i>> <asbrsummary> <external> <instance <i>instance-name</i>> <link-local> <logical-system (all <i>logical-system-name</i>)> <lsa-id <i>lsa-id</i>> <netsummary> <network> <nssa> <opaque-area> <router></pre> |
| Syntax (EX Series Switches and QFX Series) | <pre>show ospf database <brief detail extensive summary> <advertising-router (<i>address</i> self)> <area <i>area-id</i>> <asbrsummary> <external> <instance <i>instance-name</i>> <link-local> <lsa-id <i>lsa-id</i>> <netsummary> <network> <nssa> <opaque-area> <router></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>advertising-router self (<i>address</i> self) option introduced in Junos OS Release 9.5.</p> <p>advertising-router self (<i>address</i> self) option introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> |
| Description | Display the entries in the OSPF version 2 (OSPFv2) link-state database, which contains data about link-state advertisement (LSA) packets. |
| Options | <p>none—Display standard information about entries in the OSPFv2 link-state database for all routing instances.</p> <p>brief detail extensive summary—(Optional) Display the specified level of output.</p> <p>advertising-router (<i>address</i> self)—(Optional) Display the LSAs advertised either by a particular routing device or by this routing device.</p> <p>area <i>area-id</i>—(Optional) Display the LSAs in a particular area.</p> |

asbrsummary—(Optional) Display summary AS boundary router LSA entries.

external—(Optional) Display external LSAs.

instance *instance-name*—(Optional) Display all OSPF database information under the named routing instance.

link-local—(Optional) Display information about link-local LSAs.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

lsa-id *lsa-id*—(Optional) Display the LSA with the specified LSA identifier.

netsummary—(Optional) Display summary network LSAs.

network—(Optional) Display information about network LSAs.

nssa—(Optional) Display information about not-so-stubby area (NSSA) LSAs.

opaque-area—(Optional) Display opaque area-scope LSAs.

router—(Optional) Display information about router LSAs.

Required Privilege Level view

Related Documentation • [clear \(ospf | ospf3\) database on page 393](#)

List of Sample Output [show ospf database on page 424](#)
[show ospf database brief on page 424](#)
[show ospf database detail on page 424](#)
[show ospf database extensive on page 425](#)
[show ospf database summary on page 428](#)

Output Fields [Table 111 on page 421](#) describes the output fields for the **show ospf database** command. Output fields are listed in the approximate order in which they appear.

Table 111: show ospf database Output Fields

| Field Name | Field Description | Level of Output |
|------------|--|-----------------|
| area | Area number. Area 0.0.0.0 is the backbone area. | All levels |
| Type | Type of link advertisement: ASBRSum, Extern, Network, NSSA, OpaqArea, Router, or Summary. | All levels |
| ID | LSA identifier included in the advertisement. An asterisk preceding the identifier marks database entries that originated from the local routing device. | All levels |
| Adv Rtr | Address of the routing device that sent the advertisement. | All levels |
| Seq | Link sequence number of the advertisement. | All levels |

Table 111: show ospf database Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---|--|-------------------------|
| Age | Time elapsed since the LSA was originated, in seconds. | All levels |
| Opt | Optional OSPF capabilities associated with the LSA. | All levels |
| Cksum | Checksum value of the LSA. | All levels |
| Len | Length of the advertisement, in bytes. | All levels |
| Router | Router link-state advertisement information: <ul style="list-style-type: none"> bits—Flags describing the routing device that generated the LSP. link count—Number of links in the advertisement. id—ID of a routing device or subnet on the link. data—For stub networks, the subnet mask. Otherwise, the IP address of the routing device that generated the LSP. type—Type of link. It can be PointToPoint, Transit, Stub, or Virtual. TOS count—Number of type-of-service (ToS) entries in the advertisement. TOS 0 metric—Metric for ToS 0. TOS—Type-of-service (ToS) value. metric—Metric for the ToS. | detail extensive |
| Network | Network link-state advertisement information: <ul style="list-style-type: none"> mask—Network mask. attached router—ID of the attached neighbor. | detail extensive |
| Summary | Summary link-state advertisement information: <ul style="list-style-type: none"> mask—Network mask. TOS—Type-of-service (ToS) value. metric—Metric for the ToS. | detail extensive |
| Gen timer | How long until the LSA is regenerated. | extensive |
| Aging timer | How long until the LSA expires. | extensive |
| Installed <i>hh:mm:ss</i> ago | How long ago the route was installed. | extensive |
| expires in <i>hh:mm:ss</i> | How long until the route expires. | extensive |
| sent <i>hh:mm:ss</i> ago | How long ago the LSA was sent. | extensive |
| Last changed <i>hh:mm:ss</i> ago | How long ago the route was changed. | extensive |
| Change count | Number of times the route has changed. | extensive |

Table 111: show ospf database Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------|--|------------------|
| Ours | Indicates that this is a local advertisement. | extensive |
| Router LSAs | Number of router link-state advertisements in the link-state database. | summary |
| Network LSAs | Number of network link-state advertisements in the link-state database. | summary |
| Summary LSAs | Number of summary link-state advertisements in the link-state database. | summary |
| NSSA LSAs | Number of not-so-stubby area link-state advertisements in the link-state database. | summary |

Sample Output

show ospf database

```

user@host> show ospf database
OSPF link state database, Area 0.0.0.1
  Type      ID                Adv Rtr          Seq      Age  Opt  Cksum  Len
Router     10.255.70.103        10.255.70.103    0x80000002 215  0x20 0x4112 48
Router     *10.255.71.242        10.255.71.242    0x80000002 214  0x20 0x11b1 48
Summary    *23.1.1.0             10.255.71.242    0x80000002 172  0x20 0x6d72 28
Summary    *24.1.1.0             10.255.71.242    0x80000002 177  0x20 0x607e 28
NSSA       *33.1.1.1             10.255.71.242    0x80000002 217  0x28 0x73bd 36

      OSPF link state database, Area 0.0.0.2
  Type      ID                Adv Rtr          Seq      Age  Opt  Cksum  Len
Router     10.255.71.52         10.255.71.52     0x80000004 174  0x20 0xd021 36
Router     *10.255.71.242        10.255.71.242    0x80000003 173  0x20 0xe191 36
Network    *23.1.1.1             10.255.71.242    0x80000002 173  0x20 0x9c76 32
Summary    *12.1.1.0             10.255.71.242    0x80000001 217  0x20 0xfeec 28
Summary    *24.1.1.0             10.255.71.242    0x80000002 177  0x20 0x607e 28
NSSA       *33.1.1.1             10.255.71.242    0x80000001 222  0x28 0xe047 36

      OSPF link state database, Area 0.0.0.3
  Type      ID                Adv Rtr          Seq      Age  Opt  Cksum  Len
Router     10.255.71.238        10.255.71.238    0x80000003 179  0x20 0x3942 36
Router     *10.255.71.242        10.255.71.242    0x80000003 177  0x20 0xf37d 36
Network    *24.1.1.1             10.255.71.242    0x80000002 177  0x20 0xc591 32
Summary    *12.1.1.0             10.255.71.242    0x80000001 217  0x20 0xfeec 28
Summary    *23.1.1.0             10.255.71.242    0x80000002 172  0x20 0x6d72 28
NSSA       *33.1.1.1             10.255.71.242    0x80000001 222  0x28 0xeb3b 36

```

show ospf database brief

The output for the **show ospf database brief** command is identical to that for the **show ospf database** command. For sample output, see [show ospf database on page 424](#).

show ospf database detail

```

user@host> show ospf database detail
      OSPF link state database, Area 0.0.0.1
  Type      ID                Adv Rtr          Seq      Age  Opt  Cksum  Len
Router     10.255.70.103        10.255.70.103    0x80000002 261  0x20 0x4112 48
  bits 0x0, link count 2
  id 10.255.71.242, data 12.1.1.1, Type PointToPoint (1)
  TOS count 0, TOS 0 metric 1
  id 12.1.1.0, data 255.255.255.0, Type Stub (3)
  TOS count 0, TOS 0 metric 1
Router     *10.255.71.242        10.255.71.242    0x80000002 260  0x20 0x11b1 48
  bits 0x3, link count 2
  id 10.255.70.103, data 12.1.1.2, Type PointToPoint (1)
  TOS count 0, TOS 0 metric 1
  id 12.1.1.0, data 255.255.255.0, Type Stub (3)
  TOS count 0, TOS 0 metric 1
Summary    *23.1.1.0             10.255.71.242    0x80000002 218  0x20 0x6d72 28
  mask 255.255.255.0
  TOS 0x0, metric 1
Summary    *24.1.1.0             10.255.71.242    0x80000002 223  0x20 0x607e 28
  mask 255.255.255.0
  TOS 0x0, metric 1
NSSA       *33.1.1.1             10.255.71.242    0x80000002 263  0x28 0x73bd 36
  mask 255.255.255.255
  Type 2, TOS 0x0, metric 0, fwd addr 12.1.1.2, tag 0.0.0.0

      OSPF link state database, Area 0.0.0.2

```

```

Type      ID          Adv Rtr      Seq      Age  Opt  Cksum  Len
Router  10.255.71.52    10.255.71.52  0x80000004  220  0x20 0xd021  36
  bits 0x0, link count 1
  id 23.1.1.1, data 23.1.1.2, Type Transit (2)
  TOS count 0, TOS 0 metric 1
Router *10.255.71.242 10.255.71.242  0x80000003  219  0x20 0xe191  36
  bits 0x3, link count 1
  id 23.1.1.1, data 23.1.1.1, Type Transit (2)
  TOS count 0, TOS 0 metric 1
Network *23.1.1.1    10.255.71.242  0x80000002  219  0x20 0x9c76  32
  mask 255.255.255.0
  attached router 10.255.71.242
  attached router 10.255.71.52
Summary *12.1.1.0    10.255.71.242  0x80000001  263  0x20 0xfeec  28
  mask 255.255.255.0
  TOS 0x0, metric 1
Summary *24.1.1.0    10.255.71.242  0x80000002  223  0x20 0x607e  28
  mask 255.255.255.0
  TOS 0x0, metric 1
NSSA   *33.1.1.1      10.255.71.242  0x80000001  268  0x28 0xe047  36
  mask 255.255.255.255
  Type 2, TOS 0x0, metric 0, fwd addr 23.1.1.1, tag 0.0.0.0

```

OSPF link state database, Area 0.0.0.3

```

Type      ID          Adv Rtr      Seq      Age  Opt  Cksum  Len
Router  10.255.71.238  10.255.71.238  0x80000003  225  0x20 0x3942  36
  bits 0x0, link count 1
  id 24.1.1.1, data 24.1.1.2, Type Transit (2)
  TOS count 0, TOS 0 metric 1
Router *10.255.71.242 10.255.71.242  0x80000003  223  0x20 0xf37d  36
  bits 0x3, link count 1
  id 24.1.1.1, data 24.1.1.1, Type Transit (2)
  TOS count 0, TOS 0 metric 1
Network *24.1.1.1    10.255.71.242  0x80000002  223  0x20 0xc591  32
  mask 255.255.255.0
  attached router 10.255.71.242
  attached router 10.255.71.238
Summary *12.1.1.0    10.255.71.242  0x80000001  263  0x20 0xfeec  28
  mask 255.255.255.0
  TOS 0x0, metric 1
Summary *23.1.1.0    10.255.71.242  0x80000002  218  0x20 0x6d72  28
  mask 255.255.255.0
  TOS 0x0, metric 1
NSSA   *33.1.1.1      10.255.71.242  0x80000001  268  0x28 0xeb3b  36
  mask 255.255.255.255
  Type 2, TOS 0x0, metric 0, fwd addr 24.1.1.1, tag 0.0.0.0

```

show ospf database extensive

```

user@host> show ospf database extensive
  OSPF link state database, Area 0.0.0.1
Type      ID          Adv Rtr      Seq      Age  Opt  Cksum  Len
Router  10.255.70.103  10.255.70.103  0x80000002  286  0x20 0x4112  48
  bits 0x0, link count 2
  id 10.255.71.242, data 12.1.1.1, Type PointToPoint (1)
  TOS count 0, TOS 0 metric 1
  id 12.1.1.0, data 255.255.255.0, Type Stub (3)
  TOS count 0, TOS 0 metric 1
  Aging timer 00:55:14
  Installed 00:04:43 ago, expires in 00:55:14
  Last changed 00:04:43 ago, Change count: 2
Router *10.255.71.242 10.255.71.242  0x80000002  285  0x20 0x11b1  48
  bits 0x3, link count 2

```

```

id 10.255.70.103, data 12.1.1.2, Type PointToPoint (1)
TOS count 0, TOS 0 metric 1
id 12.1.1.0, data 255.255.255.0, Type Stub (3)
TOS count 0, TOS 0 metric 1
Gen timer 00:45:15
Aging timer 00:55:15
Installed 00:04:45 ago, expires in 00:55:15, sent 00:04:43 ago
Last changed 00:04:45 ago, Change count: 2, Ours
Summary *23.1.1.0      10.255.71.242    0x80000002    243    0x20 0x6d72    28
mask 255.255.255.0
TOS 0x0, metric 1
Gen timer 00:45:57
Aging timer 00:55:57
Installed 00:04:03 ago, expires in 00:55:57, sent 00:04:01 ago
Last changed 00:04:48 ago, Change count: 1, Ours
Summary *24.1.1.0      10.255.71.242    0x80000002    248    0x20 0x607e    28
mask 255.255.255.0
TOS 0x0, metric 1
Gen timer 00:45:52
Aging timer 00:55:52
Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:06 ago
Last changed 00:04:48 ago, Change count: 1, Ours
NSSA *33.1.1.1        10.255.71.242    0x80000002    288    0x28 0x73bd    36
mask 255.255.255.255
Type 2, TOS 0x0, metric 0, fwd addr 12.1.1.2, tag 0.0.0.0
Gen timer 00:45:12
Aging timer 00:55:12
Installed 00:04:48 ago, expires in 00:55:12, sent 00:04:48 ago
Last changed 00:04:48 ago, Change count: 2, Ours

    OSPF link state database, Area 0.0.0.2
Type      ID          Adv Rtr      Seq      Age  Opt  Cksum  Len
Router 10.255.71.52    10.255.71.52 0x80000004  245  0x20 0xd021  36
bits 0x0, link count 1
id 23.1.1.1, data 23.1.1.2, Type Transit (2)
TOS count 0, TOS 0 metric 1
Aging timer 00:55:55
Installed 00:04:02 ago, expires in 00:55:55
Last changed 00:04:02 ago, Change count: 2
Router *10.255.71.242  10.255.71.242 0x80000003  244  0x20 0xe191  36
bits 0x3, link count 1
id 23.1.1.1, data 23.1.1.1, Type Transit (2)
TOS count 0, TOS 0 metric 1
Gen timer 00:45:56
Aging timer 00:55:56
Installed 00:04:04 ago, expires in 00:55:56, sent 00:04:02 ago
Last changed 00:04:04 ago, Change count: 2, Ours
Network *23.1.1.1      10.255.71.242  0x80000002    244  0x20 0x9c76  32
mask 255.255.255.0
attached router 10.255.71.242
attached router 10.255.71.52
Gen timer 00:45:56
Aging timer 00:55:56
Installed 00:04:04 ago, expires in 00:55:56, sent 00:04:02 ago
Last changed 00:04:04 ago, Change count: 1, Ours
Summary *12.1.1.0      10.255.71.242  0x80000001    288  0x20 0xfeec  28
mask 255.255.255.0
TOS 0x0, metric 1
Gen timer 00:45:12
Aging timer 00:55:12
Installed 00:04:48 ago, expires in 00:55:12, sent 00:04:04 ago

```

```

Last changed 00:04:48 ago, Change count: 1, Ours
Summary *24.1.1.0      10.255.71.242    0x80000002    248    0x20 0x607e    28
mask 255.255.255.0
TOS 0x0, metric 1
Gen timer 00:45:52
Aging timer 00:55:52
Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:04 ago
Last changed 00:04:48 ago, Change count: 1, Ours
NSSA *33.1.1.1      10.255.71.242    0x80000001    293    0x28 0xe047    36
mask 255.255.255.255
Type 2, TOS 0x0, metric 0, fwd addr 23.1.1.1, tag 0.0.0.0
Gen timer 00:45:07
Aging timer 00:55:07
Installed 00:04:53 ago, expires in 00:55:07, sent 00:04:04 ago
Last changed 00:04:53 ago, Change count: 1, Ours

```

```

OSPF link state database, Area 0.0.0.3
Type      ID      Adv Rtr      Seq      Age  Opt  Cksum  Len
Router  10.255.71.238  10.255.71.238  0x80000003  250  0x20 0x3942  36
bits 0x0, link count 1
id 24.1.1.1, data 24.1.1.2, Type Transit (2)
TOS count 0, TOS 0 metric 1
Aging timer 00:55:50
Installed 00:04:07 ago, expires in 00:55:50
Last changed 00:04:07 ago, Change count: 2
Router *10.255.71.242  10.255.71.242  0x80000003  248  0x20 0xf37d  36
bits 0x3, link count 1
id 24.1.1.1, data 24.1.1.1, Type Transit (2)
TOS count 0, TOS 0 metric 1
Gen timer 00:45:52
Aging timer 00:55:52
Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:06 ago
Last changed 00:04:08 ago, Change count: 2, Ours
Network *24.1.1.1      10.255.71.242    0x80000002    248    0x20 0xc591    32
mask 255.255.255.0
attached router 10.255.71.242
attached router 10.255.71.238
Gen timer 00:45:52
Aging timer 00:55:52
Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:06 ago
Last changed 00:04:08 ago, Change count: 1, Ours
Summary *12.1.1.0      10.255.71.242    0x80000001    288    0x20 0xfeec    28
mask 255.255.255.0
TOS 0x0, metric 1
Gen timer 00:45:12
Aging timer 00:55:12
Installed 00:04:48 ago, expires in 00:55:12, sent 00:04:13 ago
Last changed 00:04:48 ago, Change count: 1, Ours
Summary *23.1.1.0      10.255.71.242    0x80000002    243    0x20 0x6d72    28
mask 255.255.255.0
TOS 0x0, metric 1
Gen timer 00:45:57
Aging timer 00:55:57
Installed 00:04:03 ago, expires in 00:55:57, sent 00:04:01 ago
Last changed 00:04:48 ago, Change count: 1, Ours
NSSA *33.1.1.1      10.255.71.242    0x80000001    293    0x28 0xeb3b    36
mask 255.255.255.255
Type 2, TOS 0x0, metric 0, fwd addr 24.1.1.1, tag 0.0.0.0
Gen timer 00:45:07
Aging timer 00:55:07
Installed 00:04:53 ago, expires in 00:55:07, sent 00:04:13 ago

```

Last changed 00:04:53 ago, Change count: 1, Ours

**show ospf database
summary**

user@host> show ospf database summary

```
Area 0.0.0.1:
  2 Router LSAs
  2 Summary LSAs
  1 NSSA LSAs
Area 0.0.0.2:
  2 Router LSAs
  1 Network LSAs
  2 Summary LSAs
  1 NSSA LSAs
Area 0.0.0.3:
  2 Router LSAs
  1 Network LSAs
  2 Summary LSAs
  1 NSSA LSAs
Externals:
Interface fe-2/2/1.0:
Interface ge-0/3/2.0:
Interface so-0/1/2.0:
Interface so-0/1/2.0:
```

show ospf3 database

| | |
|---|--|
| Syntax | <pre>show ospf3 database <brief detail extensive summary> <advertising-router (address self)> <area area-id> <external> <instance instance-name> <inter-area-prefix> <inter-area-router> <intra-area-prefix> <link> <link-local> <logical-system (all logical-system-name)> <lsa-id lsa-id> <network> <nssa> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)> <router></pre> |
| Syntax (EX Series Switches and QFX Series) | <pre>show ospf3 database <brief detail extensive summary> <advertising-router (address self)> <area area-id> <external> <instance instance-name> <inter-area-prefix> <inter-area-router> <intra-area-prefix> <link> <link-local> <lsa-id lsa-id> <network> <nssa> <router></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>realm option introduced in Junos OS Release 9.2.</p> <p>advertising-router (address self) option introduced in Junos Release 9.5.</p> <p>advertising-router (address self) option introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> |
| Description | Display the entries in the OSPF version 3 (OSPFv3) link-state database, which contains data about link-state advertisement (LSA) packets. |
| Options | <p>none—Display standard information about all entries in the OSPFv3 link-state database.</p> <p>brief detail extensive summary—(Optional) Display the specified level of output.</p> <p>advertising-router (address self)—(Optional) Display the LSAs advertised either by a particular routing device or by this routing device.</p> |

area *area-id*—(Optional) Display the LSAs in a particular area.

external—(Optional) Display external LSAs.

instance *instance-name*—(Optional) Display all OSPF database information under the named routing instance.

inter-area-prefix—(Optional) Display information about interarea-prefix LSAs.

inter-area-router—(Optional) Display information about interarea-router LSAs.

intra-area-prefix—(Optional) Display information about intra-area-prefix LSAs.

link—(Optional) Display information about link LSAs.

link-local—(Optional) Display information about link-local LSAs.

logical-system (**all** | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

lsa-id *lsa-id*—(Optional) Display the LSA with the specified LSA identifier.

network—(Optional) Display information about network LSAs.

nssa—(Optional) Display information about not-so-stubby area (NSSA) LSAs.

realm (**ipv4-multicast** | **ipv4-unicast** | **ipv6-multicast**)—(Optional) Display information about the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family other than IPv6 unicast, which is the default.

router—(Optional) Display information about router LSAs.

Required Privilege Level view

Related Documentation • [clear \(ospf | ospf3\) database on page 393](#)

List of Sample Output [show ospf3 database brief on page 436](#)
[show ospf3 database extensive on page 436](#)
[show ospf3 database summary on page 439](#)

Output Fields [Table 112 on page 430](#) lists the output fields for the **show ospf3 database** command. Output fields are listed in the approximate order in which they appear.

Table 112: show ospf3 database Output Fields

| Field Name | Field Description | Level of Output |
|---|---|------------------------|
| OSPF link state database, area <i>area-number</i> | Entries in the link-state database for this area. | brief detail extensive |
| OSPF AS SCOPE link state database | Entries in the AS scope link-state database. | brief detail extensive |

Table 112: show ospf3 database Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--|--|------------------------|
| OSPF Link-Local link state database, interface <i>interface-name</i> | Entries in the link-local link-state database for this interface. | brief detail extensive |
| area | Area number. Area 0.0.0.0 is the backbone area. | All levels |
| Type | Type of link advertisement: Extern , InterArPfx , InterArRtr , IntraArPrx , Link , Network , NSSA , or Router . | brief detail extensive |
| ID | Link identifier included in the advertisement. An asterisk (*) preceding the identifier marks database entries that originated from the local routing device. | brief detail extensive |
| Adv Rtr | Address of the routing device that sent the advertisement. | brief detail extensive |
| Seq | Link sequence number of the advertisement. | brief detail extensive |
| Age | Time elapsed since the LSA was originated, in seconds. | brief detail extensive |
| Cksum | Checksum value of the LSA. | brief detail extensive |
| Len | Length of the advertisement, in bytes. | brief detail extensive |
| Router (Router Link-State Advertisements) | | |
| bits | Flags describing the routing device that generated the LSP. | detail extensive |
| Options | Option bits carried in the router LSA. | detail extensive |
| For Each Router Link | | |
| Type | Type of interface. The value of all other output fields describing a routing device interface depends on the interface's type: <ul style="list-style-type: none"> • PointToPoint (1)—Point-to-point connection to another routing device. • Transit (2)—Connection to a transit network. • Virtual (4)—Virtual link. | detail extensive |
| Loc-if-id | Local interface ID assigned to the interface that uniquely identifies the interface with the routing device. | detail extensive |
| Nbr-if-id | Interface ID of the neighbor's interface for this routing device link. | detail extensive |
| Nbr-rtr-id | Router ID of the neighbor routing device (for type 2 interfaces, the attached link's designated router). | detail extensive |
| Metric | Cost of the router link. | detail extensive |
| Gen timer | How long until the LSA is regenerated, in the format <i>hours:minutes:seconds</i> . | extensive |

Table 112: show ospf3 database Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--|--|------------------|
| Aging timer | How long until the LSA expires, in the format <i>hours:minutes:seconds</i> . | extensive |
| Installed <i>nn:nn:nn</i> ago | How long ago the route was installed, in the format <i>hours:minutes:seconds</i> . | extensive |
| expires in <i>nn:nn:nn</i> | How long until the route expires, in the format <i>hours:minutes:seconds</i> . | extensive |
| sent <i>nn:nn:nn</i> ago | Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format <i>hours:minutes:seconds</i> . | extensive |
| Ours | Indicates that this is a local advertisement. | extensive |
| Network (Network Link-State Advertisements) | | |
| Options | Option bits carried in the network LSA. | detail extensive |
| Attached Router | Router IDs of each of the routing devices attached to the link. Only routing devices that are fully adjacent to the designated router are listed. The designated router includes itself in this list. | detail extensive |
| InterArPfx (Interarea-Prefix Link-State Advertisements) | | |
| Prefix | IPv6 address prefix. | detail extensive |
| Prefix-options | Option bit associated with the prefix. | detail extensive |
| Metric | Cost of this route. Expressed in the same units as the interface costs in the router LSAs. When the interarea-prefix LSA is describing a route to a range of addresses, the cost is set to the maximum cost to any reachable component of the address range. | detail extensive |
| Gen timer | How long until the LSA is regenerated, in the format <i>hours:minutes:seconds</i> . | extensive |
| Aging timer | How long until the LSA expires, in the format <i>hours:minutes:seconds</i> . | extensive |
| Installed <i>nn:nn:nn</i> ago | How long ago the route was installed, in the format <i>hours:minutes:seconds</i> . | extensive |
| expires in <i>nn:nn:nn</i> | How long until the route expires, in the format <i>hours:minutes:seconds</i> . | extensive |
| sent <i>nn:nn:nn</i> ago | Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format <i>hours:minutes:seconds</i> . | extensive |
| Ours | Indicates that this is a local advertisement. | extensive |
| InterArRtr (Interarea-Router Link-State Advertisements) | | |
| Dest-router-id | Router ID of the routing device described by the LSA. | detail extensive |
| options | Optional capabilities supported by the routing device. | detail extensive |

Table 112: show ospf3 database Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--|--|-------------------------|
| Metric | Cost of this route. Expressed in the same units as the interface costs in the router LSAs. When the interarea-prefix LSA is describing a route to a range of addresses, the cost is set to the maximum cost to any reachable component of the address range. | detail extensive |
| Prefix | IPv6 address prefix. | extensive |
| Prefix-options | Option bit associated with the prefix. | extensive |
| Extern (External Link-State Advertisements) | | |
| Prefix | IPv6 address prefix. | detail extensive |
| Prefix-options | Option bit associated with the prefix. | detail extensive |
| Metric | Cost of the route, which depends on the value of Type . | detail extensive |
| Type <i>n</i> | Type of external metric: Type 1 or Type 2 . | detail extensive |
| Aging timer | How long until the LSA expires, in the format <i>hours:minutes:seconds</i> . | extensive |
| Installed <i>nn:nn:nn</i> ago | How long ago the route was installed, in the format <i>hours:minutes:seconds</i> . | extensive |
| expires in <i>nn:nn:nn</i> | How long until the route expires, in the format <i>hours:minutes:seconds</i> . | extensive |
| sent <i>nn:nn:nn</i> ago | Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format <i>hours:minutes:seconds</i> . | extensive |
| Link (Link-State Advertisements) | | |
| IPv6-Address | IPv6 link-local address on the link for which this link LSA originated. | detail extensive |
| Options | Option bits carried in the link LSA. | detail extensive |
| priority | Router priority of the interface attaching the originating routing device to the link. | detail extensive |
| Prefix-count | Number of IPv6 address prefixes contained in the LSA. The rest of the link LSA contains a list of IPv6 prefixes to be associated with the link. | detail extensive |
| Prefix | IPv6 address prefix. | detail extensive |
| Prefix-options | Option bit associated with the prefix. | detail extensive |
| Gen timer | How long until the LSA is regenerated, in the format <i>hours:minutes:seconds</i> . | extensive |
| Aging timer | How long until the LSA expires, in the format <i>hours:minutes:seconds</i> . | extensive |

Table 112: show ospf3 database Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---|---|------------------|
| Installed <i>nn:nn:nn</i> ago | How long ago the route was installed, in the format <i>hours:minutes:seconds</i> . | extensive |
| expires in <i>nn:nn:nn</i> | How long until the route expires, in the format <i>hours:minutes:seconds</i> . | extensive |
| sent <i>nn:nn:nn</i> ago | Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format <i>hours:minutes:seconds</i> . | extensive |
| Ours | Indicates that this is a local advertisement. | extensive |
| IntraArPfx (Intra-Area-Prefix Link-State Advertisements) | | |
| Ref-lsa-type | LSA type of the referenced LSA. <ul style="list-style-type: none"> Router—Address prefixes are associated with a router LSA. Network—Address prefixes are associated with a network LSA. | detail extensive |
| Ref-lsa-id | Link-state ID of the referenced LSA. | detail extensive |
| Ref-router-id | Advertising router ID of the referenced LSA. | detail extensive |
| Prefix-count | Number of IPv6 address prefixes contained in the LSA. The rest of the link LSA contains a list of IPv6 prefixes to be associated with the link. | detail extensive |
| Prefix | IPv6 address prefix. | detail extensive |
| Prefix-options | Option bit associated with the prefix. | detail extensive |
| Metric | Cost of this prefix. Expressed in the same units as the interface costs in the router LSAs. | detail extensive |
| Gen timer | How long until the LSA is regenerated, in the format <i>hours:minutes:seconds</i> . | extensive |
| Aging timer | How long until the LSA expires, in the format <i>hours:minutes:seconds</i> . | extensive |
| Installed <i>hh:mm:ss</i> ago | How long ago the route was installed, in the format <i>hours:minutes:seconds</i> . | extensive |
| expires in <i>hh:mm:ss</i> | How long until the route expires, in the format <i>hours:minutes:seconds</i> . | extensive |
| sent <i>hh:mm:ss</i> ago | Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format <i>hours:minutes:seconds</i> . | extensive |
| <i>n</i> Router LSAs | Number of router LSAs in the link-state database. | summary |
| <i>n</i> Network LSAs | Number of network LSAs in the link-state database. | summary |
| <i>n</i> InterArPfx LSAs | Number of interarea-prefix LSAs in the link-state database. | summary |

Table 112: show ospf3 database Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|------------------------------------|--|-----------------|
| <i>n</i> InterArRtr LSAs | Number of interarea-router LSAs in the link-state database. | summary |
| <i>n</i> IntraArPfx LSAs | Number of intra-area-prefix LSAs in the link-state database. | summary |
| Externals | Display of the external LSA database. | summary |
| <i>n</i> Extern LSAs | Number of external LSAs in the link-state database. | summary |
| Interface <i>interface-name</i> | Name of the interface for which link-local LSA information is displayed. | summary |
| <i>n</i> Link LSAs | Number of link LSAs in the link-state database. | summary |

Sample Output

show ospf3 database brief

```
user@host> show ospf3 database brief
    OSPF3 link state database, area 0.0.0.0
  Type      ID          Adv Rtr      Seq          Age  Cksum  Len
Router      0.0.0.1        10.255.4.85  0x80000003   885  0xa697  40
Router      *0.0.0.1        10.255.4.93  0x80000002   953  0xc677  40
InterArPfx  *0.0.0.2        10.255.4.93  0x80000001   910  0xb96f  44
InterArRtr  *0.0.0.1        10.255.4.93  0x80000001   910  0xe159  32
IntraArPfx  *0.0.0.1        10.255.4.93  0x80000002   432  0x788f  72

    OSPF3 link state database, area 0.0.0.1
  Type      ID          Adv Rtr      Seq          Age  Cksum  Len
Router      *0.0.0.1        10.255.4.93  0x80000003   916  0xea40  40
Router      0.0.0.1        10.255.4.97  0x80000006   851  0xc95b  40
Network     0.0.0.2        10.255.4.97  0x80000002   916  0x4598  32
InterArPfx  *0.0.0.1        10.255.4.93  0x80000002   117  0xa980  44
InterArPfx  *0.0.0.2        10.255.4.93  0x80000002    62  0xd47e  44
NSSA        0.0.0.1        10.255.4.97  0x80000002   362  0x45ee  44
IntraArPfx  0.0.0.1        10.255.4.97  0x80000006   851  0x2f77  52

    OSPF3 AS SCOPE link state database
  Type      ID          Adv Rtr      Seq          Age  Cksum  Len
Extern     0.0.0.1        10.255.4.85  0x80000002    63  0x9b86  44
Extern     *0.0.0.1        10.255.4.93  0x80000001   910  0x59c9  44

    OSPF3 Link-Local link state database, interface ge-1/3/0.0
  Type      ID          Adv Rtr      Seq          Age  Cksum  Len
Link       *0.0.0.2        10.255.4.93  0x80000003   916  0x4dab  64
```

show ospf3 database extensive

```
user@host> show ospf3 database extensive
    OSPF3 link state database, area 0.0.0.0
  Type      ID          Adv Rtr      Seq          Age  Cksum  Len
Router      0.0.0.1        10.255.4.85  0x80000003  1028  0xa697  40
  bits 0x2, Options 0x13
  Type PointToPoint (1), Metric 10
    Loc-If-Id 2, Nbr-If-Id 3, Nbr-Rtr-Id 10.255.4.93
  Aging timer 00:42:51
  Installed 00:17:05 ago, expires in 00:42:52, sent 02:37:54 ago
Router      *0.0.0.1        10.255.4.93  0x80000002  1096  0xc677  40
  bits 0x3, Options 0x13
  Type PointToPoint (1), Metric 10
    Loc-If-Id 3, Nbr-If-Id 2, Nbr-Rtr-Id 10.255.4.85
  Gen timer 00:00:40
  Aging timer 00:41:44
  Installed 00:18:16 ago, expires in 00:41:44, sent 00:18:14 ago
  Ours
InterArPfx  *0.0.0.2        10.255.4.93  0x80000001  1053  0xb96f  44
  Prefix feee::10:10:2:0/126
  Prefix-options 0x0, Metric 10
  Gen timer 00:17:02
  Aging timer 00:42:26
  Installed 00:17:33 ago, expires in 00:42:27, sent 00:17:31 ago
  Ours
InterArPfx  *0.0.0.3        10.255.4.93  0x80000001  1053  0x71d3  44
  Prefix feee::10:255:4:97/128
  Prefix-options 0x0, Metric 10
  Gen timer 00:21:07
  Aging timer 00:42:26
```

```

    Installed 00:17:33 ago, expires in 00:42:27, sent 00:17:31 ago
    Ours
InterArRtr *0.0.0.1          10.255.4.93      0x80000001  1053  0xe159  32
  Dest-router-id 10.255.4.97, Options 0x19, Metric 10
  Gen timer 00:29:18
  Aging timer 00:42:26
  Installed 00:17:33 ago, expires in 00:42:27, sent 00:17:31 ago
  Ours
IntraArPfx 0.0.0.1          10.255.4.85      0x80000002  1028  0x2403  72
  Ref-lsa-type Router, Ref-lsa-id 0.0.0.0, Ref-router-id 10.255.4.85
  Prefix-count 2
  Prefix feee::10:255:4:85/128
    Prefix-options 0x2, Metric 0
  Prefix feee::10:10:1:0/126
    Prefix-options 0x0, Metric 10
  Aging timer 00:42:51
  Installed 00:17:05 ago, expires in 00:42:52, sent 02:37:54 ago
IntraArPfx *0.0.0.1          10.255.4.93      0x80000002   575  0x788f  72
  Ref-lsa-type Router, Ref-lsa-id 0.0.0.0, Ref-router-id 10.255.4.93
  Prefix-count 2
  Prefix feee::10:255:4:93/128
    Prefix-options 0x2, Metric 0
  Prefix feee::10:10:1:0/126
    Prefix-options 0x0, Metric 10
  Gen timer 00:33:23
  Aging timer 00:50:24
  Installed 00:09:35 ago, expires in 00:50:25, sent 00:09:33 ago
  OSPF3 link state database, area 0.0.0.1
Type      ID              Adv Rtr          Seq            Age  Cksum  Len
Router    *0.0.0.1          10.255.4.93      0x80000003     1059  0xea40  40
  bits 0x3, Options 0x19
  Type Transit (2), Metric 10
    Loc-If-Id 2, Nbr-If-Id 2, Nbr-Rtr-Id 10.255.4.97
  Gen timer 00:08:51
  Aging timer 00:42:20
  Installed 00:17:39 ago, expires in 00:42:21, sent 00:17:37 ago
Router     0.0.0.1          10.255.4.97      0x80000006     994  0xc95b  40
  bits 0x2, Options 0x19
  Type Transit (2), Metric 10
    Loc-If-Id 2, Nbr-If-Id 2, Nbr-Rtr-Id 10.255.4.97
  Aging timer 00:43:25
  Installed 00:16:31 ago, expires in 00:43:26, sent 02:37:54 ago
Network    0.0.0.2          10.255.4.97      0x80000002     1059  0x4598  32
  Options 0x11
  Attached router 10.255.4.97
  Attached router 10.255.4.93
  Aging timer 00:42:20
  Installed 00:17:36 ago, expires in 00:42:21, sent 02:37:54 ago
InterArPfx *0.0.0.1          10.255.4.93      0x80000002   260  0xa980  44
  Prefix feee::10:10:1:0/126
  Prefix-options 0x0, Metric 10
  Gen timer 00:45:39
  Aging timer 00:55:39
  Installed 00:04:20 ago, expires in 00:55:40, sent 00:04:18 ago
  Ours
InterArPfx *0.0.0.2          10.255.4.93      0x80000002   205  0xd47e  44
  Prefix feee::10:255:4:93/128
  Prefix-options 0x0, Metric 0
  Gen timer 00:46:35
  Aging timer 00:56:35
  Installed 00:03:25 ago, expires in 00:56:35, sent 00:03:23 ago

```

```

Ours
InterArPfx *0.0.0.3          10.255.4.93      0x80000001 1089 0x9bbb 44
Prefix feee::10:255:4:85/128
Prefix-options 0x0, Metric 10
Gen timer 00:04:46
Aging timer 00:41:51
Installed 00:18:09 ago, expires in 00:41:51, sent 00:17:43 ago
Ours
NSSA      0.0.0.1          10.255.4.97      0x80000002 505 0x45ee 44
Prefix feee::200:200:1:0/124
Prefix-options 0x8, Metric 10, Type 2,
Aging timer 00:51:35
Installed 00:08:22 ago, expires in 00:51:35, sent 02:37:54 ago
IntraArPfx 0.0.0.1          10.255.4.97      0x80000006 994 0x2f77 52
Ref-lsa-type Router, Ref-lsa-id 0.0.0.0, Ref-router-id 10.255.4.97
Prefix-count 1
Prefix feee::10:255:4:97/128
Prefix-options 0x2, Metric 0
Aging timer 00:43:25
Installed 00:16:31 ago, expires in 00:43:26, sent 02:37:54 ago
IntraArPfx 0.0.0.3          10.255.4.97      0x80000002 1059 0x4446 52
Ref-lsa-type Network, Ref-lsa-id 0.0.0.2, Ref-router-id 10.255.4.97
Prefix-count 1
Prefix feee::10:10:2:0/126
Prefix-options 0x0, Metric 0
Aging timer 00:42:20
Installed 00:17:36 ago, expires in 00:42:21, sent 02:37:54 ago
OSPF3 AS SCOPE link state database
Type      ID              Adv Rtr          Seq              Age  Cksum  Len
Extern    0.0.0.1              10.255.4.85      0x80000002      206  0x9b86 44
Prefix feee::100:100:1:0/124
Prefix-options 0x0, Metric 20, Type 2,
Aging timer 00:56:34
Installed 00:03:23 ago, expires in 00:56:34, sent 02:37:54 ago
Extern    *0.0.0.1              10.255.4.93      0x80000001      1053 0x59c9 44
Prefix feee::200:200:1:0/124
Prefix-options 0x0, Metric 10, Type 2,
Gen timer 00:25:12
Aging timer 00:42:26
Installed 00:17:33 ago, expires in 00:42:27, sent 00:17:31 ago

OSPF3 Link-Local link state database, interface ge-1/3/0.0
Type      ID              Adv Rtr          Seq              Age  Cksum  Len
Link      *0.0.0.2              10.255.4.93      0x80000003      1059 0x4dab 64
fe80::290:69ff:fe39:1cdb
Options 0x11, priority 128
Prefix-count 1
Prefix feee::10:10:2:0/126 Prefix-options 0x0
Gen timer 00:12:56
Aging timer 00:42:20
Installed 00:17:39 ago, expires in 00:42:21, sent 00:17:37 ago
Link      0.0.0.2              10.255.4.97      0x80000003      205  0xa87d 64
fe80::290:69ff:fe38:883e
Options 0x11, priority 128
Prefix-count 1
Prefix feee::10:10:2:0/126 Prefix-options 0x0
Aging timer 00:56:35
Installed 00:03:22 ago, expires in 00:56:35, sent 02:37:54 ago

OSPF3 Link-Local link state database, interface so-2/2/0.0
Type      ID              Adv Rtr          Seq              Age  Cksum  Len

```



```

Link          0.0.0.2          10.255.4.85          0x80000002    506 0x42bb 64
fe80::280:42ff:fe10:f169
Options 0x13, priority 128
Prefix-count 1
Prefix feee::10:10:1:0/126 Prefix-options 0x0
Aging timer 00:51:34
Installed 00:08:23 ago, expires in 00:51:34, sent 02:37:54 ago
Link          *0.0.0.3          10.255.4.93          0x80000002    505 0x6b7a 64
fe80::280:42ff:fe10:f177
Options 0x13, priority 128
Prefix-count 1
Prefix feee::10:10:1:0/126 Prefix-options 0x0
Gen timer 00:37:28
Aging timer 00:51:35
Installed 00:08:25 ago, expires in 00:51:35, sent 00:08:23 ago
Ours

```

show ospf3 database summary

```

user@host> show ospf3 database summary
Area 0.0.0.0:
  2 Router LSAs
  1 InterArPfx LSAs
  1 InterArRtr LSAs
  1 IntraArPfx LSAs
Area 0.0.0.1:
  2 Router LSAs
  1 Network LSAs
  2 InterArPfx LSAs
  1 NSSA LSAs
  1 IntraArPfx LSAs
Externals:
  2 Extern LSAs
Interface ge-1/3/0.0:
  1 Link LSAs
Interface lo0.0:
Interface so-2/2/0.0:
  1 Link LSAs

```

show (ospf | ospf3) interface

| | |
|---|---|
| Syntax | <code>show (ospf ospf3) interface</code> <code><brief detail extensive></code> <code><area <i>area-id</i>></code> <code><interface-name></code> <code><instance <i>instance-name</i>></code> <code><logical-system (all <i>logical-system-name</i>)></code> <code><realm (ipv4-multicast ipv4-unicast ipv6-multicast)></code> |
| Syntax (EX Series Switches and QFX Series) | <code>show (ospf ospf3) interface</code> <code><brief detail extensive></code> <code><area <i>area-id</i>></code> <code><interface-name></code> <code><instance <i>instance-name</i>></code> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. area option introduced in Junos OS Release 9.2. area option introduced in Junos OS Release 9.2 for EX Series switches. realm option introduced in Junos OS Release 9.2. Command introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | Display the status of OSPF interfaces. |
| Options | none —Display standard information about the status of all OSPF interfaces for all routing instances brief detail extensive —(Optional) Display the specified level of output. area <i>area-id</i> —(Optional) Display information about the interfaces that belong to the specified area. <i>interface-name</i> —(Optional) Display information for the specified interface. instance <i>instance-name</i> —(Optional) Display all OSPF interfaces under the named routing instance. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. realm (ipv4-multicast ipv4-unicast ipv6-multicast) —(OSPFv3 only) (Optional) Display information about the interfaces for the specified OSPFv3 realm, or address family. Use the realm option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default. |
| Required Privilege Level | view |
| List of Sample Output | show ospf interface brief on page 443 show ospf interface detail on page 443 show ospf3 interface detail on page 443 |

[show ospf interface detail \(When Multiarea Adjacency Is Configured\) on page 443](#)
[show ospf interface area area-id on page 444](#)
[show ospf interface extensive \(When Flooding Reduction Is Enabled\) on page 444](#)
[show ospf interface extensive \(When LDP Synchronization Is Configured\) on page 445](#)

Output Fields Table 113 on page 441 lists the output fields for the **show (ospf | ospf3) interface** command. Output fields are listed in the approximate order in which they appear.

Table 113: show (ospf | ospf3) interface Output Fields

| Field Name | Field Description | Level of Output |
|-------------------------|---|-------------------------|
| Interface | Name of the interface running OSPF version 2 or OSPF version 3. | All levels |
| State | State of the interface: BDR , Down , DR , DRother , Loop , PtToPt , or Waiting . | All levels |
| Area | Number of the area that the interface is in. | All levels |
| DR ID | Address of the area's designated router. | All levels |
| BDR ID | Backup designated router for a particular subnet. | All levels |
| Nbrs | Number of neighbors on this interface. | All levels |
| Type | Type of interface: LAN , NBMA , P2MP , P2P , or Virtual . | detail extensive |
| Address | IP address of the neighbor. | detail extensive |
| Mask | Netmask of the neighbor. | detail extensive |
| Prefix-length | (OSPFv3) IPv6 prefix length, in bits. | detail extensive |
| OSPF3-Intf-Index | (OSPFv3) OSPF version 3 interface index. | detail extensive |
| MTU | Interface maximum transmission unit (MTU). | detail extensive |
| Cost | Interface cost (metric). | detail extensive |
| DR addr | Address of the designated router. | detail extensive |
| BDR addr | Address of the backup designated router. | detail extensive |
| Adj count | Number of adjacent neighbors. | detail extensive |
| Secondary | Indicates that this interface is configured as a secondary interface for this area. This interface can belong to more than one area, but can be designated as a primary interface for only one area. | detail extensive |
| Flood Reduction | Indicates that this interface is configured with flooding reduction. All self-originated LSAs from this interface are initially sent with the DoNotAge bit set. As a result, LSAs are refreshed only when a change occurs. | extensive |

Table 113: show (ospf | ospf3) interface Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-------------------------------------|---|-------------------------|
| Priority | Router priority used in designated router (DR) election on this interface. | detail extensive |
| Flood list | List of link-state advertisements (LSAs) that might be about to flood this interface. | extensive |
| Ack list | Acknowledgment list. List of pending acknowledgments on this interface. | extensive |
| Descriptor list | List of packet descriptors. | extensive |
| Hello | Configured value for the hello timer. | detail extensive |
| Dead | Configured value for the dead timer. | detail extensive |
| Auth type | (OSPFv2) Authentication mechanism for sending and receiving OSPF protocol packets: <ul style="list-style-type: none"> • MD5—The MD5 mechanism is configured in accordance with RFC 2328. • None—No authentication method is configured. • Password—A simple password (RFC 2328) is configured. | detail extensive |
| Topology | (Multiarea adjacency) Name of topology: default or name . | |
| LDP sync state | (OSPFv2 and LDP synchronization) Current state of LDP synchronization: in sync , in holddown , and not supported . | extensive |
| reason | (OSPFv2 and LDP synchronization) Reason for the current state of LDP synchronization. The LDP session might be up or down, or adjacency might be up or down. | extensive |
| config holdtime | (OSPFv2 and LDP synchronization) Configured value of the hold timer. If the state is not synchronized, and the hold time is not infinity, the remaining field displays the number of seconds that remain until the configured hold timer expires. | extensive |
| IPSec SA name | (OSPFv2) Name of the IPSec security association name. | detail extensive |
| Active key ID | (OSPFv2 and MD5) Number from 0 to 255 that uniquely identifies an MD5 key. | detail extensive |
| Start time | (OSPFv2 and MD5) Time at which the routing device starts using an MD5 key to authenticate OSPF packets transmitted on the interface on which this key is configured. To authenticate received OSPF protocol packets, the key becomes effective immediately after the configuration is committed. If the start time option is not configured, the key is effective immediately for send and receive and is displayed as Start time 1970 Jan 01 00:00:00 PST . | detail extensive |
| ReXmit | Configured value for the Retransmit timer. | detail extensive |
| Stub, Not Stub, or Stub NSSA | Type of area. | detail extensive |

Sample Output

show ospf interface brief

```
user@host> show ospf interface brief

```

| Intf | State | Area | DR ID | BDR ID | Nbrs |
|------------|--------|---------|--------------|--------------|------|
| at-5/1/0.0 | PtToPt | 0.0.0.0 | 0.0.0.0 | 0.0.0.0 | 1 |
| ge-2/3/0.0 | DR | 0.0.0.0 | 192.168.4.16 | 192.168.4.15 | 1 |
| lo0.0 | DR | 0.0.0.0 | 192.168.4.16 | 0.0.0.0 | 0 |
| so-0/0/0.0 | Down | 0.0.0.0 | 0.0.0.0 | 0.0.0.0 | 0 |
| so-6/0/1.0 | PtToPt | 0.0.0.0 | 0.0.0.0 | 0.0.0.0 | 1 |
| so-6/0/2.0 | Down | 0.0.0.0 | 0.0.0.0 | 0.0.0.0 | 0 |
| so-6/0/3.0 | PtToPt | 0.0.0.0 | 0.0.0.0 | 0.0.0.0 | 1 |

show ospf interface detail

```
user@host> show ospf interface detail

```

| Interface | State | Area | DR ID | BDR ID | Nbrs |
|------------|-------|---------|---------------|----------------|------|
| fe-0/0/1.0 | BDR | 0.0.0.0 | 192.168.37.12 | 10.255.245.215 | 1 |

Type LAN, address 192.168.37.11, Mask 255.255.255.248, MTU 4460, Cost 40
 DR addr 192.168.37.12, BDR addr 192.168.37.11, Adj count 1, Priority 128
 Hello 10, Dead 40, ReXmit 5, Not Stub

| Interface | State | Area | DR ID | BDR ID | Nbrs |
|------------|--------|---------|---------|---------|------|
| tl-0/2/1.0 | PtToPt | 0.0.0.0 | 0.0.0.0 | 0.0.0.0 | 0 |

Type P2P, Address 0.0.0.0, Mask 0.0.0.0, MTU 1500, Cost 2604
 Adj count 0
 Hello 10, Dead 40, ReXmit 5, Not Stub
 Auth type: MD5, Active key ID 3, Start time 2002 Nov 19 10:00:00 PST
 IPsec SA Name: sa

show ospf3 interface detail

```
user@host> show ospf3 interface so-0/0/3.0 detail

```

| Interface | State | Area | DR-ID | BDR-ID | Nbrs |
|------------|--------|---------|---------|---------|------|
| so-0/0/3.0 | PtToPt | 0.0.0.0 | 0.0.0.0 | 0.0.0.0 | 1 |

Address fe80::2a0:a5ff:fe28:1dfc, Prefix-length 64
 OSPF3-Intf-index 1, Type P2P, MTU 4470, Cost 12, Adj-count 1
 Hello 10, Dead 40, ReXmit 5, Not Stub

show ospf interface detail (When Multiarea Adjacency Is Configured)

```
user@host> show ospf interface detail
regress@router> show ospf interface detail

```

| Interface | State | Area | DR ID | BDR ID | Nbrs |
|-----------|-------|---------|--------------|---------|------|
| lo0.0 | DR | 0.0.0.0 | 10.255.245.2 | 0.0.0.0 | 0 |

Type: LAN, Address: 127.0.0.1, Mask: 255.255.255.255, MTU: 65535, Cost: 0
 DR addr: 127.0.0.1, Adj count: 0, Priority: 128
 Hello: 10, Dead: 40, ReXmit: 5, Not Stub
 Auth type: None
 Topology default (ID 0) -> Cost: 0

| Interface | State | Area | DR ID | BDR ID | Nbrs |
|-----------|-------|---------|--------------|---------|------|
| lo0.0 | DR | 0.0.0.0 | 10.255.245.2 | 0.0.0.0 | 0 |

Type: LAN, Address: 10.255.245.2, Mask: 255.255.255.255, MTU: 65535, Cost: 0
 DR addr: 10.255.245.2, Adj count: 0, Priority: 128
 Hello: 10, Dead: 40, ReXmit: 5, Not Stub
 Auth type: None
 Topology default (ID 0) -> Cost: 0

| Interface | State | Area | DR ID | BDR ID | Nbrs |
|------------|--------|---------|---------|---------|------|
| so-0/0/0.0 | PtToPt | 0.0.0.0 | 0.0.0.0 | 0.0.0.0 | 1 |

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
 Adj count: 1
 Hello: 10, Dead: 40, ReXmit: 5, Not Stub
 Auth type: None
 Topology default (ID 0) -> Cost: 1

| Interface | State | Area | DR ID | BDR ID | Nbrs |
|------------|--------|---------|---------|---------|------|
| so-0/0/0.0 | PtToPt | 0.0.0.0 | 0.0.0.0 | 0.0.0.0 | 0 |

```

Type: P2P, Address: 192.168.37.46, Mask: 255.255.255.254, MTU: 4470, Cost: 1
Adj count: 0, , Passive
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Passive, Cost: 1
so-1/0/0.0      PtToPt  0.0.0.0      0.0.0.0      0.0.0.0      1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-1/0/0.0      PtToPt  0.0.0.0      0.0.0.0      0.0.0.0      0

Type: P2P, Address: 192.168.37.54, Mask: 255.255.255.254, MTU: 4470, Cost: 1
Adj count: 0, , Passive
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Passive, Cost: 1
so-0/0/0.0      PtToPt  1.1.1.1      0.0.0.0      0.0.0.0      1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1, Secondary
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-1/0/0.0      PtToPt  1.1.1.1      0.0.0.0      0.0.0.0      1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1, Secondary
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-0/0/0.0      PtToPt  2.2.2.2      0.0.0.0      0.0.0.0      1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1, Secondary
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-1/0/0.0      PtToPt  2.2.2.2      0.0.0.0      0.0.0.0      1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1, Secondary
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1

```

show ospf interface area area-id

```

user@host> show ospf interface area 1.1.1.1
Interface      State  Area      DR ID      BDR ID      Nbrs
so-0/0/0.0     PtToPt 1.1.1.1   0.0.0.0    0.0.0.0     1
so-1/0/0.0     PtToPt 1.1.1.1   0.0.0.0    0.0.0.0     1

```

show ospf interface extensive (When Flooding Reduction Is Enabled)

```

user@host> show ospf interface extensive
Interface      State  Area      DR ID      BDR ID      Nbrs
fe-0/0/0.0     PtToPt 0.0.0.0    0.0.0.0    0.0.0.0     0

Type: P2P, Address: 10.10.10.1, Mask: 255.255.255.0, MTU: 1500, Cost: 1
Adj count: 0

```

Secondary, Flood Reduction
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1

**show ospf interface
extensive
(When LDP
Synchronization Is
Configured)**

```
user@host> show ospf interface extensive
Interface          State      Area      DR ID      BDR ID
Nbrs
so-1/0/3.0         Down      0.0.0.0    0.0.0.0    0.0.0.0
0
  Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 65535
  Adj count: 0
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  LDP sync state: in holddown, for: 00:00:08, reason: LDP down during config
                    config holdtime: 10 seconds, remaining: 1
```

show (ospf | ospf3) io-statistics

| | |
|---|--|
| Syntax | show (ospf ospf3) io-statistics <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switch and QFX Series) | show (ospf ospf3) io-statistics |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | Display Open Shortest Path First (OSPF) input and output statistics. |
| Options | <p>none—Display OSPF input and output statistics.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> clear (ospf ospf3) statistics on page 401 |
| List of Sample Output | show ospf io-statistics on page 447 |
| Output Fields | <p>Table 114 on page 446 lists the output fields for the show ospf io-statistics command. Output fields are listed in the approximate order in which they appear.</p> |

Table 114: show (ospf | ospf3) io-statistics Output Fields

| Field Name | Field Description |
|------------------------|---|
| Packets read | Number of OSPF packets read since the last time the routing protocol was started. |
| average per run | Total number of packets divided by the total number of times the OSPF read operation is scheduled to run. |
| max run | Maximum number of packets for a given run among all scheduled runs. |
| Receive errors | Number of faulty packets received with errors. |

Sample Output

```
show ospf io-statistics  user@host> show ospf io-statistics

Packets read: 7361, average per run: 1.00, max run: 1
Receive errors:
  None
```

show (ospf | ospf3) log

| | |
|---|---|
| Syntax | <pre>show (ospf ospf3) log <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)> <topology <i>topology-name</i>></pre> |
| Syntax (EX Series Switch and QFX Series) | <pre>show (ospf ospf3) log <instance <i>instance-name</i>> <topology <i>topology-name</i>></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>topology option introduced in Junos OS Release 9.0.</p> <p>topology option introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>realm option introduced in Junos OS Release 9.2.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> |
| Description | Display the entries in the Open Shortest Path First (OSPF) log of SPF calculations. |
| Options | <p>none—Display entries in the OSPF log of SPF calculations for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Display entries for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>topology <i>topology-name</i>—(Optional) (OSPFv2 only) Display entries for the specified topology.</p> <p>realm (ipv4-multicast ipv4-unicast ipv6-multicast)—(OSPFv3 only) (Optional) Display entries for the specified OSPFv3 realm, or address family. Use the realm option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.</p> |
| Required Privilege Level | view |
| List of Sample Output | <p>show ospf log on page 450</p> <p>show ospf log topology voice on page 450</p> |
| Output Fields | <p>Table 115 on page 448 lists the output fields for the show (ospf ospf3) log command. Output fields are listed in the approximate order in which they appear.</p> |

Table 115: show (ospf | ospf3) log Output Fields

| Field Name | Field Description |
|-------------|--|
| When | Time, in weeks (w) and days (d), since the SPF calculation was made. |

Table 115: show (ospf | ospf3) log Output Fields (*continued*)

| Field Name | Field Description |
|------------|---|
| Type | Type of calculation: Cleanup , External , Interarea , NSSA , Redist , SPF , Stub , Total , or Virtuallink . |
| Elapsed | Amount of time, in seconds, that elapsed during the operation, or the time required to complete the SPF calculation. The start time is the time displayed in the When field. |

Sample Output

show ospf log

```
user@host> show ospf log
When                Type                Elapsed

1w4d 17:25:58      Stub                0.000017
1w4d 17:25:58      SPF                0.000070
1w4d 17:25:58      Stub                0.000019
1w4d 17:25:58      Interarea          0.000054
1w4d 17:25:58      External           0.000005
1w4d 17:25:58      Cleanup            0.000203
1w4d 17:25:58      Total              0.000537
1w4d 17:24:48      SPF                0.000125
1w4d 17:24:48      Stub                0.000017
1w4d 17:24:48      SPF                0.000100
1w4d 17:24:48      Stub                0.000016
1w4d 17:24:48      Interarea          0.000056
1w4d 17:24:48      External           0.000005
1w4d 17:24:48      Cleanup            0.000238
1w4d 17:24:48      Total              0.000600
...
```

show ospf log topology voice

```
user@host> show ospf log topology voice
Topology voice SPF log:
```

Last instance of each event type

| When | Type | Elapsed |
|----------|-----------|----------|
| 00:06:11 | SPF | 0.000116 |
| 00:06:11 | Stub | 0.000114 |
| 00:06:11 | Interarea | 0.000126 |
| 00:06:11 | External | 0.000067 |
| 00:06:11 | NSSA | 0.000037 |
| 00:06:11 | Cleanup | 0.000186 |

Maximum length of each event type

| When | Type | Elapsed |
|----------|-----------|----------|
| 00:13:43 | SPF | 0.000140 |
| 00:13:33 | Stub | 0.000116 |
| 00:13:43 | Interarea | 0.000128 |
| 00:13:33 | External | 0.000075 |
| 00:13:38 | NSSA | 0.000039 |
| 00:13:53 | Cleanup | 0.000657 |

Last 100 events

| When | Type | Elapsed |
|----------|-----------|----------|
| 00:13:53 | SPF | 0.000090 |
| 00:13:53 | Stub | 0.000041 |
| 00:13:53 | Interarea | 0.000123 |
| 00:13:53 | External | 0.000040 |
| 00:13:53 | NSSA | 0.000038 |
| 00:13:53 | Cleanup | 0.000657 |
| 00:13:53 | Total | 0.001252 |
| . | | |
| . | | |
| 00:06:11 | SPF | 0.000116 |
| 00:06:11 | Stub | 0.000114 |
| 00:06:11 | Interarea | 0.000126 |
| 00:06:11 | External | 0.000067 |

| | | |
|----------|---------|----------|
| 00:06:11 | NSSA | 0.000037 |
| 00:06:11 | Cleanup | 0.000186 |
| 00:06:11 | Total | 0.000818 |

show (ospf | ospf3) neighbor

| | |
|---|---|
| Syntax | <pre>show (ospf ospf3) neighbor <brief detail extensive> <area <i>area-id</i>> <instance (all <i>instance-name</i>)> <interface <i>interface-name</i>> <logical-system (all <i>logical-system-name</i>)> <neighbor> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)></pre> |
| Syntax (EX Series Switches and QFX Series) | <pre>show (ospf ospf3) neighbor <brief detail extensive> <area <i>area-id</i>> <instance (all <i>instance-name</i>)> <interface <i>interface-name</i>> <neighbor></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>instance all option introduced in Junos OS Release 9.1.</p> <p>instance all option introduced in Junos OS Release 9.1 for EX Series switches.</p> <p>area, interface, and realm options introduced in Junos OS Release 9.2.</p> <p>area and interface options introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> |
| Description | <p>Display information about OSPF neighbors.</p> <p>CPU utilization might increase while the device learns its OSPF neighbors. We recommend that you use the show (ospf ospf3) neighbor command after the device learns and establishes OSPF neighbor adjacencies. Depending on the size of your network, this might take several minutes. If you receive a “timeout communicating with routing daemon” error when using the show (ospf ospf3) neighbor command, wait several minutes before attempting to use the command again. This is not a critical system error, but you might experience a delay in using the CLI.</p> |
| Options | <p>none—Display standard information about all OSPF neighbors for all routing instances.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>area <i>area-id</i>—(Optional) Display information about the OSPF neighbors for the specified area.</p> <p>instance (all <i>instance-name</i>)—(Optional) Display all OSPF interfaces for all routing instances or under the named routing instance.</p> <p>interface <i>interface-name</i>—(Optional) Display information about OSPF neighbors for the specified logical interface.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |

neighbor—(Optional) Display information about the specified OSPF neighbor.

realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)—(OSPFv3 only) (Optional) Display information about the OSPF neighbors for the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

Required Privilege Level view

Related Documentation

- [clear \(ospf | ospf3\) neighbor on page 398](#)

List of Sample Output

- [show ospf neighbor brief on page 456](#)
- [show ospf neighbor detail on page 456](#)
- [show ospf neighbor extensive on page 456](#)
- [show ospf3 neighbor detail on page 457](#)
- [show ospf neighbor area area-id on page 457](#)
- [show ospf neighbor interface interface-name on page 457](#)
- [show ospf3 neighbor instance all \(OSPFv3 Multiple Family Address Support Enabled\) on page 458](#)

Output Fields [Table 116 on page 453](#) lists the output fields for the **show (ospf | ospf3) neighbor** command. Output fields are listed in the approximate order in which they appear.

Table 116: show (ospf | ospf3) neighbor Output Fields

| Field Name | Field Description | Level of Output |
|------------------|--|-----------------|
| Address | Address of the neighbor. | All levels |
| Interface | Interface through which the neighbor is reachable. | All levels |

Table 116: show (ospf | ospf3) neighbor Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------------------------|--|-------------------------|
| State | <p>State of the neighbor:</p> <ul style="list-style-type: none"> • Attempt—Valid only for neighbors attached to nonbroadcast networks. It indicates that no recent information has been received from the neighbor, but that a more concerted effort must be made to contact the neighbor. • Down—Initial state of a neighbor conversation. It indicates that no recent information has been received from the neighbor. Hello packets might continue to be sent to neighbors in the Down state, although at a reduced frequency. • Exchange—Routing device is describing its entire link-state database by sending database description packets to the neighbor. Each packet has a sequence number and is explicitly acknowledged. • ExStart—First step in creating an adjacency between the two neighboring routing devices. The goal of this step is to determine which routing device is the master, and to determine the initial sequence number. • Full—Neighboring routing devices are fully adjacent. These adjacencies appear in router link and network link advertisements. • Init—A hello packet has recently been sent by the neighbor. However, bidirectional communication has not yet been established with the neighbor. This state might occur, for example, because the routing device itself did not appear in the neighbor's hello packet. • Loading—Link-state request packets are sent to the neighbor to acquire more recent advertisements that have been discovered (but not yet received) in the Exchange state. • 2Way—Communication between the two routing devices is bidirectional. This state has been ensured by the operation of the Hello Protocol. This is the most advanced state short of beginning adjacency establishment. The (backup) designated router is selected from the set of neighbors in state 2Way or greater. | All levels |
| ID | Router ID of the neighbor. | All levels |
| Pri | Priority of the neighbor to become the designated router. | All levels |
| Dead | Number of seconds until the neighbor becomes unreachable. | All levels |
| Link state acknowledgment list | Number of link-state acknowledgments received. | extensive |
| Link state retransmission list | <p>Total number of link-state advertisements retransmitted. For extensive output only, the following information is also displayed:</p> <ul style="list-style-type: none"> • Type—Type of link advertisement: ASBR, Sum, Extern, Network, NSSA, OpagArea, Router, or Summary. • LSA ID—LSA identifier included in the advertisement. An asterisk preceding the identifier marks database entries that originated from the local routing device. • Adv rtr—Address of the routing device that sent the advertisement. • Seq—Link sequence number of the advertisement. | detail extensive |

Table 116: show (ospf | ospf3) neighbor Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-------------------------|--|-------------------------|
| Neighbor-address | (OSPFv3 only) If the neighbor uses virtual links, the Neighbor-address is the site-local, local, or global address. If the neighbor uses a physical interface, the Neighbor-address is an IPv6 link-local address. | detail extensive |
| area | Area that the neighbor is in. | detail extensive |
| OSPF3-Intf-Index | (OSPFv3 only) Displays the OSPFv3 interface index. | detail extensive |
| opt | Option bits received in the hello packets from the neighbor. | detail extensive |
| DR or DR-ID | Address of the designated router. | detail extensive |
| BDR or BDR-ID | Address of the backup designated router. | detail extensive |
| Up | Length of time since the neighbor came up. | detail extensive |
| adjacent | Length of time since the adjacency with the neighbor was established. | detail extensive |

Sample Output

show ospf neighbor brief

```

user@host> show ospf neighbor brief
  Address      Intf      State      ID          Pri  Dead
192.168.254.225 fxp3.0    2Way       10.250.240.32 128  36
192.168.254.230 fxp3.0    Full       10.250.240.8  128  38
192.168.254.229 fxp3.0    Full       10.250.240.35 128  33
10.1.1.129      fxp2.0    Full       10.250.240.12 128  37
10.1.1.131      fxp2.0    Full       10.250.240.11 128  38
10.1.2.1        fxp1.0    Full       10.250.240.9  128  32
10.1.2.81       fxp0.0    Full       10.250.240.10 128  33

```

show ospf neighbor detail

```

user@host> show ospf neighbor detail
  Address      Interface      State      ID          Pri  Dead
10.5.1.2       ge-1/2/0.1     Full       10.5.1.2    128  37
  area 0.0.0.1, opt 0x42, DR 10.5.1.2, BDR 10.5.1.1
  Up 06:09:28, adjacent 05:17:36
  Link state acknowledgment list:  3 entries

  Link state retransmission list:  9 entries

10.5.10.2      ge-1/2/0.10    ExStart    10.5.1.38   128  34
  area 0.0.0.1, opt 0x42, DR 10.5.10.2, BDR 10.5.10.1
  Up 06:09:28
  master, seq 0xac1530f8, rexmit DBD in 3 sec
  rexmit LSREQ in 0 sec
10.5.11.2      ge-1/2/0.11    Full       10.5.1.42   128  38
  area 0.0.0.1, opt 0x42, DR 10.5.11.2, BDR 10.5.11.1
  Up 06:09:28, adjacent 05:26:46
  Link state retransmission list:  1 entries

10.5.12.2      ge-1/2/0.12    ExStart    10.5.1.46   128  33
  area 0.0.0.1, opt 0x42, DR 10.5.12.2, BDR 10.5.12.1
  Up 06:09:28
  master, seq 0xac188a68, rexmit DBD in 2 sec
  rexmit LSREQ in 0 sec

```

show ospf neighbor extensive

```

user@host> show ospf neighbor extensive
  Address      Interface      State      ID          Pri  Dead
10.5.1.2       ge-1/2/0.1     Full       10.5.1.2    128  33
  area 0.0.0.1, opt 0x42, DR 10.5.1.2, BDR 10.5.1.1
  Up 06:09:42, adjacent 05:17:50
  Link state retransmission list:

  Type      LSA ID          Adv rtr          Seq
  Summary   10.8.56.0       172.25.27.82    0x8000004d
  Router     10.5.1.94       10.5.1.94       0x8000005c
  Network    10.5.24.2       10.5.1.94       0x80000036
  Summary    10.8.57.0       172.25.27.82    0x80000024
  Extern     1.10.90.0       10.8.1.2        0x80000041
  Extern     1.4.109.0       10.6.1.2        0x80000041

```

```

Router 10.5.1.190      10.5.1.190      0x8000005f
Network 10.5.48.2      10.5.1.190      0x8000003d
Summary 10.8.58.0      172.25.27.82    0x8000004d
Extern 1.10.91.0      10.8.1.2        0x80000041
Extern 1.4.110.0      10.6.1.2        0x80000041
Router 10.5.1.18      10.5.1.18      0x8000005f
Network 10.5.5.2      10.5.1.18      0x80000033
Summary 10.8.59.0      172.25.27.82    0x8000003a
Summary 10.8.62.0      172.25.27.82    0x80000025

10.5.10.2      ge-1/2/0.10      ExStart 10.5.1.38      128 38
area 0.0.0.1, opt 0x42, DR 10.5.10.2, BDR 10.5.10.1
Up 06:09:42
master, seq 0xac1530f8, retransmit DBD in 2 sec
retransmit LSREQ in 0 sec
10.5.11.2      ge-1/2/0.11      Full 10.5.1.42      128 33
area 0.0.0.1, opt 0x42, DR 10.5.11.2, BDR 10.5.11.1
Up 06:09:42, adjacent 05:27:00
Link state retransmission list:

Type      LSA ID      Adv rtr      Seq
Summary 10.8.58.0      172.25.27.82    0x8000004d
Extern 1.10.91.0      10.8.1.2        0x80000041
Extern 1.1.247.0      10.5.1.2        0x8000003f
Extern 1.4.110.0      10.6.1.2        0x80000041
Router 10.5.1.18      10.5.1.18      0x8000005f
Network 10.5.5.2      10.5.1.18      0x80000033
Summary 10.8.59.0      172.25.27.82    0x8000003a

```

show ospf3 neighbor detail

```

user@host> show ospf3 neighbor detail
ID      Interface      State      Pri  Dead
10.255.71.13  fe-0/0/2.0      Full      128  30
Neighbor-address fe80::290:69ff:fe9b:e002
area 0.0.0.0, opt 0x13, OSPF3-Intf-Index 2
DR-ID 10.255.71.13, BDR-ID 10.255.71.12
Up 02:51:43, adjacent 02:51:43

```

show ospf neighbor area area-id

```

user@host >show ospf neighbor area 1.1.1.1
Address      Interface      State      ID      Pri  Dead
192.168.37.47  so-0/0/0.0      Full      10.255.245.4  128  33
Area 1.1.1.1
192.168.37.55  so-1/0/0.0      Full      10.255.245.5  128  37
Area 1.1.1.1

```

**show ospf neighbor
interface
interface-name**

```
user@host > show ospf neighbor interface so-0/0/0.0
Address      Interface      State    ID           Pri    Dead
192.168.37.47 so-0/0/0.0     Full    10.255.245.4 128    37
  Area 0.0.0.0
192.168.37.47 so-0/0/0.0     Full    10.255.245.4 128    33
  Area 1.1.1.1
192.168.37.47 so-0/0/0.0     Full    10.255.245.4 128    32
  Area 2.2.2.2
```

**show ospf3 neighbor
instance all (OSPFv3
Multiple Family
Address Support
Enabled)**

```
user @host > show ospf3 neighbor instance all
Instance: ina
  Realm: ipv6-unicast
    ID      Interface      State    Pri    Dead
    100.1.1.1 fe-0/0/2.0     Full    128    37
    Neighbor-address fe80::217:cb00:c87c:8c03
  Instance: inb
  Realm: ipv4-unicast
    ID      Interface      State    Pri    Dead
    100.1.2.1 fe-0/0/2.1     Full    128    33
    Neighbor-address fe80::217:cb00:c97c:8c03
```

show (ospf | ospf3) overview

| | |
|---|--|
| Syntax | show (ospf ospf3) overview <brief extensive> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)> |
| Syntax (EX Series Switch and QFX Series) | show (ospf ospf3) overview <brief extensive> <instance <i>instance-name</i> > |
| Release Information | Command introduced in Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. realm option introduced in Junos OS Release 9.2. Database protection introduced in Junos 10.2. Command introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | Display Open Shortest Path First (OSPF) overview information. |
| Options | <p>none—Display standard information about all OSPF neighbors for all routing instances.</p> <p>brief extensive—(Optional) Display the specified level of output.</p> <p>instance <i>instance-name</i>—(Optional) Display all OSPF interfaces under the named routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>realm (ipv4-multicast ipv4-unicast ipv6-multicast)—(Optional) (OSPFv3 only) Display information about the specified OSPFv3 realm, or address family. Use the realm option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.</p> |
| Required Privilege Level | view |
| List of Sample Output | show ospf overview on page 462 show ospf overview (With Database Protection) on page 462 show ospf3 overview (With Database Protection) on page 462 show ospf overview extensive on page 463 |
| Output Fields | Table 117 on page 459 lists the output fields for the show ospf overview command. Output fields are listed in the approximate order in which they appear. |

Table 117: show ospf overview Output Fields

| Field name | Field Description | Level of Output |
|------------|------------------------|-----------------|
| Instance | OSPF routing instance. | All levels |

Table 117: show ospf overview Output Fields (*continued*)

| Field name | Field Description | Level of Output |
|----------------------------------|--|-----------------|
| Router ID | Router ID of the routing device. | All levels |
| Route table index | Route table index. | All levels |
| Configured overload | Overload capability is enabled. If the overload timer is also configured, display the time that remains before it is set to expire. This field is not displayed after the timer expires. | All levels |
| Topology | Topology identifier. | All levels |
| Prefix export count | Number of prefixes exported into OSPF. | All levels |
| Full SPF runs | Number of complete Shortest Path First calculations. | All levels |
| SPF delay | Delay before performing consecutive Shortest Path First calculations. | All levels |
| SPF holddown | Delay before performing additional Shortest Path First (SPF) calculations after the maximum number of consecutive SPF calculations is reached. | All levels |
| SPF rapid runs | Maximum number of Shortest Path First calculations that can be performed in succession before the hold-down timer begins. | All levels |
| LSA refresh time | Refresh period for link-state advertisement (in minutes). | All levels |
| Database protection state | Current state of database protection. | All levels |
| Warning threshold | Threshold at which a warning message is logged (percentage of maximum LSA count). | All levels |
| Non self-generated LSAs | Number of LSAs whose router ID is not equal to the local router ID: Current , Warning (threshold), and Allowed . | All levels |
| Ignore time | How long the database has been in the ignore state. | All levels |
| Reset time | How long the database must stay out of the ignore or isolated state before it returns to normal operations. | All levels |
| Ignore count | Number of times the database has been in the ignore state: Current and Allowed . | All levels |
| Restart | Graceful restart capability: enabled or disabled . | All levels |
| Restart duration | Time period for complete reacquisition of OSPF neighbors. | All levels |
| Restart grace period | Time period for which the neighbors should consider the restarting routing device as part of the topology. | All levels |

Table 117: show ospf overview Output Fields (*continued*)

| Field name | Field Description | Level of Output |
|-------------------------------|---|------------------|
| Graceful restart helper mode | (OSPFv2) Standard graceful restart helper capability (based on RFC 3623): enabled or disabled . | All levels |
| Restart-signaling helper mode | (OSPFv2) Restart signaling-based graceful restart helper capability (based on RFC 4811, RFC 4812, and RFC 4813): enabled or disabled . | All levels |
| Helper mode | (OSPFv3) Graceful restart helper capability: enabled or disabled . | All levels |
| Trace options | OSPF-specific trace options. | extensive |
| Trace file | Name of the file to receive the output of the tracing operation. | extensive |
| Area | Area number. Area 0.0.0.0 is the backbone area. | All levels |
| Stub type | Stub type of area: Normal Stub , Not Stub , or Not so Stubby Stub . | All levels |
| Authentication Type | Type of authentication: None , Password , or MD5 . NOTE: The Authentication Type field refers to the authentication configured at the [edit protocols ospf area area-id] level. Any authentication configured for an interface in this area will not affect the value of this field. | All levels |
| Area border routers | Number of area border routers. | All levels |
| Neighbors | Number of autonomous system boundary routers. | All levels |

Sample Output

show ospf overview

```
user@host> show ospf overview
Instance: master
  Router ID: 10.255.245.6
  Route table index: 0
  Configured overload, expires in 118 seconds
  LSA refresh time: 50 minutes
  Restart: Enabled
    Restart duration: 20 sec
    Restart grace period: 40 sec
    Helper mode: enabled
  Area: 0.0.0.0
    Stub type: Not Stub
    Authentication Type: None
    Area border routers: 0, AS boundary routers: 0
  Neighbors
    Up (in full state): 0
  Topology: default (ID 0)
  Prefix export count: 0
  Full SPF runs: 1
  SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
```

show ospf overview (With Database Protection)

```
user@host> show ospf overview
Instance: master
  Router ID: 10.255.112.218
  Route table index: 0
  LSA refresh time: 50 minutes
  Traffic engineering
  Restart: Enabled
    Restart duration: 180 sec
    Restart grace period: 210 sec
    Graceful restart helper mode: Enabled
    Restart-signaling helper mode: Enabled
  Database protection state: Normal
    Warning threshold: 70 percent
    Non self-generated LSAs: Current 582, Warning 700, Allowed 1000
    Ignore time: 30, Reset time: 60
    Ignore count: Current 0, Allowed 1
  Area: 0.0.0.0
    Stub type: Not Stub
    Authentication Type: None
    Area border routers: 0, AS boundary routers: 0
  Neighbors
    Up (in full state): 160
  Topology: default (ID 0)
  Prefix export count: 0
  Full SPF runs: 70
  SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
  Backup SPF: Not Needed
```

show ospf3 overview (With Database Protection)

```
user@host> show ospf3 overview
Instance: master
  Router ID: 10.255.112.128
  Route table index: 0
  LSA refresh time: 50 minutes
  Database protection state: Normal
```



```

Warning threshold: 80 percent
Non self-generated LSAs: Current 3, Warning 8, Allowed 10
Ignore time: 30, Reset time: 60
Ignore count: Current 0, Allowed 2
Area: 0.0.0.0
Stub type: Not Stub
Area border routers: 0, AS boundary routers: 0
Neighbors
  Up (in full state): 1
Topology: default (ID 0)
Prefix export count: 0
Full SPF runs: 7
SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
Backup SPF: Not Needed

```

show ospf overview extensive

```

user@host> show ospf overview extensive
Instance: master
Router ID: 1.1.1.103
Route table index: 0
Full SPF runs: 13, SPF delay: 0.200000 sec
LSA refresh time: 50 minutes
Restart: Disabled
Trace options: lsa
Trace file: /var/log/ospf size 131072 files 10
Area: 0.0.0.0
Stub type: Not Stub
Authentication Type: None
Area border routers: 0, AS boundary routers: 0
Neighbors
  Up (in full state): 1

```

show (ospf | ospf3) route

| | |
|---|--|
| Syntax | <pre>show (ospf ospf3) route <brief detail extensive> <abr asbr extern inter intra> <destination> <instance (default ipv4-multicast <i>instance-name</i>)> <logical-system (default ipv4-multicast <i>logical-system-name</i>)> <network> <no-backup-coverage> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)> <router> <topology (default ipv4-multicast <i>topology-name</i>)> <transit></pre> |
| Syntax (EX Series Switch and QFX Series) | <pre>show (ospf ospf3) route <brief detail extensive> <abr asbr extern inter intra> <destination> <instance <i>instance-name</i> <network> <no-backup-coverage> <router> <topology (default ipv4-multicast <i>topology-name</i>)> <transit></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>topology option introduced in Junos OS Release 9.0.</p> <p>realm option introduced in Junos OS Release 9.2.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> |
| Description | Display the entries in the Open Shortest Path First (OSPF) routing table. |
| Options | <p>none—Display standard information about all entries in the OSPF routing table for all routing instances and all topologies.</p> <p>destination—Display routes to the specified IP address (with optional destination prefix length).</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>abr—(Optional) Display routes to area border routers.</p> <p>asbr—(Optional) Display routes to autonomous system border routers.</p> <p>extern—(Optional) Display external routes.</p> <p>inter—(Optional) Display interarea routes.</p> <p>intra—(Optional) Display intra-area routes.</p> |

instance (**default** | **ipv4-multicast** | *instance-name*)—(Optional) Display entries for the default routing instance, the IPv4 multicast routing instance, or for the specified routing instance.

logical-system (**default** | **ipv4-multicast** | *logical-system-name*)—(Optional) Perform this operation on the default logical system, the IPv4 multicast logical system, or on a particular logical system.

network—(Optional) Display routes to networks.

no-backup-coverage—(Optional) Display routes with no backup coverage.

realm (**ipv4-multicast** | **ipv4-unicast** | **ipv6-multicast**)—(OSPFv3 only) (Optional) Display entries in the routing table for the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

router—(Optional) Display routes to all routers.

topology (**default** | **ipv4-multicast** | *topology-name*)—(OSPFv2 only) (Optional) Display routes for the default OSPF topology, IPv4 multicast topology, or for a particular topology.

transit—(Optional) (OSPFv3 only) Display OSPFv3 routes to pseudonodes.

Required Privilege Level

view

List of Sample Output

[show ospf route on page 468](#)
[show ospf route detail on page 468](#)
[show ospf3 route on page 468](#)
[show ospf3 route detail on page 468](#)
[show ospf route topology voice on page 469](#)

Output Fields

[Table 118 on page 465](#) list the output fields for the **show (ospf | ospf3) route** command. Output fields are listed in the approximate order in which they appear.

Table 118: show (ospf | ospf3) route Output Fields

| Field Name | Field Description | Output Level |
|------------------|--|--------------|
| Topology | Name of the topology. | All levels |
| Prefix | Destination of the route. | All levels |
| Path type | How the route was learned: <ul style="list-style-type: none"> Inter—Interarea route Ext1—External type 1 route Ext2—External type 2 route Intra—Intra-area route | All levels |

Table 118: show (ospf | ospf3) route Output Fields (*continued*)

| Field Name | Field Description | Output Level |
|----------------------------|---|---------------|
| Route type | The type of routing device from which the route was learned: <ul style="list-style-type: none"> • AS BR—Route to AS border router. • Area BR—Route to area border router. • Area/AS BR—Route to router that is both an Area BR and AS BR. • Network—Network router. • Router—Route to a router that is neither an Area BR nor an AS BR. • Transit—(OSPFv3 only) Route to a pseudonode representing a transit network, LAN, or nonbroadcast multiaccess (NBMA) link. • Discard—Route to a summary discard. | All levels |
| NH Type | Next-hop type: LSP or IP . | All levels |
| Metric | Route's metric value. | All levels |
| NH-interface | (OSPFv3 only) Interface through which the route's next hop is reachable. | All levels |
| NH-addr | (OSPFv3 only) IPv6 address of the next hop. | All levels |
| NextHop Interface | (OSPFv2 only) Interface through which the route's next hop is reachable. | All levels |
| Nexthop addr/label | (OSPFv2 only) If the NH Type is IP , then it is the address of the next hop. If the NH Type is LSP , then it is the name of the label-switched path. | All levels |
| Area | Area ID of the route. | detail |
| Origin | Router from which the route was learned. | detail |
| Type 7 | Route was learned through a not-so-stubby area (NSSA) link-state advertisement (LSA). | detail |
| P-bit | Route was learned through NSSA LSA and the propagate bit was set. | detail |
| Fwd NZ | Forwarding address is nonzero. Fwd NZ is only displayed if the route is learned through an NSSA LSA. | detail |
| optional-capability | Optional capabilities propagated in the router LSA. This field is in the output for intra-area router routes only (when Route Type is Area BR , AS BR , Area/AS BR , or Router), not for interarea router routes or network routes. Three bits in this field are defined as follows: <ul style="list-style-type: none"> • 0x4 (V)—Routing device is at the end of a virtual active link. • 0x2 (E)—Routing device is an autonomous system boundary router. • 0x1 (B)—Routing device is an area border router. | detail |

Table 118: show (ospf | ospf3) route Output Fields (*continued*)

| Field Name | Field Description | Output Level |
|-----------------|--|---------------|
| priority | The priority assigned to the prefix: <ul style="list-style-type: none">• high• medium• low <p>NOTE: The priority field applies only to routes of type Network.</p> | detail |

Sample Output

show ospf route

```
user@host> show ospf route
Prefix                Path    Route    NH    Metric  NextHop    Nexthop
                    Type    Type      Type
10.255.71.12          Intra  Router    IP     1       fe-0/0/2.0 192.16.22.86
10.255.71.13/32       Intra  Network    IP     0       lo0.0
192.168.222.84/30     Intra  Network    LSP    1       fe-0/0/2.0 1sp-ab
```

show ospf route detail

```
user@host> show ospf route detail
Topology default Route Table:

Prefix                Path    Route    NH    Metric  NextHop    Nexthop
                    Type    Type      Type
10.255.14.174          Inter  AS BR      IP     210      t1-3/0/1.0
    area 0.0.0.2, origin 10.255.14.185
10.255.14.178          Intra  Router      IP     200      t3-3/1/3.0
    area 0.0.0.2, origin 10.255.14.178, optional-capability 0x0
10.210.1.0/30          Intra  Network      IP     10       t3-3/1/2.0
    area 0.0.0.2, origin 10.255.14.172, priority medium
100.1.1.1/32           Inter  Network      IP     210      t1-3/0/1.0
    area 0.0.0.2, origin 10.255.14.185, priority low
112.3.1.0/24           Ext2   Network      IP      0       t1-3/0/1.0
    area 0.0.0.0, origin 10.255.14.174, priority high
200.3.3.0/30           Inter  Network      IP     220      t1-3/0/1.0
    area 0.0.0.2, origin 10.255.14.185, priority high
```

show ospf3 route

```
user@host> show ospf3 route
Prefix                Path    Route    NH    Metric  NextHop    Nexthop
                    Type    Type      Type
10.255.71.13          Intra  Router      IP      1
    NH-interface fe-0/0/2.0, NH-addr fe80::290:69ff:fe9b:e002
10.255.71.13;0.0.0.2
10.255.245.1           Intra  Router      IP     40      fxp1.1      192.168.36.17
    area 0.0.0.0, origin 10.255.245.1 optional-capability 0x0,
10.255.245.3           Intra  AS BR      IP      1      fxp2.3      192.168.36.34
    area 0.0.0.0, origin 10.255.245.3 optional-capability 0x0,
10.255.245.1/32        Intra  Network      IP     40      fxp1.1      192.168.36.17
    area 0.0.0.0, origin 10.255.245.1, priority high
10.255.245.2/32        Intra  Network      IP      0      lo0.0
    area 0.0.0.0, origin 10.255.245.2, priority medium
10.255.245.3/32        Intra  Network      IP      1      fxp2.3      192.168.36.34
    area 0.0.0.0, origin 10.255.245.3, priority low
    Intra  Transit      IP      1
    NH-interface fe-0/0/2.0
192::168:222:84/126    Intra  Network      IP      1
    NH-interface fe-0/0/2.0
abcd::71:12/128        Intra  Network      IP      0
    NH-interface lo0.0
abcd::71:13/128        Intra  Network      LSP     1
    NH-interface fe-0/0/2.0, NH-addr 1sp-cd
```

show ospf3 route detail

```

user@host> show ospf3 route detail
Prefix                                Path    Route    NH    Metric
                                     type    type    type
10.255.14.174                         Intra   Area/AS BR IP    110
    NH-interface so-1/2/2.0
    Area 0.0.0.0, Origin 10.255.14.174, Optional-capability 0x3
10.255.14.178                         Intra   Router   IP    200
    NH-interface t3-3/1/3.0
    Area 0.0.0.0, Origin 10.255.14.178, Optional-capability 0x0
10.255.14.185;0.0.0.2                 Intra   Transit  IP    200
    NH-interface t1-3/0/1.0
    NH-interface so-1/2/2.0
    Area 0.0.0.0, Origin 10.255.14.185
1000:1:1::1/128                      Inter   Network  IP    110
    NH-interface so-1/2/2.0
    Area 0.0.0.0, Origin 10.255.14.174, Priority low
1001:2:1::/48                        Ext1    Network  IP    110
    NH-interface so-1/2/2.0
    Area 0.0.0.0, Origin 10.255.14.174, Fwd NZ, Priority medium
1002:1:7::/48                        Ext2    Network  IP    0
    NH-interface so-1/2/2.0
    Area 0.0.0.0, Origin 10.255.14.174, Fwd NZ, Priority low
1002:3:4::/48                        Ext2    Network  IP    0
    NH-interface so-1/2/2.0
    Area 0.0.0.0, Origin 10.255.14.174, Fwd NZ, Priority high
abcd::10:255:14:172/128              Intra   Network  IP    0
    NH-interface lo0.0
    Area 0.0.0.0, Origin 10.255.14.172, Priority low

```

show ospf route topology voice

```

user@host show ospf route topology voice
Topology voice Route Table:
Prefix          Path    Route    NH    Metric  NextHop    Nexthop
                Type    Type     Type
10.255.8.2      Intra   Router   IP    1        so-0/2/0.0
10.255.8.3      Intra   Router   IP    2        so-0/2/0.0
10.255.8.1/32   Intra   Network  IP    0        lo0.0
10.255.8.2/32   Intra   Network  IP    1        so-0/2/0.0
10.255.8.3/32   Intra   Network  IP    2        so-0/2/0.0
192.168.8.0/29  Intra   Network  IP    2        so-0/2/0.0
192.168.8.44/30 Intra   Network  IP    2        so-0/2/0.0
192.168.8.46/32 Intra   Network  IP    1        so-0/2/0.0
192.168.8.48/30 Intra   Network  IP    1        so-0/2/1.0
192.168.8.52/30 Intra   Network  IP    2        so-0/2/0.0
192.168.9.44/30 Intra   Network  IP    1        so-0/2/0.0
192.168.9.45/32 Intra   Network  IP    2        so-0/2/0.0

```

show (ospf | ospf3) statistics

| | |
|---|--|
| Syntax | show (ospf ospf3) statistics <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)> |
| Syntax (EX Series Switch and QFX Series) | show (ospf ospf3) statistics <instance <i>instance-name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. realm option introduced in Junos OS Release 9.2. Command introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | Display OSPF statistics. |
| Options | <p>none—Display OSPF statistics for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Display all statistics for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>realm (ipv4-multicast ipv4-unicast ipv6-multicast)—(Optional) (OSPFv3 only) Display all statistics for the specified OSPFv3 realm, or address family. Use the realm option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> clear (ospf ospf3) statistics on page 401 |
| List of Sample Output | show ospf statistics on page 472 show ospf statistics logical-system all on page 472 show ospf3 statistics on page 473 |
| Output Fields | Table 119 on page 470 lists the output fields for the show (ospf ospf3) statistics command. Output fields are listed in the approximate order in which they appear. |

Table 119: show (ospf | ospf3) statistics Output Fields

| Field Name | Field Description |
|---|--|
| Packet type | Type of OSPF packet. |
| Total Sent/Total Received | Total number of packets sent and received. |
| Last 5 seconds Sent/Last 5 seconds Received | Total number of packets sent and received in the last 5 seconds. |

Table 119: show (ospf | ospf3) statistics Output Fields (*continued*)

| Field Name | Field Description |
|--------------------------------|---|
| DBDs retransmitted | Total number of database description packets retransmitted, and number retransmitted in the last 5 seconds. |
| LSAs flooded | Total number of link-state advertisements flooded, and number flooded in the last 5 seconds. |
| LSAs flooded high-prio | <p>Total number of high priority link-state advertisements flooded, and number flooded in the last 5 seconds.</p> <p>A link-state advertisement is deemed a high priority if it has changed since it was last sent.</p> |
| LSAs retransmitted | Total number of link-state advertisements retransmitted, and number retransmitted in the last 5 seconds. |
| LSAs transmitted to nbr | Total number of link-state advertisements transmitted to a neighbor, and number transmitted in the last 5 seconds. |
| LSAs requested | Total number of link-state advertisements requested by neighboring devices, and number requested in the last 5 seconds. |
| LSAs acknowledged | Total number of link-state advertisements acknowledged, and number acknowledged in the last 5 seconds. |
| Flood queue depth | Total number of entries in the extended queue. |
| Total rexmit entries | Total number of retransmission entries waiting to be sent from the OSPF routing instance. |
| db summaries | Total number of database description summaries waiting to be sent from the OSPF routing instance. |
| lsreq entries | Total number of link-state request entries waiting to be sent from the OSPF routing instance. |
| Receive errors | <p>Number and type of receive errors. Some sample receive errors include:</p> <ul style="list-style-type: none"> • mtu mismatches • no interface found • no virtual link found • nssa mismatches • stub area mismatches • subnet mismatches <p>If there are no receive errors, the output displays none.</p> |

Sample Output

show ospf statistics

```

user@host> show ospf statistics
Packet type          Total          Last 5 seconds
                   Sent      Received      Sent      Received
Hello                31         14           2           2
DbD                  9          10           0           0
LSReq                2           2           0           0
LSUpdate             8          16           0           0
LSAck                9           9           0           0

DBDs retransmitted   :           3, last 5 seconds :           0
LSAs flooded         :          12, last 5 seconds :           0
LSAs flooded high-prio :           0, last 5 seconds :           0
LSAs retransmitted   :           0, last 5 seconds :           0
LSAs transmitted to nbr:           3, last 5 seconds :           0
LSAs requested       :           5, last 5 seconds :           0
LSAs acknowledged    :          19, last 5 seconds :           0

Flood queue depth    :           0
Total rexmit entries :           0
db summaries         :           0
lsreq entries        :           0

Receive errors:
  862 no interface found
 115923 no virtual link found

```

show ospf statistics logical-system all

```

user@host> show ospf statistics logical-system all
Logical-system: C
OSPF instance is not running
-----

Logical-system: B

Packet type          Total          Last 5 seconds
                   Sent      Received      Sent      Received
Hello              313740      313653         1           0
DbD                 3           2           0           0
LSReq               1           1           0           0
LSUpdate           2752      1825           0           0
LSAck              1821      2747           0           0

DBDs retransmitted   :           0, last 5 seconds :           0
LSAs flooded         :          2741, last 5 seconds :           0
LSAs flooded high-prio :          10, last 5 seconds :           0
LSAs retransmitted   :           0, last 5 seconds :           0
LSAs transmitted to nbr:           2, last 5 seconds :           0
LSAs requested       :           1, last 5 seconds :           0
LSAs acknowledged    :          1831, last 5 seconds :           0

Flood queue depth    :           0
Total rexmit entries :           0
db summaries         :           0
lsreq entries        :           0

Receive errors:
  None
-----

```

logical-system: A

| Packet type | Total | | Last 5 seconds | |
|-------------|--------|----------|----------------|----------|
| | Sent | Received | Sent | Received |
| Hello | 313698 | 313695 | 0 | 0 |
| DbD | 2 | 3 | 0 | 0 |
| LSReq | 1 | 1 | 0 | 0 |
| LSUpdate | 1825 | 2752 | 0 | 0 |
| LSAck | 2747 | 1821 | 0 | 0 |

| | | | | |
|--------------------------|---|----------------------|---|---|
| DBDs retransmitted | : | 0, last 5 seconds | : | 0 |
| LSAs flooded | : | 1825, last 5 seconds | : | 0 |
| LSAs flooded high-prio | : | 10, last 5 seconds | : | 0 |
| LSAs retransmitted | : | 0, last 5 seconds | : | 0 |
| LSAs transmitted to nbr: | : | 1, last 5 seconds | : | 0 |
| LSAs requested | : | 2, last 5 seconds | : | 0 |
| LSAs acknowledged | : | 2748, last 5 seconds | : | 0 |

| | | |
|----------------------|---|---|
| Flood queue depth | : | 0 |
| Total rexmit entries | : | 0 |
| db summaries | : | 0 |
| lsreq entries | : | 0 |

Receive errors:

None

show ospf3 statistics

```
user@host> show ospf3 statistics
```

| Packet type | Total | | Last 5 seconds | |
|-------------|-------|----------|----------------|----------|
| | Sent | Received | Sent | Received |
| Hello | 0 | 0 | 0 | 0 |
| DbD | 0 | 0 | 0 | 0 |
| LSReq | 0 | 0 | 0 | 0 |
| LSUpdate | 0 | 0 | 0 | 0 |
| LSAck | 0 | 0 | 0 | 0 |

| | | | | |
|--------------------------|---|-------------------|---|---|
| DBDs retransmitted | : | 0, last 5 seconds | : | 0 |
| LSAs flooded | : | 0, last 5 seconds | : | 0 |
| LSAs flooded high-prio | : | 0, last 5 seconds | : | 0 |
| LSAs retransmitted | : | 0, last 5 seconds | : | 0 |
| LSAs transmitted to nbr: | : | 0, last 5 seconds | : | 0 |
| LSAs requested | : | 0, last 5 seconds | : | 0 |
| LSAs acknowledged | : | 0, last 5 seconds | : | 0 |

| | | |
|----------------------|---|---|
| Flood queue depth | : | 0 |
| Total rexmit entries | : | 0 |
| db summaries | : | 0 |
| lsreq entries | : | 0 |

Receive errors:

None

Protocol-Independent Routing Operational Mode Commands

Table 120 on page 475 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot protocol-independent routing properties. Commands are listed in alphabetical order.



NOTE: The `show route` command has a lengthy set of options. Therefore, this chapter describes each option as a separate command. You can, however, combine several options and issue them as single `show route` command. For example, `show route ccc exact`.

The exceptions to this convention are the `show as-path`, `show route damping`, `show route export`, `show route export-vrf-target`, `show route forwarding-table`, `show route instance`, and `show route martians` commands, which cannot be used with any other options (other than level of output options, such as `detail` and `extensive`).

The `show route flow validation` command can only be used with the `table` option.

Table 120: Protocol-Independent Routing Operational Mode Commands

| Task | Command |
|--|--|
| Display known autonomous system (AS) paths. | <code>show as-path</code> |
| Display AS path domain information. | <code>show as-path domain</code> |
| Display AS path summary information. | <code>show as-path summary</code> |
| Display information about the entries in the routing tables. | <code>show route</code> |
| Display routes that are currently active. | <code>show route active-path</code> |
| Display routes transmitted by a particular routing protocol. | <code>show route advertising-protocol</code> |
| Display all information about all routes. | <code>show route all</code> |

Table 120: Protocol-Independent Routing Operational Mode Commands (*continued*)

| Task | Command |
|---|---|
| Display routes containing a specified AS path. | <code>show route aspath-regex</code> |
| Display the best route to the specified address or range of addresses. | <code>show route best</code> |
| Display brief information about the entries in the routing table. | <code>show route brief</code> |
| Display circuit cross-connect (CCC) entries in the Multiprotocol Link Switching (MPLS) routing table. | <code>show route ccc</code> |
| Display routes containing members of a specified BGP community. | <code>show route community</code> |
| Display routes containing members of a specified BGP community based on a particular community name. | <code>show route community-name</code> |
| Display routes that have been damped. | <code>show route damping</code> |
| Display detailed information about the entries in the routing table. | <code>show route detail</code> |
| Display routes that exactly match the specified address or range of addresses. | <code>show route exact</code> |
| Display list of instances or routing tables that are importers or exporters of routes. | <code>show route export</code> |
| Display target communities for which autoexport is currently distributing routes. | <code>show route export vrf-target</code> |
| Display extensive information about the entries in the routing table. | <code>show route extensive</code> |
| Display the best route to an address. | <code>show route flow validation</code> |
| Display the Junos OS forwarding table. | <code>show route forwarding-table</code> |
| Display information about the interfaces in the Junos OS forwarding table. | <code>show route forwarding-table interface-name</code> |
| Display hidden routes only. | <code>show route hidden</code> |
| Display routes that are not preferred. | <code>show route inactive-path</code> |
| Display routes that are currently inactive. | <code>show route inactive-prefix</code> |
| Display routing instance information. | <code>show route instance</code> |

Table 120: Protocol-Independent Routing Operational Mode Commands (*continued*)

| Task | Command |
|--|---|
| Display routes corresponding to a specified label value. | <code>show route label</code> |
| Display routes that form a label-switched path. | <code>show route label-switched-path</code> |
| Display route localization information. | <code>show route localization</code> |
| Display information about martian addresses. | <code>show route martians</code> |
| Display routes that contain the specified next hop. | <code>show route next-hop</code> |
| Display routes not associated with any BGP community. | <code>show route no-community</code> |
| Display routes exiting the router through the specified interface. | <code>show route output</code> |
| Display routes learned by the specified protocol. | <code>show route protocol</code> |
| Display routes in a range of destination prefixes. | <code>show route range</code> |
| Display routes received by a particular routing protocol. | <code>show route receive-protocol</code> |
| Display entries in the next-hop resolution database. | <code>show route resolution</code> |
| Display routes learned from snooping. | <code>show route snooping</code> |
| Display routes learned from the specified source. | <code>show route source-gateway</code> |
| Display statistics about the routes in all routing tables. | <code>show route summary</code> |
| Display routes in a particular routing table. | <code>show route table</code> |
| Display high-level summary of routing table information. | <code>show route terse</code> |



NOTE: For information about how to configure protocol-independent features, see the *Junos Routing Protocols Configuration Guide* and the *Junos Policy Framework Configuration Guide*.

show as-path

| | |
|------------------------------------|---|
| Syntax | show as-path <brief detail> <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switches) | show as-path <brief detail> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | <p>Display the distribution of autonomous system (AS) paths that the local routing device is using (usually through the routing table). Use this command to debug problems for AS paths and to understand how AS paths have been manipulated through a policy (through the as-path-prepend action) or through aggregation.</p> <p>AS paths are stored in a hash table. A hash table is one method for fast lookup. Each entry in the table is called a bucket. Junos OS computes a hash value that indicates in which bucket the AS path is stored. The AS paths are dispersed among the hash buckets so that a manageable number of AS paths is stored in each bucket. Only unique AS paths are stored. Duplicate AS paths increase a reference count, but do not increase the number of AS paths stored in the hash table.</p> |
| Options | <p>none—Display basic information about AS paths that the local routing device is using (same as brief).</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> show as-path summary on page 484 |
| List of Sample Output | show as-path on page 480 show as-path detail on page 480 |
| Output Fields | Table 121 on page 478 lists the output fields for the show as-path command. Output fields are listed in the approximate order in which they appear. |

Table 121: show as-path Output Fields

| Field Name | Field Description | Level of Output |
|----------------|---------------------------|-----------------|
| Total AS paths | Total number of AS paths. | brief none |
| Bucket | Bucket number. | All levels |

Table 121: show as-path Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------|---|-----------------|
| Count | Number of AS path entries in this bucket. | All levels |
| AS path | <p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> • I—IGP. • E—EGP. • ?—Incomplete; typically, the AS path was aggregated. • Atomic—Route is an aggregate of several route prefixes. • Aggregator—Routing device has summarized a range of prefixes. | All levels |
| domain | Number of independent AS domains. The AS paths of an independent AS domain are not shared with the AS paths and AS path attributes of other domains, including the master routing instance domain. | detail |
| neighbor as | AS peer address. | detail |
| length | Length of the AS path. | detail |
| segments | Length of the AS segment descriptor. | detail |
| references | Path reference count. | detail |

Sample Output

show as-path

```

user@host> show as-path
Total AS paths: 30382
Bucket 0      Count: 36
I
14203 2914 174 31752 I
14203 2914 701 21512 I
14203 2914 1239 26632 I
14203 2914 1239 29704 I
14203 2914 4323 10248 I
14203 2914 4766 23560 I
14203 2914 6395 32776 I
14203 2914 7911 11272 I
14203 2914 12180 18440 I
14203 2914 17408 17416 I
14203 2914 701 702 24586 I
14203 2914 1239 4657 9226 I
14203 2914 1239 7132 16394 I
14203 2914 1299 8308 34826 I
14203 2914 3320 5603 28682 I
14203 2914 3491 1680 33802 I
14203 2914 3549 7908 27658 I
14203 2914 3549 20804 30730 I
14203 2914 7018 2687 9226 I
14203 2914 174 9318 9318 23564 I
14203 2914 701 3786 3786 23564 I
14203 2914 701 4761 4795 9228 I
14203 2914 1239 7132 5673 18444 I
14203 2914 3491 20485 24588 24588 I
14203 2914 5511 2200 1945 2060 I
14203 2914 7911 14325 14325 14348 I
14203 2914 701 4637 9230 9230 9230 I
14203 2914 6395 14 14 14 14 I
14203 2914 9299 6163 6163 6163 6163 9232 I
14203 2914 3356 3356 3356 3356 3356 11955 21522 I
14203 2914 9837 9837 9219 I Aggregator: 9219 202.27.91.253
14203 2914 174 30209 30222 30222 30222 ?
14203 2914 1299 5377 I (Atomic) Aggregator: 5377 193.219.192.22
14203 2914 4323 36097 I (Atomic) Aggregator: 36097 216.69.252.254
14203 2914 209 2516 17676 23813 I (Atomic) Aggregator: 23813 219.127.233.66
Bucket 1      Count: 28
14203 2914 35847 I
14203 2914 174 19465 I
14203 2914 174 35849 I
14203 2914 2828 32777 I
14203 2914 4323 14345 I
14203 2914 4323 29705 I
14203 2914 6395 32777 I

...

```

show as-path detail

```

user@host> show as-path detail
Total AS paths: 30410
Bucket 0      Count: 36
AS path: I
    domain 0, length 0, segments 0, references 54
AS path: 14203 2914 174 31752 I
    domain 1, neighbor as: 14203, length 4, segments 1, references 2

```

```

AS path: 14203 2914 701 21512 I
    domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 1239 26632 I
    domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 1239 29704 I
    domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 4323 10248 I
    domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 4766 23560 I
    domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 6395 32776 I
    domain 1, neighbor as: 14203, length 4, segments 1, references 3
AS path: 14203 2914 7911 11272 I
    domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 12180 18440 I
    domain 1, neighbor as: 14203, length 4, segments 1, references 3
AS path: 14203 2914 17408 17416 I
    domain 1, neighbor as: 14203, length 4, segments 1, references 3
AS path: 14203 2914 701 702 24586 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 3
AS path: 14203 2914 1239 4657 9226 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 7
AS path: 14203 2914 1239 7132 16394 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 1299 8308 34826 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 3320 5603 28682 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 3491 1680 33802 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 3549 7908 27658 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 3549 20804 30730 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 7018 2687 9226 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 3
AS path: 14203 2914 174 9318 9318 23564 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 701 3786 3786 23564 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 701 4761 4795 9228 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 14
AS path: 14203 2914 1239 7132 5673 18444 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 3491 20485 24588 24588 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 4
AS path: 14203 2914 5511 2200 1945 2060 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 7911 14325 14325 14348 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 701 4637 9230 9230 9230 I
    domain 1, neighbor as: 14203, length 7, segments 1, references 3
AS path: 14203 2914 6395 14 14 14 14 I
    domain 1, neighbor as: 14203, length 7, segments 1, references 10
...

```

show as-path domain

| | |
|------------------------------------|--|
| Syntax | show as-path domain <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switches) | show as-path domain |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. |
| Description | Display autonomous system (AS) path domain information. |
| Options | none —(Optional) Display AS path domain information for all routing instances. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | view |
| List of Sample Output | show as-path domain on page 483 |
| Output Fields | Table 122 on page 482 lists the output fields for the show as-path domain command. Output fields are listed in the approximate order in which they appear |

Table 122: show as-path domain Output Fields

| Field Name | Field Description |
|---------------------|---|
| Domain | Number of independent AS domains. The AS paths of an independent AS domain are not shared with the AS paths and AS path attributes of other domains, including the master routing instance domain. |
| Primary | Primary AS number. |
| References | Path reference count. |
| Number Paths | Number of known AS paths. |
| Flags | Information about the AS path: <ul style="list-style-type: none"> • ASLoop—Path contains an AS loop. • Atomic—Path includes the ATOMIC_AGGREGATE path attribute. • Local—Path was created by local aggregation. • Master—Path was created by the master routing instance. |
| Local AS | AS number of the local routing device. |
| Loops | How many times this AS number can appear in an AS path. |

Sample Output

```
show as-path domain    user@host> show as-path domain
Domain: 1               Primary: 10458
References:             3 Paths:      30383
Flags: Master
Local AS: 10458  Loops: 1
```

show as-path summary

| | |
|------------------------------------|--|
| Syntax | show as-path summary <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switches) | show as-path summary |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. |
| Description | Display autonomous system (AS) path summary information. AS paths are stored in a hash table. A hash table is one method for fast lookup. Each entry in the table is called a bucket. Junos OS computes a hash value that indicates in which bucket the AS path is stored. The AS paths are dispersed among the hash buckets so that a manageable number of AS paths is stored in each bucket. Only unique AS paths are stored. Duplicate AS paths increase a reference count, but do not increase the number of AS paths stored in the hash table. |
| Options | none —(Optional) Display AS path summary information for all routing instances. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> show as-path on page 478 |
| List of Sample Output | show as-path summary on page 485 |
| Output Fields | Table 123 on page 484 lists the output fields for the show as-path summary command. Output fields are listed in the approximate order in which they appear. |

Table 123: show as-path summary Output Fields

| Field Name | Field Description |
|---------------|---|
| AS Paths | Number of AS paths. |
| Buckets | Number of hash buckets in use. |
| Max | Maximum number of AS path entries per bucket. |
| Min | Minimum number of AS path entries per bucket. |
| Avg | Average number of AS path entries per bucket. |
| Std deviation | Standard deviation of AS path entries per bucket. |

Sample Output

`show as-path
summary`

```
user@host> show as-path summary
AS Paths Buckets    Max    Min    Avg    Std deviation
  30425     1024      95     12     29      6.481419
```

show route

| | |
|------------------------------------|---|
| Syntax | <code>show route</code> <code><all></code> <code><destination-prefix></code> <code><logical-system (all <i>logical-system-name</i>)></code> <code><private></code> |
| Syntax (EX Series Switches) | <code>show route</code> <code><all></code> <code><destination-prefix></code> <code><private></code> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Option private introduced in Junos OS Release 9.5. Option private introduced in Junos OS Release 9.5 for EX Series switches. |
| Description | Display the active entries in the routing tables. |
| Options | none —Display brief information about all active entries in the routing tables. all —(Optional) Display information about all routing tables, including private, or internal, routing tables. destination-prefix —(Optional) Display active entries for the specified address or range of addresses. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. private —(Optional) Display information only about all private, or internal, routing tables. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">• Example: Configuring RIP• Example: Configuring RIPng• Example: Configuring IS-IS• Examples: Configuring Internal BGP Peering• Examples: Configuring External BGP Peering• Examples: Configuring OSPF Routing Policy |
| List of Sample Output | show route on page 490 show route destination-prefix on page 490 show route extensive on page 490 |

Output Fields Table 124 on page 487 describes the output fields for the **show route** command. Output fields are listed in the approximate order in which they appear.

Table 124: show route Output Fields

| Field Name | Field Description |
|---------------------------------|--|
| <i>routing-table-name</i> | Name of the routing table (for example, inet.0). |
| <i>number destinations</i> | Number of destinations for which there are routes in the routing table. |
| <i>number routes</i> | <p>Number of routes in the routing table and total number of routes in the following states:</p> <ul style="list-style-type: none"> • active (routes that are active). • holddown (routes that are in the pending state before being declared inactive). A holddown route was once the active route and is no longer the active route. The route is in the holddown state because a protocol still has interest in the route, meaning that the interest bit is set. A protocol might have its interest bit set on the previously active route because the protocol is still advertising the route. The route will be deleted after all protocols withdraw their advertisement of the route and remove their interest bit. A persistent holddown state often means that the interested protocol is not releasing its interest bit properly. <p>However, if you have configured advertisement of multiple routes (with the add-path or advertise-inactive statement), the holddown bit is most likely set because BGP is advertising the route as an active route. In this case, you can ignore the holddown state because nothing is wrong.</p> <ul style="list-style-type: none"> • hidden (routes that are not used because of a routing policy). |
| <i>destination-prefix</i> | <p>Route destination (for example:10.0.0.1/24). Sometimes the route information is presented in another format, such as:</p> <ul style="list-style-type: none"> • MPLS-label (for example, 80001). • interface-name (for example, ge-1/0/2). • neighbor-address:control-word-status:encapsulation type:vc-id:source (Layer 2 circuit only. For example, 10.1.1.195:NoCtrlWord:1:1:Local/96): <ul style="list-style-type: none"> • neighbor-address—Address of the neighbor. • control-word-status—Whether the use of the control word has been negotiated for this virtual circuit: NoCtrlWord or CtrlWord. • encapsulation type—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport. • vc-id—Virtual circuit identifier. • source—Source of the advertisement: Local or Remote. |
| [protocol, preference] | <p>Protocol from which the route was learned and the preference value for the route.</p> <ul style="list-style-type: none"> • +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • -—A hyphen indicates the last active route. • *—An asterisk indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route. <p>In every routing metric except for the BGP LocalPref attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the LocalPref value in the Preference2 field. For example, if the LocalPref value for Route 1 is 100, the Preference2 value is -101. If the LocalPref value for Route 2 is 155, the Preference2 value is -156. Route 2 is preferred because it has a higher LocalPref value and a lower Preference2 value.</p> |

Table 124: show route Output Fields (*continued*)

| Field Name | Field Description |
|---|---|
| <i>weeks:days</i> <i>hours:minutes:seconds</i> | How long the route been known (for example, 2w4d 13:11:14 , or 2 weeks, 4 days, 13 hours, 11 minutes, and 14 seconds). |
| metric | Cost value of the indicated route. For routes within an AS, the cost is determined by the IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value. |
| localpref | Local preference value included in the route. |
| from | Interface from which the route was received. |
| AS path | <p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> • I—IGP. • E—EGP. • ?—Incomplete; typically, the AS path was aggregated. <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> • []—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device, or if AS path prepending is configured. • { }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • ()—Parentheses enclose a confederation. • ([])—Parentheses and brackets enclose a confederation set. <p>NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p> |
| validation-state | <p>(BGP-learned routes) Validation status of the route:</p> <ul style="list-style-type: none"> • Invalid—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • Unknown—Indicates that the prefix is not among the prefixes or prefix ranges in the database. • Valid—Indicates that the prefix and autonomous system pair are found in the database. |
| to | <p>Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.</p> <p>If the destination is Discard, traffic is dropped.</p> |

Table 124: show route Output Fields (*continued*)

| Field Name | Field Description |
|------------|---|
| via | <p>Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word Selected. This field can also contain the following information:</p> <ul style="list-style-type: none">• Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.• Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.• lsp-path-name—Name of the LSP used to reach the next hop.• label-action—MPLS label and operation occurring at the next hop. The operation can be pop (where a label is removed from the top of the stack), push (where another label is added to the label stack), or swap (where a label is replaced by another label). |

Sample Output

show route

```

user@host> show route
inet.0: 11 destinations, 12 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:65500:1:10.0.0.20/240
    *[MVPN/70] 19:53:41, metric2 1
    Indirect
1:65500:1:10.0.0.40/240
    *[BGP/170] 19:53:29, localpref 100, from 10.0.0.30
    AS path: I
    > to 10.0.24.4 via lt-0/3/0.24, label-switched-path toD
    [BGP/170] 19:53:26, localpref 100, from 10.0.0.33
    AS path: I
    > to 10.0.24.4 via lt-0/3/0.24, label-switched-path toD
1:65500:1:10.0.0.60/240
    *[BGP/170] 19:53:29, localpref 100, from 10.0.0.30
    AS path: I
    > to 10.0.28.8 via lt-0/3/0.28, label-switched-path toF
    [BGP/170] 19:53:25, localpref 100, from 10.0.0.33
    AS path: I
    > to 10.0.28.8 via lt-0/3/0.28, label-switched-path toF

```

show route destination-prefix

```

user@host> show route 172.16.0.0/12
inet.0: 10 destinations, 10 routes (9 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.0.0/12    *[Static/5] 2w4d 12:54:27
    > to 192.168.167.254 via fxp0.0

```

show route extensive

```

user@host> show route extensive
v1.mvpn.0: 5 destinations, 8 routes (5 active, 1 holddown, 0 hidden)
1:65500:1:10.0.0.40/240 (1 entry, 1 announced)
    *BGP    Preference: 170/-101
    PMSI: Flags 0x0: Label[0:0:0]: PIM-SM: Sender 10.0.0.40 Group 225.1.1.1

    Next hop type: Indirect
    Address: 0x92455b8
    Next-hop reference count: 2
    Source: 10.0.0.30
    Protocol next hop: 10.0.0.40
    Indirect next hop: 2 no-forward
    State: <Active Int Ext>
        Local AS: 65500 Peer AS: 65500
    Age: 3 Metric2: 1
    Task: BGP_65500.10.0.0.30+179
    Announcement bits (2): 0-PIM.v1 1-mvpn global task
    AS path: I (Originator) Cluster list: 10.0.0.30
    AS path: Originator ID: 10.0.0.40
    Communities: target:65520:100
    Import Accepted
    Localpref: 100
    Router ID: 10.0.0.30
    Primary Routing Table bgp.mvpn.0
    Indirect next hops: 1
        Protocol next hop: 10.0.0.40 Metric: 1

```

```
Indirect next hop: 2 no-forward
Indirect path forwarding next hops: 1
  Next hop type: Router
  Next hop: 10.0.24.4 via lt-0/3/0.24 weight 0x1
10.0.0.40/32 Originating RIB: inet.3
  Metric: 1                      Node path count: 1
  Forwarding nexthops: 1
    Nexthop: 10.0.24.4 via lt-0/3/0.24
```

show route active-path

| | |
|------------------------------------|---|
| Syntax | show route active-path <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switches) | show route active-path <brief detail extensive terse> |
| Release Information | Command introduced in Junos OS Release 8.0. Command introduced in Junos OS Release 9.0 for EX Series switches. |
| Description | Display all active routes for destinations. An active route is a route that is selected as the best path. Inactive routes are not displayed. |
| Options | none —Display all active routes. brief detail extensive terse —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief . logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | view |
| List of Sample Output | show route active-path on page 493 show route active-path brief on page 493 show route active-path detail on page 493 show route active-path extensive on page 494 show route active-path terse on page 496 |
| Output Fields | For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command. |

Sample Output

show route active-path user@host> **show route active-path**

```
inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.70.19/32    *[Direct/0] 21:33:52
                  > via lo0.0
10.255.71.50/32    *[IS-IS/15] 00:18:13, metric 10
                  > to 100.1.2.1 via so-2/1/3.0
100.1.2.0/24      *[Direct/0] 00:18:36
                  > via so-2/1/3.0
100.1.2.2/32      *[Local/0] 00:18:41
                  Local via so-2/1/3.0
192.168.64.0/21   *[Direct/0] 21:33:52
                  > via fxp0.0
192.168.70.19/32  *[Local/0] 21:33:52
                  Local via fxp0.0
```

show route active-path brief The output for the **show route active-path brief** command is identical to that for the **show route active-path** command. For sample output, see [show route active-path on page 493](#).

show route active-path detail user@host> **show route active-path detail**

```
inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)

10.255.70.19/32 (1 entry, 1 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 3
    Next hop: via lo0.0, selected
    State: <Active Int>
    Local AS: 200
    Age: 21:37:10
    Task: IF
    Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3
    AS path: I

10.255.71.50/32 (1 entry, 1 announced)
  *IS-IS Preference: 15
    Level: 1
    Next hop type: Router, Next hop index: 397
    Next-hop reference count: 4
    Next hop: 100.1.2.1 via so-2/1/3.0, selected
    State: <Active Int>
    Local AS: 200
    Age: 21:31 Metric: 10
    Task: IS-IS
    Announcement bits (4): 0-KRT 2-IS-IS 5-Resolve tree 2 6-Resolve
tree 3
    AS path: I

100.1.2.0/24 (1 entry, 1 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 3
```

```

Next hop: via so-2/1/3.0, selected
State: <Active Int>
Local AS: 200
Age: 21:54
Task: IF
Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3

AS path: I

100.1.2.2/32 (1 entry, 1 announced)
  *Local Preference: 0
    Next hop type: Local
    Next-hop reference count: 11
    Interface: so-2/1/3.0
    State: <Active NoReadvrt Int>
    Local AS: 200
    Age: 21:59
    Task: IF
    Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
    AS path: I

192.168.64.0/21 (1 entry, 1 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 3
    Next hop: via fxp0.0, selected
    State: <Active Int>
    Local AS: 200
    Age: 21:37:10
    Task: IF
    Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
    AS path: I

192.168.70.19/32 (1 entry, 1 announced)
  *Local Preference: 0
    Next hop type: Local
    Next-hop reference count: 11
    Interface: fxp0.0
    State: <Active NoReadvrt Int>
    Local AS: 200
    Age: 21:37:10
    Task: IF
    Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
    AS path: I

```

show route active-path extensive

```

user@host> show route active-path extensive

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
10.255.70.19/32 (1 entry, 1 announced)
TSI:
IS-IS level 1, LSP fragment 0
IS-IS level 2, LSP fragment 0
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 3
    Next hop: via lo0.0, selected
    State: <Active Int>
    Local AS: 200
    Age: 21:39:47
    Task: IF
    Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3

```



```

AS path: I

10.255.71.50/32 (1 entry, 1 announced)
TSI:
KRT in-kerne1 10.255.71.50/32 -> {100.1.2.1}
IS-IS level 2, LSP fragment 0
    *IS-IS Preference: 15
        Level: 1
        Next hop type: Router, Next hop index: 397
        Next-hop reference count: 4
        Next hop: 100.1.2.1 via so-2/1/3.0, selected
        State: <Active Int>
        Local AS: 200
        Age: 24:08 Metric: 10
        Task: IS-IS
        Announcement bits (4): 0-KRT 2-IS-IS 5-Resolve tree 2 6-Resolve
tree 3
AS path: I

100.1.2.0/24 (1 entry, 1 announced)
TSI:
IS-IS level 1, LSP fragment 0
IS-IS level 2, LSP fragment 0
    *Direct Preference: 0
        Next hop type: Interface
        Next-hop reference count: 3
        Next hop: via so-2/1/3.0, selected
        State: <Active Int>
        Local AS: 200
        Age: 24:31
        Task: IF
        Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3
AS path: I

100.1.2.2/32 (1 entry, 1 announced)
    *Local Preference: 0
        Next hop type: Local
        Next-hop reference count: 11
        Interface: so-2/1/3.0
        State: <Active NoReadvrt Int>
        Local AS: 200
        Age: 24:36
        Task: IF
        Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

192.168.64.0/21 (1 entry, 1 announced)
    *Direct Preference: 0
        Next hop type: Interface
        Next-hop reference count: 3
        Next hop: via fxp0.0, selected
        State: <Active Int>
        Local AS: 200
        Age: 21:39:47
        Task: IF
        Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

192.168.70.19/32 (1 entry, 1 announced)

```

```

*Local Preference: 0
  Next hop type: Local
  Next-hop reference count: 11
  Interface: fxp0.0
  State: <Active NoReadvrt Int>
  Local AS: 200
  Age: 21:39:47
  Task: IF
  Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
  AS path: I

```

show route active-path terse user@host> **show route active-path terse**

```

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

```

| A | Destination | P | Prf | Metric 1 | Metric 2 | Next hop | AS path |
|---|------------------|---|-----|----------|----------|-------------|---------|
| * | 10.255.70.19/32 | D | 0 | | | >lo0.0 | |
| * | 10.255.71.50/32 | I | 15 | 10 | | >100.1.2.1 | |
| * | 100.1.2.0/24 | D | 0 | | | >so-2/1/3.0 | |
| * | 100.1.2.2/32 | L | 0 | | | Local | |
| * | 192.168.64.0/21 | D | 0 | | | >fxp0.0 | |
| * | 192.168.70.19/32 | L | 0 | | | Local | |

show route advertising-protocol

| | |
|---------------------------------|--|
| Syntax | <code>show route advertising-protocol <i>protocol</i> <i>neighbor-address</i></code> <code><brief detail extensive terse></code> <code><logical-system (all <i>logical-system-name</i>)></code> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display the routing information as it has been prepared for advertisement to a particular neighbor of a particular dynamic routing protocol. |
| Options | <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><i>neighbor-address</i>—Address of the neighboring router to which the route entry is being transmitted.</p> <p><i>protocol</i>—Protocol transmitting the route:</p> <ul style="list-style-type: none"> • bgp—Border Gateway Protocol • dvmrp—Distance Vector Multicast Routing Protocol • msdp—Multicast Source Discovery Protocol • pim—Protocol Independent Multicast • rip—Routing Information Protocol • ripng—Routing Information Protocol next generation |
| Additional Information | Routes displayed are routes that the routing table has exported into the routing protocol and that have been filtered by the associated protocol's export routing policy statements. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • Example: Configuring the MED Attribute Directly |
| List of Sample Output | show route advertising-protocol bgp (Layer 3 VPN) on page 500 show route advertising-protocol bgp detail on page 500 show route advertising-protocol bgp detail (Layer 2 VPN) on page 500 show route advertising-protocol bgp detail (Layer 3 VPN) on page 500 show route advertising-protocol bgp extensive all (Next Hop Self with RIB-out IP Address) on page 501 |
| Output Fields | Table 125 on page 498 lists the output fields for the show route advertising-protocol command. Output fields are listed in the approximate order in which they appear. |

Table 125: show route advertising-protocol Output Fields

| Field Name | Field Description | Level of Output |
|--|---|-------------------------|
| <i>routing-table-name</i> | Name of the routing table—for example, inet.0. | All levels |
| <i>number destinations</i> | Number of destinations for which there are routes in the routing table. | All levels |
| <i>number routes</i> | Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> • active (routes that are active) • holddown (routes that are in the pending state before being declared inactive) • hidden (routes that are not used because of a routing policy) | All levels |
| Prefix | Destination prefix. | brief none |
| <i>destination-prefix (entry, announced)</i> | Destination prefix. The entry value is the number of routes for this destination, and the announced value is the number of routes being announced for this destination. | detail extensive |
| BGP group and type | BGP group name and type (Internal or External). | detail extensive |
| Route Distinguisher | Unique 64-bit prefix augmenting each IP subnet. | detail extensive |
| Advertised Label | Incoming label advertised by the LDP. When an IP packet enters a label-switched path (LSP), the ingress router examines the packet and assigns it a label based on its destination, placing the label in the packet's header. The label transforms the packet from one that is forwarded based on its IP routing information to one that is forwarded based on information associated with the label. | detail extensive |
| Label-Base, range | First label in a block of labels and label block size. A remote PE router uses this first label when sending traffic toward the advertising PE router. | detail extensive |
| VPN Label | Virtual private network (VPN) label. Packets are sent between CE and PE routers by advertising VPN labels. VPN labels transit over either an RSVP or an LDP LSP tunnel. | detail extensive |
| Nexthop | Next hop to the destination. An angle bracket (>) indicates that the route is the selected route. If the next-hop advertisement to the peer is Self , and the RIB-out next hop is a specific IP address, the RIB-out IP address is included in the extensive output. See show route advertising-protocol bgp extensive all (Next Hop Self with RIB-out IP Address) on page 501. | All levels |
| MED | Multiple exit discriminator value included in the route. | brief |
| Lclpref or Localpref | Local preference value included in the route. | All levels |

Table 125: show route advertising-protocol Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|----------------------------|--|------------------|
| AS path | <p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> • I—IGP. • E—EGP. • ?—Incomplete; typically, the AS path was aggregated. <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> • []—Brackets enclose the local AS number associated with the AS path if configured on the router, or if AS path prepending is configured. • { }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • ()—Parentheses enclose a confederation. • ([])—Parentheses and brackets enclose a confederation set. <p>NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p> | All levels |
| Communities | Community path attribute for the route. See the output field table for the show route detail command for all possible values for this field. | detail extensive |
| AIGP | Accumulated interior gateway protocol (AIGP) BGP attribute. | detail extensive |
| Attrset AS | Number, local preference, and path of the autonomous system (AS) that originated the route. These values are stored in the Attrset attribute at the originating router. | detail extensive |
| Layer2-info: encaps | Layer 2 encapsulation (for example, VPLS). | detail extensive |
| control flags | Control flags: none or Site Down . | detail extensive |
| mtu | Maximum transmission unit (MTU) of the Layer 2 circuit. | detail extensive |

Sample Output

**show route
advertising-protocol
bgp (Layer 3 VPN)**

```
user@host> show route advertising-protocol bgp 10.255.14.171
VPN-A.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
10.255.14.172/32 Self              1      100 I
VPN-B.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
10.255.14.181/32 Self              2      100 I
```

**show route
advertising-protocol
bgp detail**

```
user@host> show route advertising-protocol bgp 111.222.1.3 detail
bgp20.inet.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
111.222.1.11/32 (1 entry, 1 announced)
  BGP group pe-pe type Internal
  Route Distinguisher: 111.255.14.11:69
  Advertised Label: 100000
  next hop: Self
  Localpref: 100
  AS path: 2 I
  Communities: target:69:20
  AIGP 210
111.8.0.0/16 (1 entry, 1 announced)
  BGP group pe-pe type Internal
  Route Distinguisher: 111.255.14.11:69
  Advertised Label: 100000
  Next hop: Self
  Localpref: 100
  AS path: 2 I
  Communities: target:69:20
  AIGP 210
```

**show route
advertising-protocol
bgp detail (Layer 2
VPN)**

```
user@host> show route advertising-protocol bgp 192.168.24.1 detail
vpn-a.12vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
192.168.16.1:1:1:1/96 (1 entry, 1 announced)
  BGP group int type Internal
  Route Distinguisher: 192.168.16.1:1
  Label-base : 32768, range : 3
  Nexthop: Self
  Localpref: 100
  AS path: I
  Communities: target:65412:100
  AIGP 210
  Layer2-info: encaps:VLAN, control flags:, mtu:
```

**show route
advertising-protocol**

```
user@host> show route advertising-protocol bgp 10.255.14.176 detail
vpna.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
* 10.49.0.0/30 (1 entry, 1 announced)
```

bgp detail (Layer 3 VPN)

```

BGP group ibgp type Internal
Route Distinguisher: 10.255.14.174:2
VPN Label: 101264
Nexthop: Self
Localpref: 100
AS path: I
Communities: target:200:100
AIGP 210
AttrSet AS: 100
    Localpref: 100
    AS path: I
...

```

**show route
advertising-protocol
bgp extensive all (Next
Hop Self with RIB-out
IP Address)**

```

user@host> show route advertising-protocol bgp 200.0.0.2 170.0.1.0/24 extensive all
inet.0: 13 destinations, 19 routes (13 active, 0 holddown, 6 hidden)
  170.0.1.0/24 (2 entries, 1 announced)
    BGP group eBGP-INTEROP type External
      Nexthop: Self (rib-out 10.100.3.2)
      AS path: [4713] 200 I
...

```

show route all

| | |
|------------------------------------|---|
| Syntax | <code>show route all</code> <code><logical-system (all <i>logical-system-name</i>)></code> |
| Syntax (EX Series Switches) | <code>show route all</code> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. |
| Description | Display information about all routes in all routing tables, including private, or internal, tables. |
| Options | none —Display information about all routes in all routing tables, including private, or internal, tables. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | view |
| List of Sample Output | show route all on page 503 |
| Output Fields | In Junos OS Release 9.5 and later, only the output fields for the show route all command display all routing tables, including private, or hidden, routing tables. The output field table of the show route command does not display entries for private, or hidden, routing tables in Junos OS Release 9.5 and later. |

Sample Output

show route all

The following example displays a snippet of output from the **show route** command and then displays the same snippet of output from the **show route all** command:

```
user@host> show route
mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
0          *[MPLS/0] 2d 02:24:39, metric 1
            Receive
1          *[MPLS/0] 2d 02:24:39, metric 1
            Receive
2          *[MPLS/0] 2d 02:24:39, metric 1
            Receive
800017     *[VPLS/7] 1d 14:00:16
            > via vt-3/2/0.32769, Pop
800018     *[VPLS/7] 1d 14:00:26
            > via vt-3/2/0.32772, Pop

user@host> show route all
mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
0          *[MPLS/0] 2d 02:19:12, metric 1
            Receive
1          *[MPLS/0] 2d 02:19:12, metric 1
            Receive
2          *[MPLS/0] 2d 02:19:12, metric 1
            Receive
800017     *[VPLS/7] 1d 13:54:49
            > via vt-3/2/0.32769, Pop
800018     *[VPLS/7] 1d 13:54:59
            > via vt-3/2/0.32772, Pop
vt-3/2/0.32769 [VPLS/7] 1d 13:54:49
                Unusable
vt-3/2/0.32772 [VPLS/7] 1d 13:54:59
                Unusable
```

show route aspath-regex

| | |
|------------------------------------|---|
| Syntax | <code>show route aspath-regex <i>regular-expression</i></code> <code><logical-system (all <i>logical-system-name</i>)></code> |
| Syntax (EX Series Switches) | <code>show route aspath-regex <i>regular-expression</i></code> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. |
| Description | Display the entries in the routing table that match the specified autonomous system (AS) path regular expression. |
| Options | <p><i>regular-expression</i>—Regular expression that matches an entire AS path.</p> <p><i>logical-system (all logical-system-name)</i>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Additional Information | <p>You can specify a regular expression as:</p> <ul style="list-style-type: none"> • An individual AS number • A period wildcard used in place of an AS number • An AS path regular expression that is enclosed in parentheses <p>You also can include the operators described in the table of AS path regular expression operators in the <i>Junos Policy Framework Configuration Guide</i>. The following list summarizes these operators:</p> <ul style="list-style-type: none"> • <i>{m,n}</i>—At least <i>m</i> and at most <i>n</i> repetitions of the AS path term. • <i>{m}</i>—Exactly <i>m</i> repetitions of the AS path term. • <i>{m,}</i>—<i>m</i> or more repetitions of the AS path term. • <i>*</i>—Zero or more repetitions of an AS path term. • <i>+</i>—One or more repetitions of an AS path term. • <i>?</i>—Zero or one repetition of an AS path term. • <i>aspath_term aspath_term</i>—Match one of the two AS path terms. <p>When you specify more than one AS number or path term, or when you include an operator in the regular expression, enclose the entire regular expression in quotation marks. For example, to match any path that contains AS number 234, specify the following command:</p> <pre>show route aspath-regex ". * 234 . *"</pre> |
| Required Privilege Level | view |

List of Sample Output [show route aspath-regex \(Matching a Specific AS Number\) on page 505](#)
[show route aspath-regex \(Matching Any Path with Two AS Numbers\) on page 505](#)

Output Fields For information about output fields, see the output field table for the [show route](#) command.

Sample Output

**show route
aspath-regex
(Matching a Specific
AS Number)**

```
user@host> show route aspath-regex 65477
inet.0: 46411 destinations, 46411 routes (46409 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

111.222.1.0/25      *[BGP/170] 00:08:48, localpref 100, from 111.222.2.24
                   AS Path: [65477] ({65488 65535}) IGP
                   to 111.222.18.225 via fpa0.0(111.222.18.233)
111.222.1.128/25   *[IS-IS/15] 09:15:37, metric 37, tag 1
                   to 111.222.18.225 via fpa0.0(111.222.18.233)
                   [BGP/170] 00:08:48, localpref 100, from 111.222.2.24
                   AS Path: [65477] ({65488 65535}) IGP
                   to 111.222.18.225 via fpa0.0(111.222.18.233)
...
```

**show route
aspath-regex
(Matching Any Path
with Two AS Numbers)**

```
user@host> show route aspath-regex ?.* 234 3561.*?
inet.0: 46351 destinations, 46351 routes (46349 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

9.20.0.0/17        *[BGP/170] 01:35:00, localpref 100, from 131.103.20.49
                   AS Path: [666] 234 3561 2685 2686 Incomplete
                   to 192.156.169.1 via 192.156.169.14(so-0/0/0)
12.10.231.0/24     *[BGP/170] 01:35:00, localpref 100, from 131.103.20.49
                   AS Path: [666] 234 3561 5696 7369 IGP
                   to 192.156.169.1 via 192.156.169.14(so-0/0/0)
24.64.32.0/19      *[BGP/170] 01:34:59, localpref 100, from 131.103.20.49
                   AS Path: [666] 234 3561 6327 IGP
                   to 192.156.169.1 via 192.156.169.14(so-0/0/0)
...
```

show route best

| | |
|------------------------------------|--|
| Syntax | <code>show route best <i>destination-prefix</i></code> <code><brief detail extensive terse></code> <code><logical-system (all <i>logical-system-name</i>)></code> |
| Syntax (EX Series Switches) | <code>show route best <i>destination-prefix</i></code> <code><brief detail extensive terse></code> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. |
| Description | Display the route in the routing table that is the best route to the specified address or range of addresses. The best route is the longest matching route. |
| Options | brief detail extensive terse —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief . <i>destination-prefix</i> —Address or range of addresses. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | view |
| List of Sample Output | show route best on page 507 show route best detail on page 507 show route best extensive on page 508 show route best terse on page 508 |
| Output Fields | For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command. |

Sample Output

show route best

```

user@host> show route best 10.255.70.103
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.255.70.103/32    *[OSPF/10] 1d 13:19:20, metric 2
                  > to 10.31.1.6 via ge-3/1/0.0
                  via so-0/3/0.0

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.255.70.103/32    *[RSVP/7] 1d 13:20:13, metric 2
                  > via so-0/3/0.0, label-switched-path green-r1-r3

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.0.0.0/8          *[Direct/0] 2d 01:43:34
                  > via fxp2.0
                  [Direct/0] 2d 01:43:34
                  > via fxp1.0

```

show route best detail

```

user@host> show route best 10.255.70.103 detail
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
10.255.70.103/32 (1 entry, 1 announced)
    *OSPF    Preference: 10
             Next-hop reference count: 9
             Next hop: 10.31.1.6 via ge-3/1/0.0, selected
             Next hop: via so-0/3/0.0
             State: <Active Int>
             Local AS:    69
             Age: 1d 13:20:06      Metric: 2
             Area: 0.0.0.0
             Task: OSPF
             Announcement bits (2): 0-KRT 3-Resolve tree 2
             AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
10.255.70.103/32 (1 entry, 1 announced)
    State: <FlashAll>
    *RSVP    Preference: 7
             Next-hop reference count: 5
             Next hop: via so-0/3/0.0 weight 0x1, selected
             Label-switched-path green-r1-r3
             Label operation: Push 100016
             State: <Active Int>
             Local AS:    69
             Age: 1d 13:20:59      Metric: 2
             Task: RSVP
             Announcement bits (1): 1-Resolve tree 2
             AS path: I

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
10.0.0.0/8 (2 entries, 0 announced)
    *Direct Preference: 0
             Next hop type: Interface

```

```

Next-hop reference count: 1
Next hop: via fxp2.0, selected
State: <Active Int>
Age: 2d 1:44:20
Task: IF
AS path: I
Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 1
Next hop: via fxp1.0, selected
State: <NotBest Int>
Inactive reason: No difference
Age: 2d 1:44:20
Task: IF
AS path: I

```

show route best extensive

The output for the **show route best extensive** command is identical to that for the **show route best detail** command. For sample output, see [show route best detail on page 507](#).

show route best terse

```

user@host> show route best 10.255.70.103 terse
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1   Metric 2   Next hop      AS path
* 10.255.70.103/32  0  10           2           >10.31.1.6
                                     so-0/3/0.0

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1   Metric 2   Next hop      AS path
* 10.255.70.103/32  R   7           2           >so-0/3/0.0

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1   Metric 2   Next hop      AS path
* 10.0.0.0/8        D   0           0           >fxp2.0
                   D   0           0           >fxp1.0

```

show route brief

| | |
|------------------------------------|--|
| Syntax | show route brief <destination-prefix> <logical-system (all logical-system-name)> |
| Syntax (EX Series Switches) | show route brief <destination-prefix> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. |
| Description | Display brief information about the active entries in the routing tables. |
| Options | <p>none—Display all active entries in the routing table.</p> <p>destination-prefix—(Optional) Display active entries for the specified address or range of addresses.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show route brief on page 510 |
| Output Fields | For information about output fields, see the Output Field table of the show route command. |

Sample Output

show route brief

```

user@host> show route brief
inet.0: 10 destinations, 10 routes (9 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0      *[Static/5] 1w5d 20:30:29
                Discard
10.255.245.51/32 *[Direct/0] 2w4d 13:11:14
                > via lo0.0
172.16.0.0/12  *[Static/5] 2w4d 13:11:14
                > to 192.168.167.254 via fxp0.0
192.168.0.0/18 *[Static/5] 1w5d 20:30:29
                > to 192.168.167.254 via fxp0.0
192.168.40.0/22 *[Static/5] 2w4d 13:11:14
                > to 192.168.167.254 via fxp0.0
192.168.64.0/18 *[Static/5] 2w4d 13:11:14
                > to 192.168.167.254 via fxp0.0
192.168.164.0/22 *[Direct/0] 2w4d 13:11:14
                > via fxp0.0
192.168.164.51/32 *[Local/0] 2w4d 13:11:14
                Local via fxp0.0
207.17.136.192/32 *[Static/5] 2w4d 13:11:14
                > to 192.168.167.254 via fxp0.0
green.inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
100.101.0.0/16  *[Direct/0] 1w5d 20:30:28
                > via fe-0/0/3.0
100.101.2.3/32  *[Local/0] 1w5d 20:30:28
                Local via fe-0/0/3.0
224.0.0.5/32   *[OSPF/10] 1w5d 20:30:29, metric 1
                MultiRecv

```


show route ccc

| | |
|---------------------------------|---|
| Syntax | show route ccc ccc <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display circuit cross-connect (CCC) entries in the Multiprotocol Link Switching (MPLS) routing table. |
| Options | <p>ccc—Name of an entry with a circuit cross-connect interface.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • show connections on page 813 |
| List of Sample Output | show route ccc extensive on page 511 |
| Output Fields | For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command. |

Sample Output

show route ccc extensive

```

user@host> show route ccc fe-0/1/0.600 extensive
mpls.0: 19 destinations, 19 routes (19 active, 0 holddown, 0 hidden)
fe-0/1/2.600 (1 entry, 1 announced)
TSI:
KRT in-kernel fe-0/1/2.600.0      /16 -> {0.0.0.0}
      *CCC      Preference: 7
                Next-hop reference count: 2
                Next hop: via so-0/0/3.0 weight 0x1, selected
                Label operation: Push 101424
                State: <Active Int>
                Local AS: 100
                Age: 28:13 Metric: 3
                Task: MPLS
                Announcement bits (1): 0-KRT
                AS path: I

```

show route community

| | |
|-----------------------------|--|
| Syntax | show route community <i>as-number:community-value</i> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switches) | show route community <i>as-number:community-value</i> <brief detail extensive terse> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. |
| Description | Display the route entries in each routing table that are members of a Border Gateway Protocol (BGP) community. |
| Options | <p><i>as-number:community-value</i>—One or more community identifiers. <i>as-number</i> is the AS number, and <i>community-value</i> is the community identifier. When you specify more than one community identifier, enclose the identifiers in double quotation marks. Community identifiers can include wildcards.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Additional Information | Specifying the community option displays all routes matching the community found within the routing table. The community option does not limit the output to only the routes being advertised to the neighbor after any egress routing policy. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">• show route detail on page 522 |
| List of Sample Output | show route community on page 513 |
| Output Fields | For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command. |

Sample Output

```
show route community user@host> show route community 234:80
inet.0: 46511 destinations, 46511 routes (46509 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

4.0.0.0/8          *[BGP/170] 03:33:07, localpref 100, from 131.103.20.49
                   AS Path: {666} 234 2548 1 IGP
                   to 192.156.169.1 via 192.156.169.14(so-0/0/0)
6.0.0.0/8          *[BGP/170] 03:33:07, localpref 100, from 131.103.20.49
                   AS Path: {666} 234 2548 568 721 Incomplete
                   to 192.156.169.1 via 192.156.169.14(so-0/0/0)
9.2.0.0/16         *[BGP/170] 03:33:06, localpref 100, from 131.103.20.49
                   AS Path: {666} 234 2548 1673 1675 1747 IGP
                   to 192.156.169.1 via 192.156.169.14(so-0/0/0)
```

show route community-name

| | |
|------------------------------------|---|
| Syntax | show route community-name <i>community-name</i> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switches) | show route community-name <i>community-name</i> <brief detail extensive terse> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. |
| Description | Display the route entries in each routing table that are members of a Border Gateway Protocol (BGP) community, specified by a community name. |
| Options | <i>community-name</i> —Name of the community. brief detail extensive terse—(Optional) Display the specified level of output. logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | view |
| List of Sample Output | show route community-name on page 515 |
| Output Fields | For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command. |

Sample Output

show route
community-name

```

user@host> show route community-name red-com
inet.0: 17 destinations, 17 routes (16 active, 0 holddown, 1 hidden)

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

instance1.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.245.212/32  *[BGP/170] 00:04:40, localpref 100, from 10.255.245.204
                  AS path: 300 I
                  > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix
20.20.20.20/32    *[BGP/170] 00:04:40, localpref 100, from 10.255.245.204
                  AS path: I
                  > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix
100.1.4.0/24     *[BGP/170] 00:04:40, localpref 100, from 10.255.245.204
                  AS path: I
                  > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

bgp.l3vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.245.204:10:10.255.245.212/32
                  *[BGP/170] 00:06:40, localpref 100, from 10.255.245.204
                  AS path: 300 I
                  > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix
10.255.245.204:10:20.20.20.20/32
                  *[BGP/170] 00:36:02, localpref 100, from 10.255.245.204
                  AS path: I
                  > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix
10.255.245.204:10:100.1.4.0/24
                  *[BGP/170] 00:36:02, localpref 100, from 10.255.245.204
                  AS path: I
                  > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

instance1.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

show route damping

| | | |
|---|---|--|
| Syntax | show route damping (decayed history suppressed) <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)> | |
| Syntax (EX Series Switch and QFX Series) | show route damping (decayed history suppressed) <brief detail extensive terse> | |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. | |
| Description | Display the BGP routes for which updates might have been reduced because of route flap damping. | |
| Options | <p>brief detail extensive terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>decayed—Display route damping entries that might no longer be valid, but are not suppressed.</p> <p>history—Display entries that have already been withdrawn, but have been logged.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>suppressed—Display entries that have been suppressed and are no longer being installed into the forwarding table or exported by routing protocols.</p> | |
| Required Privilege Level | view | |
| Related Documentation | <ul style="list-style-type: none"> • clear bgp damping on page 38 • show policy damping on page 74 | |
| List of Sample Output | show route damping decayed detail on page 520 show route damping history on page 520 show route damping history detail on page 520 | |
| Output Fields | Table 126 on page 516 lists the output fields for the show route damping command. Output fields are listed in the approximate order in which they appear. | |

Table 126: show route damping Output Fields

| Field Name | Field Description | Level of Output |
|---------------------------|---|-----------------|
| <i>routing-table-name</i> | Name of the routing table—for example, <i>inet.0</i> . | All levels |
| <i>destinations</i> | Number of destinations for which there are routes in the routing table. | All levels |

Table 126: show route damping Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--|---|-------------------------|
| <i>number routes</i> | Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> • active • holddown (routes that are in a pending state before being declared inactive) • hidden (the routes are not used because of a routing policy) | All levels |
| <i>destination-prefix (entry, announced)</i> | Destination prefix. The entry value is the number of routes for this destination, and the announced value is the number of routes being announced for this destination. | detail extensive |
| <i>[protocol, preference]</i> | Protocol from which the route was learned and the preference value for the route. <ul style="list-style-type: none"> • +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • -—A hyphen indicates the last active route. • *—An asterisk indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route. <p>In every routing metric except for the BGP LocalPref attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the LocalPref value in the Preference2 field. For example, if the LocalPref value for Route 1 is 100, the Preference2 value is -101. If the LocalPref value for Route 2 is 155, the Preference2 value is -156. Route 2 is preferred because it has a higher LocalPref value and a lower Preference2 value.</p> | All levels |
| Next-hop reference count | Number of references made to the next hop. | detail extensive |
| Source | IP address of the route source. | detail extensive |
| Next hop | Network layer address of the directly reachable neighboring system. | detail extensive |
| via | Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word Selected . | detail extensive |
| Protocol next hop | Network layer address of the remote routing device that advertised the prefix. This address is used to derive a forwarding next hop. | detail extensive |
| Indirect next hop | Index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops. | detail extensive |
| State | Flags for this route. For a description of possible values for this field, see the output field table for the show route detail command. | detail extensive |
| Local AS | AS number of the local routing device. | detail extensive |
| Peer AS | AS number of the peer routing device. | detail extensive |

Table 126: show route damping Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-------------------------|---|------------------|
| Age | How long the route has been known. | detail extensive |
| Metric | Metric for the route. | detail extensive |
| Task | Name of the protocol that has added the route. | detail extensive |
| Announcement bits | List of protocols that announce this route. n-Resolve inet indicates that the route is used for route resolution for next hops found in the routing table. n is an index used by Juniper Networks customer support only. | detail extensive |
| AS path | <p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> • I—IGP. • E—EGP. • ?—Incomplete; typically, the AS path was aggregated. <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> • []—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device or if AS path prepending is configured. • { }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • ()—Parentheses enclose a confederation. • ([])—Parentheses and brackets enclose a confederation set. <p>NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p> | All levels |
| to | Next hop to the destination. An angle bracket (>) indicates that the route is the selected route. | brief none |
| via | Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word Selected . | brief none |
| Communities | Community path attribute for the route. See the output field table for the show route detail command. | detail extensive |
| Localpref | Local preference value included in the route. | All levels |
| Router ID | BGP router ID as advertised by the neighbor in the open message. | detail extensive |
| Merit (last update/now) | Last updated and current figure-of-merit value. | detail extensive |

Table 126: show route damping Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------------|--|-------------------------|
| damping-parameters | Name that identifies the damping parameters used, which is defined in the damping statement at the [edit policy-options] hierarchy level. | detail extensive |
| Last update | Time of most recent change in path attributes. | detail extensive |
| First update | Time of first change in path attributes, which started the route damping process. | detail extensive |
| Flaps | Number of times the route has gone up or down or its path attributes have changed. | detail extensive |
| Suppressed | (suppressed keyword only) This route is currently suppressed. A suppressed route does not appear in the forwarding table and routing protocols do not export it. | All levels |
| Reusable in | (suppressed keyword only) Time when a suppressed route will again be available. | All levels |
| Preference will be | (suppressed keyword only) Preference value that will be applied to the route when it is again active. | All levels |

Sample Output

show route damping decayed detail

```

user@host> show route damping decayed detail
inet.0: 173319 destinations, 1533668 routes (172625 active, 4 holddown, 108083
hidden)
10.0.111.0/24 (7 entries, 1 announced)
  *BGP      Preference: 170/-101
            Next-hop reference count: 151973
            Source: 172.23.2.129
            Next hop: via so-1/2/0.0
            Next hop: via so-5/1/0.0, selected
            Next hop: via so-6/0/0.0
            Protocol next hop: 172.23.2.129
            Indirect next hop: 89a1a00 264185
            State: <Active Ext>
            Local AS: 65000 Peer AS: 65490
            Age: 3:28      Metric2: 0
            Task: BGP_65490.172.23.2.129+179
            Announcement bits (6): 0-KRT 1-RT 4-KRT 5-BGP.0.0.0.0+179

        6-Resolve tree 2 7-Resolve tree 3
        AS path: 65490 65520 65525 65525 65525 65525 I ()
        Communities: 65501:390 65501:2000 65501:3000 65504:701
        Localpref: 100
        Router ID: 172.23.2.129
        Merit (last update/now): 1934/1790
        damping-parameters: damping-high
        Last update:          00:03:28 First update:          00:06:40
        Flaps: 2

```

show route damping history

```

user@host> show route damping history
inet.0: 173320 destinations, 1533529 routes (172624 active, 6 holddown, 108122
hidden)
+ = Active Route, - = Last Active, * = Both

10.108.0.0/15      [BGP ] 2d 22:47:58, localpref 100
                   AS path: 65220 65501 65502 I
                   > to 192.168.60.85 via so-3/1/0.0

```

show route damping history detail

```

user@host> show route damping history detail
inet.0: 173319 destinations, 1533435 routes (172627 active, 2 holddown, 108105
hidden)
10.108.0.0/15 (3 entries, 1 announced)
  BGP      /-101
            Next-hop reference count: 69058
            Source: 192.168.60.85
            Next hop: 192.168.60.85 via so-3/1/0.0, selected
            State: <Hidden Ext>
            Inactive reason: Unusable path
            Local AS: 65000 Peer AS: 65220
            Age: 2d 22:48:10
            Task: BGP_65220.192.168.60.85+179
            AS path: 65220 65501 65502 I ()
            Communities: 65501:390 65501:2000 65501:3000 65504:3561
            Localpref: 100
            Router ID: 192.168.80.25
            Merit (last update/now): 1000/932
            damping-parameters: set-normal

```

Last update: 00:01:05 First update: 00:01:05
Flaps: 1

show route detail

| | |
|------------------------------------|---|
| Syntax | show route detail <destination-prefix> <logical-system (all logical-system-name)> |
| Syntax (EX Series Switches) | show route detail <destination-prefix> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. |
| Description | Display detailed information about the active entries in the routing tables. |
| Options | <p>none—Display all active entries in the routing table on all systems.</p> <p>destination-prefix—(Optional) Display active entries for the specified address or range of addresses.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show route detail on page 531 show route detail (with BGP Multipath) on page 536 |
| Output Fields | <p>Table 127 on page 522 describes the output fields for the show route detail command. Output fields are listed in the approximate order in which they appear.</p> |

Table 127: show route detail Output Fields

| Field Name | Field Description |
|----------------------------|---|
| <i>routing-table-name</i> | Name of the routing table (for example, inet.0). |
| <i>number destinations</i> | Number of destinations for which there are routes in the routing table. |
| <i>number routes</i> | Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> active (routes that are active) holddown (routes that are in the pending state before being declared inactive) hidden (routes that are not used because of a routing policy) |

Table 127: show route detail Output Fields (*continued*)

| Field Name | Field Description |
|--|---|
| <i>route-destination</i> (entry, announced) | <p>Route destination (for example:10.0.0.1/24). The entry value is the number of routes for this destination, and the announced value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:</p> <ul style="list-style-type: none"> • MPLS-label (for example, 80001). • interface-name (for example, ge-1/0/2). • neighbor-address:control-word-status:encapsulation type:vc-id:source (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96). <ul style="list-style-type: none"> • neighbor-address—Address of the neighbor. • control-word-status—Whether the use of the control word has been negotiated for this virtual circuit: NoCtrlWord or CtrlWord. • encapsulation type—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport. • vc-id—Virtual circuit identifier. • source—Source of the advertisement: Local or Remote. |
| label stacking | <p>(Next-to-the-last-hop routing device for MPLS only) Depth of the MPLS label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels.</p> <ul style="list-style-type: none"> • S=0 route indicates that a packet with an incoming label stack depth of 2 or more exits this routing device with one fewer label (the label-popping operation is performed). • If there is no S= information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed). |
| [<i>protocol, preference</i>] | <p>Protocol from which the route was learned and the preference value for the route.</p> <ul style="list-style-type: none"> • +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • - —A hyphen indicates the last active route. • *—An asterisk indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route. <p>In every routing metric except for the BGP LocalPref attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the LocalPref value in the Preference2 field. For example, if the LocalPref value for Route 1 is 100, the Preference2 value is -101. If the LocalPref value for Route 2 is 155, the Preference2 value is -156. Route 2 is preferred because it has a higher LocalPref value and a lower Preference2 value.</p> |
| Level | <p>(IS-IS only). In IS-IS, a single AS can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.</p> |
| Route Distinguisher | IP subnet augmented with a 64-bit prefix. |
| Next-hop type | Type of next hop. For a description of possible values for this field, see Table 128 on page 526 . |

Table 127: show route detail Output Fields (*continued*)

| Field Name | Field Description |
|--|---|
| Next-hop reference count | Number of references made to the next hop. |
| Flood nexthop branches exceed maximum message | Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel. |
| Source | IP address of the route source. |
| Next hop | Network layer address of the directly reachable neighboring system. |
| via | <p>Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word Selected. This field can also contain the following information:</p> <ul style="list-style-type: none"> • Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing. |
| Label-switched-path lsp-path-name | Name of the LSP used to reach the next hop. |
| Label operation | MPLS label and operation occurring at this routing device. The operation can be pop (where a label is removed from the top of the stack), push (where another label is added to the label stack), or swap (where a label is replaced by another label). |
| Interface | (Local only) Local interface name. |
| Protocol next hop | Network layer address of the remote routing device that advertised the prefix. This address is used to derive a forwarding next hop. |
| Indirect next hop | Index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops. |
| State | State of the route (a route can be in more than one state). See Table 129 on page 528 . |
| Local AS | AS number of the local routing device. |
| Age | How long the route has been known. |
| AIGP | Accumulated interior gateway protocol (AIGP) BGP attribute. |
| Metricn | Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value. |

Table 127: show route detail Output Fields (*continued*)

| Field Name | Field Description |
|--------------------------------|---|
| MED-plus-IGP | Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added. |
| TTL-Action | For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signaled and LDP-signaled LSPs or for specific VRF routing instances. For sample output, see show route table . |
| Task | Name of the protocol that has added the route. |
| Announcement bits | List of protocols that announce this route. n-Resolve inet indicates that the route is used for route resolution for next hops found in the routing table. n is an index used by Juniper Networks customer support only. |
| AS path | <p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> • I—IGP. • E—EGP. • Recorded—The AS path is recorded by the sample process (sampled). • ?—Incomplete; typically, the AS path was aggregated. <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> • []—Brackets enclose the number that precedes the AS path. This number represents the number of ASs present in the AS path, when calculated as defined in RFC 4271. This value is used in the AS-path merge process, as defined in RFC 4893. • []—If more than one AS number is configured on the routing device, or if AS path prepending is configured, brackets enclose the local AS number associated with the AS path. • { }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • ()—Parentheses enclose a confederation. • ([])—Parentheses and brackets enclose a confederation set. <p>NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p> |
| VC Label | MPLS label assigned to the Layer 2 circuit virtual connection. |
| MTU | Maximum transmission unit (MTU) of the Layer 2 circuit. |
| VLAN ID | VLAN identifier of the Layer 2 circuit. |
| Prefixes bound to route | Forwarding equivalent class (FEC) bound to this route. Applicable only to routes installed by LDP. |
| Communities | Community path attribute for the route. See Table 130 on page 530 for all possible values for this field. |
| Layer2-info: encaps | Layer 2 encapsulation (for example, VPLS). |

Table 127: show route detail Output Fields (*continued*)

| Field Name | Field Description |
|----------------------------------|--|
| control flags | Control flags: none or Site Down . |
| mtu | Maximum transmission unit (MTU) information. |
| Label-Base, range | First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device. |
| status vector | Layer 2 VPN and VPLS network layer reachability information (NLRI). |
| Accepted Multipath | Current active path when BGP multipath is configured. |
| Accepted MultipathContrib | Path currently contributing to BGP multipath. |
| Localpref | Local preference value included in the route. |
| Router ID | BGP router ID as advertised by the neighbor in the open message. |
| Primary Routing Table | In a routing table group, the name of the primary routing table in which the route resides. |
| Secondary Tables | In a routing table group, the name of one or more secondary tables in which the route resides. |

Table 128 on page 526 describes all possible values for the **Next-hop Types** output field.

Table 128: Next-hop Types Output Field Values

| Next-Hop Type | Description |
|--------------------------|--|
| Broadcast (bcast) | Broadcast next hop. |
| Deny | Deny next hop. |
| Discard | Discard next hop. |
| Flood | Flood next hop. Consists of components called branches, up to a maximum of 32 branches. Each flood next-hop branch sends a copy of the traffic to the forwarding interface. Used by point-to-multipoint RSVP, point-to-multipoint LDP, point-to-multipoint CCC, and multicast. |
| Hold | Next hop is waiting to be resolved into a unicast or multicast type. |
| Indexed (idxd) | Indexed next hop. |

Table 128: Next-hop Types Output Field Values (*continued*)

| Next-Hop Type | Description |
|---------------------------------|--|
| Indirect (indr) | Used with applications that have a protocol next hop address that is remote. You are likely to see this next-hop type for internal BGP (IBGP) routes when the BGP next hop is a BGP neighbor that is not directly connected. |
| Interface | Used for a network address assigned to an interface. Unlike the router next hop, the interface next hop does not reference any specific node on the network. |
| Local (locl) | Local address on an interface. This next-hop type causes packets with this destination address to be received locally. |
| Multicast (mcst) | Wire multicast next hop (limited to the LAN). |
| Multicast discard (mdsc) | Multicast discard. |
| Multicast group (mgrp) | Multicast group member. |
| Receive (recv) | Receive. |
| Reject (rjct) | Discard. An ICMP unreachable message was sent. |
| Resolve (rslv) | Resolving next hop. |
| Routed multicast (mcrt) | Regular multicast next hop. |
| Router | <p>A specific node or set of nodes to which the routing device forwards packets that match the route prefix.</p> <p>To qualify as next-hop type router, the route must meet the following criteria:</p> <ul style="list-style-type: none"> • Must not be a direct or local subnet for the routing device. • Must have a next hop that is directly connected to the routing device. |
| Table | Routing table next hop. |
| Unicast (ucst) | Unicast. |
| Unilist (ulst) | List of unicast next hops. A packet sent to this next hop goes to any next hop in the list. |

Table 129 on page 528 describes all possible values for the **State** output field. A route can be in more than one state (for example, **<Active NoReadvrt Int Ext>**).

Table 129: State Output Field Values

| Value | Description |
|--|--|
| Accounting | Route needs accounting. |
| Active | Route is active. |
| Always Compare MED | Path with a lower multiple exit discriminator (MED) is available. |
| AS path | Shorter AS path is available. |
| Cisco Non-deterministic MED selection | Cisco nondeterministic MED is enabled, and a path with a lower MED is available. |
| Clone | Route is a clone. |
| Cluster list length | Length of cluster list sent by the route reflector. |
| Delete | Route has been deleted. |
| Ex | Exterior route. |
| Ext | BGP route received from an external BGP neighbor. |
| FlashAll | Forces all protocols to be notified of a change to any route, active or inactive, for a prefix. When not set, protocols are informed of a prefix only when the active route changes. |
| Hidden | Route not used because of routing policy. |
| IfCheck | Route needs forwarding RPF check. |
| IGP metric | Path through next hop with lower IGP metric is available. |
| Inactive reason | Flags for this route, which was not selected as best for a particular destination. |
| Initial | Route being added. |
| Int | Interior route. |
| Int Ext | BGP route received from an internal BGP peer or a BGP confederation peer. |
| Interior > Exterior > Exterior via Interior | Direct, static, IGP, or EBGp path is available. |
| Local Preference | Path with a higher local preference value is available. |
| Martian | Route is a martian (ignored because it is obviously invalid). |

Table 129: State Output Field Values (*continued*)

| Value | Description |
|---------------------------------------|---|
| MartianOK | Route exempt from martian filtering. |
| Next hop address | Path with lower metric next hop is available. |
| No difference | Path from neighbor with lower IP address is available. |
| NoReadvrt | Route not to be advertised. |
| NotBest | Route not chosen because it does not have the lowest MED. |
| Not Best in its group | Incoming BGP AS is not the best of a group (only one AS can be the best). |
| NotInstall | Route not to be installed in the forwarding table. |
| Number of gateways | Path with a greater number of next hops is available. |
| Origin | Path with a lower origin code is available. |
| Pending | Route pending because of a hold-down configured on another route. |
| Release | Route scheduled for release. |
| RIB preference | Route from a higher-numbered routing table is available. |
| Route Distinguisher | 64-bit prefix added to IP subnets to make them unique. |
| Route Metric or MED comparison | Route with a lower metric or MED is available. |
| Route Preference | Route with lower preference value is available |
| Router ID | Path through a neighbor with lower ID is available. |
| Secondary | Route not a primary route. |
| Unusable path | Path is not usable because of one of the following conditions: <ul style="list-style-type: none"> • The route is damped. • The route is rejected by an import policy. • The route is unresolved. |
| Update source | Last tiebreaker is the lowest IP address value. |

Table 130 on page 530 describes the possible values for the **Communities** output field.

Table 130: Communities Output Field Values

| Value | Description |
|---|---|
| <i>area-number</i> | 4 bytes, encoding a 32-bit area number. For AS-external routes, the value is 0 . A nonzero value identifies the route as internal to the OSPF domain, and as within the identified area. Area numbers are relative to a particular OSPF domain. |
| <i>bandwidth: local AS number:link-bandwidth-number</i> | Link-bandwidth community value used for unequal-cost load balancing. When BGP has several candidate paths available for multipath purposes, it does not perform unequal-cost load balancing according to the link-bandwidth community unless all candidate paths have this attribute. |
| <i>domain-id</i> | Unique configurable number that identifies the OSPF domain. |
| <i>domain-id-vendor</i> | Unique configurable number that further identifies the OSPF domain. |
| <i>link-bandwidth-number</i> | Link-bandwidth number: from 0 through 4,294,967,295 (bytes per second). |
| <i>local AS number</i> | Local AS number: from 1 through 65,535 . |
| <i>options</i> | 1 byte. Currently this is only used if the route type is 5 or 7 . Setting the least significant bit in the field indicates that the route carries a type 2 metric. |
| <i>origin</i> | (Used with VPNs) Identifies where the route came from. |
| <i>ospf-route-type</i> | 1 byte, encoded as 1 or 2 for intra-area routes (depending on whether the route came from a type 1 or a type 2 LSA); 3 for summary routes; 5 for external routes (area number must be 0); 7 for NSSA routes; or 129 for sham link endpoint addresses. |
| <i>route-type-vendor</i> | Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute 0x8000 . The format is area-number:ospf-route-type:options . |
| <i>rte-type</i> | Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute 0x0306 . The format is area-number:ospf-route-type:options . |
| <i>target</i> | Defines which VPN the route participates in; target has the format 32-bit IP address:16-bit number . For example, 10.19.0.0:100. |
| <i>unknown IANA</i> | Incoming IANA codes with a value between 0x1 and 0x7fff . This code of the BGP extended community attribute is accepted, but it is not recognized. |
| <i>unknown OSPF vendor community</i> | Incoming IANA codes with a value above 0x8000 . This code of the BGP extended community attribute is accepted, but it is not recognized. |

Sample Output

show route detail

user@host> show route detail

```
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 29
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 1:31:43
    Task: RT
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I

10.31.1.0/30 (2 entries, 1 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 2
    Next hop: via so-0/3/0.0, selected
    State: <Active Int>
    Local AS: 69
    Age: 1:30:17
    Task: IF
    Announcement bits (1): 3-Resolve tree 2
    AS path: I
  OSPF Preference: 10
    Next-hop reference count: 1
    Next hop: via so-0/3/0.0, selected
    State: <Int>
    Inactive reason: Route Preference
    Local AS: 69
    Age: 1:30:17 Metric: 1
    Area: 0.0.0.0
    Task: OSPF
    AS path: I

10.31.1.1/32 (1 entry, 1 announced)
  *Local Preference: 0
    Next hop type: Local
    Next-hop reference count: 7
    Interface: so-0/3/0.0
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:30:20
    Task: IF
    Announcement bits (1): 3-Resolve tree 2
    AS path: I

...

10.31.2.0/30 (1 entry, 1 announced)
  *OSPF Preference: 10
    Next-hop reference count: 9
    Next hop: via so-0/3/0.0
    Next hop: 10.31.1.6 via ge-3/1/0.0, selected
    State: <Active Int>
    Local AS: 69
    Age: 1:29:56 Metric: 2
```

```

Area: 0.0.0.0
Task: OSPF
Announcement bits (2): 0-KRT 3-Resolve tree 2
AS path: I

...

224.0.0.2/32 (1 entry, 1 announced)
  *PIM Preference: 0
      Next-hop reference count: 18
      State: <Active NoReadvrt Int>
      Local AS: 69
      Age: 1:31:45
      Task: PIM Recv
      Announcement bits (2): 0-KRT 3-Resolve tree 2
      AS path: I

...

224.0.0.22/32 (1 entry, 1 announced)
  *IGMP Preference: 0
      Next-hop reference count: 18
      State: <Active NoReadvrt Int>
      Local AS: 69
      Age: 1:31:43
      Task: IGMP
      Announcement bits (2): 0-KRT 3-Resolve tree 2
      AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.255.70.103/32 (1 entry, 1 announced)
  State: <FlashAll>
  *RSVP Preference: 7
      Next-hop reference count: 6
      Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
      Label-switched-path green-r1-r3
      Label operation: Push 100096
      State: <Active Int>
      Local AS: 69
      Age: 1:25:49 Metric: 2
      Task: RSVP
      Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
      AS path: I

10.255.71.238/32 (1 entry, 1 announced)
  State: <FlashAll>
  *RSVP Preference: 7
      Next-hop reference count: 6
      Next hop: via so-0/3/0.0 weight 0x1, selected
      Label-switched-path green-r1-r2
      State: <Active Int>
      Local AS: 69
      Age: 1:25:49 Metric: 1
      Task: RSVP
      Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
      AS path: I

private__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

47.0005.80ff.f800.0000.0108.0001.0102.5507.1052/152 (1 entry, 0 announced)

```
*Direct Preference: 0
  Next hop type: Interface
  Next-hop reference count: 1
  Next hop: via lo0.0, selected
  State: <Active Int>
  Local AS: 69
  Age: 1:31:44
  Task: IF
  AS path: I
```

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

0 (1 entry, 1 announced)

```
*MPLS Preference: 0
  Next hop type: Receive
  Next-hop reference count: 6
  State: <Active Int>
  Local AS: 69
  Age: 1:31:45 Metric: 1
  Task: MPLS
  Announcement bits (1): 0-KRT
  AS path: I
```

...

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

299776 (1 entry, 1 announced)

TSI:

KRT in-kernel 299776 /52 -> {Flood}

```
*RSVP Preference: 7
  Next hop type: Flood
  Next-hop reference count: 130
  Flood nexthop branches exceed maximum
  Address: 0x8ea65d0
```

...

800010 (1 entry, 1 announced)

```
*VPLS Preference: 7
  Next-hop reference count: 2
  Next hop: via vt-3/2/0.32769, selected
  Label operation: Pop
  State: <Active Int>
  Age: 1:29:30
  Task: Common L2 VC
  Announcement bits (1): 0-KRT
  AS path: I
```

vt-3/2/0.32769 (1 entry, 1 announced)

```
*VPLS Preference: 7
  Next-hop reference count: 2
  Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
  Label-switched-path green-r1-r3
  Label operation: Push 800012, Push 100096(top)
  Protocol next hop: 10.255.70.103
  Push 800012
  Indirect next hop: 87272e4 1048574
  State: <Active Int>
  Age: 1:29:30 Metric2: 2
  Task: Common L2 VC
  Announcement bits (2): 0-KRT 1-Common L2 VC
```

```

AS path: I
Communities: target:11111:1 Layer2-info: encaps:VPLS,
control flags:, mtu: 0

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

abcd::10:255:71:52/128 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
    State: <Active Int>
    Local AS: 69
    Age: 1:31:44
    Task: IF
    AS path: I

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:31:44
    Task: IF
    AS path: I

ff02::2/128 (1 entry, 1 announced)
  *PIM Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:31:45
    Task: PIM Recv6
    Announcement bits (1): 0-KRT
    AS path: I

ff02::d/128 (1 entry, 1 announced)
  *PIM Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:31:45
    Task: PIM Recv6
    Announcement bits (1): 0-KRT
    AS path: I

ff02::16/128 (1 entry, 1 announced)
  *MLD Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:31:43
    Task: MLD
    Announcement bits (1): 0-KRT
    AS path: I

private.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)

```



```

*Direct Preference: 0
  Next hop type: Interface
  Next-hop reference count: 1
  Next hop: via lo0.16385, selected
  State: <Active NoReadvrt Int>
  Age: 1:31:44
  Task: IF
  AS path: I

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

10.255.70.103:1:3:1/96 (1 entry, 1 announced)
  *BGP Preference: 170/-101
    Route Distinguisher: 10.255.70.103:1
    Next-hop reference count: 7
    Source: 10.255.70.103
    Protocol next hop: 10.255.70.103
    Indirect next hop: 2 no-forward
    State: <Secondary Active Int Ext>
    Local AS: 69 Peer AS: 69
    Age: 1:25:49 Metric2: 1
    AIGP 210
    Task: BGP_69.10.255.70.103+179
    Announcement bits (1): 0-green-l2vpn
    AS path: I
    Communities: target:11111:1 Layer2-info: encaps:VPLS,
    control flags:, mtu: 0
    Label-base: 800008, range: 8
    Localpref: 100
    Router ID: 10.255.70.103
    Primary Routing Table bgp.l2vpn.0

10.255.71.52:1:1:1/96 (1 entry, 1 announced)
  *L2VPN Preference: 170/-1
    Next-hop reference count: 5
    Protocol next hop: 10.255.71.52
    Indirect next hop: 0 -
    State: <Active Int Ext>
    Age: 1:31:40 Metric2: 1
    Task: green-l2vpn
    Announcement bits (1): 1-BGP.0.0.0.0+179
    AS path: I
    Communities: Layer2-info: encaps:VPLS, control flags:Site-Down,
    mtu: 0
    Label-base: 800016, range: 8, status-vector: 0x9F

10.255.71.52:1:5:1/96 (1 entry, 1 announced)
  *L2VPN Preference: 170/-101
    Next-hop reference count: 5
    Protocol next hop: 10.255.71.52
    Indirect next hop: 0 -
    State: <Active Int Ext>
    Age: 1:31:40 Metric2: 1
    Task: green-l2vpn
    Announcement bits (1): 1-BGP.0.0.0.0+179
    AS path: I
    Communities: Layer2-info: encaps:VPLS, control flags:, mtu: 0
    Label-base: 800008, range: 8, status-vector: 0x9F

...

```

```

12circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
  *L2CKT Preference: 7
    Next hop: via so-1/1/2.0 weight 1, selected
    Label-switched-path my-lsp
    Label operation: Push 100000[0]
    Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
    State: <Active Int>
    Local AS: 99
    Age: 10:21
    Task: 12 circuit
    Announcement bits (1): 0-LDP
    AS path: I
    VC Label 100000, MTU 1500, VLAN ID 512

```

show route detail (with BGP Multipath)

```

user@host> show route detail

10.1.1.8/30 (2 entries, 1 announced)
  *BGP Preference: 170/-101
    Next hop type: Router, Next hop index: 262142
    Address: 0x901a010
    Next-hop reference count: 2
    Source: 10.1.1.2
    Next hop: 10.1.1.2 via ge-0/3/0.1, selected
    Next hop: 10.1.1.6 via ge-0/3/0.5
    State: <Active Ext>
    Local AS: 1 Peer AS: 2
    Age: 5:04:43
    Task: BGP_2.10.1.1.2+59955
    Announcement bits (1): 0-KRT
    AS path: 2 I
    Accepted Multipath
    Localpref: 100
    Router ID: 1.1.1.2
  BGP Preference: 170/-101
    Next hop type: Router, Next hop index: 678
    Address: 0x8f97520
    Next-hop reference count: 9
    Source: 10.1.1.6
    Next hop: 10.1.1.6 via ge-0/3/0.5, selected
    State: <NotBest Ext>
    Inactive reason: Not Best in its group - Active preferred
    Local AS: 1 Peer AS: 2
    Age: 5:04:43
    Task: BGP_2.10.1.1.6+58198
    AS path: 2 I
    Accepted MultipathContrib
    Localpref: 100
    Router ID: 1.1.1.3

```

show route exact

| | |
|------------------------------------|--|
| Syntax | show route exact <i>destination-prefix</i> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switches) | show route exact <i>destination-prefix</i> <brief detail extensive terse> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. |
| Description | Display only the routes that exactly match the specified address or range of addresses. |
| Options | brief detail extensive terse —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief . <i>destination-prefix</i> —Address or range of addresses. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | view |
| List of Sample Output | show route exact on page 538 show route exact detail on page 538 show route exact extensive on page 538 show route exact terse on page 538 |
| Output Fields | For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command. |

Sample Output

show route exact

```
user@host> show route exact 207.17.136.0/24

inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
207.17.136.0/24    *[Static/5] 2d 03:30:22
                  > to 192.168.71.254 via fxp0.0
```

show route exact detail

```
user@host> show route exact 207.17.136.0/24 detail

inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
207.17.136.0/24 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 29
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 2d 3:30:26
    Task: RT
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I
```

show route exact extensive

```
user@host> show route exact 207.17.136.0/24 extensive

inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
207.17.136.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.0/24 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 29
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 1:25:18
    Task: RT
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I
```

show route exact terse

```
user@host> show route exact 207.17.136.0/24 terse

inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
* 207.17.136.0/24  S  5                      >192.168.71.254
```

show route export

| | |
|------------------------------------|---|
| Syntax | show route export <brief detail> <instance <instance-name> routing-table-name> <logical-system (all logical-system-name)> |
| Syntax (EX Series Switches) | show route export <brief detail> <instance <instance-name> routing-table-name> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. |
| Description | Display policy-based route export information. Policy-based export simplifies the process of exchanging route information between routing instances. |
| Options | <p>none—(Same as brief.) Display standard information about policy-based export for all instances and routing tables on all systems.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>instance <instance-name>—(Optional) Display a particular routing instance for which policy-based export is currently enabled.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>routing-table-name—(Optional) Display information about policy-based export for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the show route export inet command).</p> |
| Required Privilege Level | view |
| List of Sample Output | show route export on page 540 show route export detail on page 540 show route export instance detail on page 540 |
| Output Fields | Table 131 on page 539 lists the output fields for the show route export command. Output fields are listed in the approximate order in which they appear. |

Table 131: show route export Output Fields

| Field Name | Field Description | Level of Output |
|-----------------------------------|---|-------------------|
| Table or table-name | Name of the routing tables that either import or export routes. | All levels |
| Routes | Number of routes exported from this table into other tables. If a particular route is exported to different tables, the counter will only increment by one. | brief none |
| Export | Whether the table is currently exporting routes to other tables: Y or N (Yes or No). | brief none |

Table 131: show route export Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|----------------------|---|-----------------|
| Import | Tables currently importing routes from the originator table. (Not displayed for tables that are not exporting any routes.) | detail |
| Flags | (instance keyword only) Flags for this feature on this instance: <ul style="list-style-type: none"> config auto-policy—The policy was deduced from the configured IGP export policies. cleanup—Configuration information for this instance is no longer valid. config—The instance was explicitly configured. | detail |
| Options | (instance keyword only) Configured option displays the type of routing tables the feature handles: <ul style="list-style-type: none"> unicast—Indicates <i>instance.inet.0</i>. multicast—Indicates <i>instance.inet.2</i>. unicast multicast—Indicates <i>instance.inet.0</i> and <i>instance.inet.2</i>. | detail |
| Import policy | (instance keyword only) Policy that route export uses to construct the import-export matrix. Not displayed if the instance type is vrf . | detail |
| Instance | (instance keyword only) Name of the routing instance. | detail |
| Type | (instance keyword only) Type of routing instance: forwarding , non-forwarding , or vrf . | detail |

Sample Output

show route export

```

user@host> show route export
Table           Export      Routes
inet.0          N           0
black.inet.0    Y           3
red.inet.0      Y           4

```

show route export detail

```

user@host> show route export detail
inet.0                      Routes:      0
black.inet.0                Routes:      3
  Import: [ inet.0 ]
red.inet.0                  Routes:      4
  Import: [ inet.0 ]

```

show route export instance detail

```

user@host> show route export instance detail
Instance: master              Type: forwarding
  Flags: <config auto-policy> Options: <unicast multicast>
  Import policy: [ (ospf-master-from-red || isis-master-from-black) ]
Instance: black               Type: non-forwarding
Instance: red                 Type: non-forwarding

```

show route export vrf-target

| | |
|---------------------------------|---|
| Syntax | show route export vrf-target <brief detail> <community <i>community--regular-expression</i> > <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display the VPN routing and forwarding (VRF) target communities for which policy-based route export is currently distributing routes. This command is relevant when there are overlapping virtual private networks (VPNs). |
| Options | <p>none—Display standard information about all target communities.</p> <p>brief detail—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>community <i>community-regular-expression</i>—(Optional) Display information about the specified community.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show route export vrf-target on page 542 show route export vrf-target community on page 542 show route export vrf-target detail on page 542 |
| Output Fields | <p>Table 132 on page 541 lists the output fields for the show route export vrf-target command. Output fields are listed in the approximate order in which they appear.</p> |

Table 132: show route export vrf-target Output Fields

| Field Name | Field Description | Level of Output |
|---------------------------------|--|-----------------|
| Route target | Target communities for which auto-export is currently distributing routes. | brief none |
| Family | Routing table entries for the specified family. | brief none |
| <i>type-of-routing-table(s)</i> | Type of routing tables the feature handles: <ul style="list-style-type: none"> unicast—Indicates <i>instance.inet.0</i>. multicast—Indicates <i>instance.inet.2</i>. unicast multicast—Indicates <i>instance.inet.0</i> and <i>instance.inet.2</i>. | brief none |
| Import | Number of routing tables that are currently importing routes with this target community. Omitted for tables that are not importing routes. | brief none |

Table 132: show route export vrf-target Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|------------------------|--|-------------------|
| Export | Number of routing tables that are currently exporting routes with this target community. Omitted for tables that are not exporting routes. | brief none |
| Target | Target communities, family, and options for which auto-export is currently distributing routes. | detail |
| Import table(s) | Name of the routing tables that are importing a particular route target. | detail |
| Export table(s) | Name of the routing tables that are exporting a particular route target. | detail |

Sample Output

**show route export
vrf-target**

```
user@host> show route export vrf-target
Route Target      Family      Import      Export
69:1              inet        unicast     2           2
69:2              inet        unicast     2           2
```

**show route export
vrf-target community**

```
user@host> show route export vrf-target community target:69:1
Route Target      Family      Import      Export
69:1              inet        unicast     2           2
```

**show route export
vrf-target detail**

```
user@host> show route export vrf-target detail
Target: 1:12              inet        unicast
  Import table(s): vrf-11.inet.0 vrf-12.inet.0
  Export table(s): vrf-12.inet.0
Target: 1:13              inet        unicast
  Import table(s): vrf-12.inet.0 vrf-13.inet.0
  Export table(s): vrf-13.inet.0
```


show route extensive

| | |
|------------------------------------|--|
| Syntax | show route extensive <destination-prefix> <logical-system (all logical-system-name)> |
| Syntax (EX Series Switches) | show route extensive <destination-prefix> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. |
| Description | Display extensive information about the active entries in the routing tables. |
| Options | <p>none—Display all active entries in the routing table.</p> <p>destination-prefix—(Optional) Display active entries for the specified address or range of addresses.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show route extensive on page 549 show route extensive (Access Route) on page 555 show route extensive (Route Reflector) on page 555 show route extensive (FRR and LFA) on page 555 show route extensive (FRR and LFA) on page 556 |
| Output Fields | Table 133 on page 543 describes the output fields for the show route extensive command. Output fields are listed in the approximate order in which they appear. |

Table 133: show route extensive Output Fields

| Field Name | Field Description |
|----------------------------|--|
| <i>routing-table-name</i> | Name of the routing table (for example, inet.0). |
| <i>number destinations</i> | Number of destinations for which there are routes in the routing table. |
| <i>number routes</i> | Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> active (routes that are active). holddown (routes that are in the pending state before being declared inactive). hidden (routes that are not used because of a routing policy). |

Table 133: show route extensive Output Fields (*continued*)

| Field Name | Field Description |
|--|---|
| <i>route-destination</i> (entry, announced) | <p>Route destination (for example: 10.0.0.1/24). The entry value is the number of route for this destination, and the announced value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:</p> <ul style="list-style-type: none"> • MPLS-label (for example, 80001). • interface-name (for example, ge-1/0/2). • neighbor-address:control-word-status:encapsulation type:vc-id:source (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96). • neighbor-address—Address of the neighbor. • control-word-status—Whether the use of the control word has been negotiated for this virtual circuit: NoCtrlWord or CtrlWord. • encapsulation type—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport. • vc-id—Virtual circuit identifier. • source—Source of the advertisement: Local or Remote. |
| TSI | Protocol header information. |
| label stacking | <p>(Next-to-the-last-hop routing device for MPLS only) Depth of the MPLS label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels.</p> <ul style="list-style-type: none"> • S=0 route indicates that a packet with an incoming label stack depth of two or more exits this router with one fewer label (the label-popping operation is performed). • If there is no S= information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed). |
| [<i>protocol, preference</i>] | <p>Protocol from which the route was learned and the preference value for the route.</p> <ul style="list-style-type: none"> • +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • -—A hyphen indicates the last active route. • *—An asterisk indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route. <p>In every routing metric except for the BGP LocalPref attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the LocalPref value in the Preference2 field. For example, if the LocalPref value for Route 1 is 100, the Preference2 value is -101. If the LocalPref value for Route 2 is 155, the Preference2 value is -156. Route 2 is preferred because it has a higher LocalPref value and a lower Preference2 value.</p> |
| Level | <p>(IS-IS only). In IS-IS, a single autonomous system (AS) can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.</p> |
| Route Distinguisher | IP subnet augmented with a 64-bit prefix. |

Table 133: show route extensive Output Fields (*continued*)

| Field Name | Field Description |
|--|---|
| Next-hop type | Type of next hop. For a description of possible values for this field, see the Output Field table in the show route detail command. |
| Next-hop reference count | Number of references made to the next hop. |
| Flood nexthop branches exceed maximum message | Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel. |
| Source | IP address of the route source. |
| Next hop | Network layer address of the directly reachable neighboring system. |
| via | <p>Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word Selected. This field can also contain the following information:</p> <ul style="list-style-type: none"> • Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing. |
| Label-switched-path <i>lsp-path-name</i> | Name of the LSP used to reach the next hop. |
| Label operation | MPLS label and operation occurring at this routing device. The operation can be pop (where a label is removed from the top of the stack), push (where another label is added to the label stack), or swap (where a label is replaced by another label). |
| Offset | Whether the metric has been increased or decreased by an offset value. |
| Interface | (Local only) Local interface name. |
| Protocol next hop | Network layer address of the remote routing device that advertised the prefix. This address is used to recursively derive a forwarding next hop. |
| <i>label-operation</i> | MPLS label and operation occurring at this routing device. The operation can be pop (where a label is removed from the top of the stack), push (where another label is added to the label stack), or swap (where a label is replaced by another label). |
| Indirect next hops | When present, a list of nodes that are used to resolve the path to the next-hop destination, in the order that they are resolved. |
| State | State of the route (a route can be in more than one state). See the Output Field table in the show route detail command. |

Table 133: show route extensive Output Fields (*continued*)

| Field Name | Field Description |
|------------------------|---|
| Session ID | The BFD session ID number that represents the protection using MPLS fast reroute (FRR) and loop-free alternate (LFA). |
| Inactive reason | <p>If the route is inactive, the reason for its current state is indicated. Typical reasons include:</p> <ul style="list-style-type: none"> • Active preferred—Currently active route was selected over this route. • Always compare MED—Path with a lower multiple exit discriminator (MED) is available. • AS path—Shorter AS path is available. • Cisco Non-deterministic MED selection—Cisco nondeterministic MED is enabled and a path with a lower MED is available. • Cluster list length—Path with a shorter cluster list length is available. • Forwarding use only—Path is only available for forwarding purposes. • IGP metric—Path through the next hop with a lower IGP metric is available. • IGP metric type—Path with a lower OSPF link-state advertisement type is available. • Interior > Exterior > Exterior via Interior—Direct, static, IGP, or EBGP path is available. • Local preference—Path with a higher local preference value is available. • Next hop address—Path with a lower metric next hop is available. • No difference—Path from a neighbor with a lower IP address is available. • Not Best in its group—Occurs when multiple peers of the same external AS advertise the same prefix and are grouped together in the selection process. When this reason is displayed, an additional reason is provided (typically one of the other reasons listed). • Number of gateways—Path with a higher number of next hops is available. • Origin—Path with a lower origin code is available. • OSPF version—Path does not support the indicated OSPF version. • RIB preference—Route from a higher-numbered routing table is available. • Route distinguisher—64-bit prefix added to IP subnets to make them unique. • Route metric or MED comparison—Route with a lower metric or MED is available. • Route preference—Route with a lower preference value is available. • Router ID—Path through a neighbor with a lower ID is available. • Unusable path—Path is not usable because of one of the following conditions: the route is damped, the route is rejected by an import policy, or the route is unresolved. • Update source—Last tiebreaker is the lowest IP address value. |
| Local AS | Autonomous system (AS) number of the local routing device. |
| Age | How long the route has been known. |
| AIGP | Accumulated interior gateway protocol (AIGP) BGP attribute. |
| Metric | Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value. |
| MED-plus-IGP | Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added. |

Table 133: show route extensive Output Fields (*continued*)

| Field Name | Field Description |
|--------------------------------------|---|
| TTL-Action | <p>For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signaled and LDP-signaled LSPs or for specific VRF routing instances.</p> <p>For sample output, see show route table.</p> |
| Task | Name of the protocol that has added the route. |
| Announcement bits | List of protocols that announce this route. n-Resolve inet indicates that the route is used for route resolution for next hops found in the routing table. n is an index used by Juniper Networks customer support only. |
| AS path | <p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> • I—IGP. • E—EGP. • Recorded—The AS path is recorded by the sample process (sampled). • ?—Incomplete; typically, the AS path was aggregated. <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> • []—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device, or if AS path prepending is configured. • { }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • ()—Parentheses enclose a confederation. • ([])—Parentheses and brackets enclose a confederation set. <p>NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p> |
| AS path: I <Originator> | (For route reflected output only) Originator ID attribute set by the route reflector. |
| VC Label | MPLS label assigned to the Layer 2 circuit virtual connection. |
| MTU | Maximum transmission unit (MTU) of the Layer 2 circuit. |
| VLAN ID | VLAN identifier of the Layer 2 circuit. |
| Cluster list | (For route reflected output only) Cluster ID sent by the route reflector. |
| Originator ID | (For route reflected output only) Address of router that originally sent the route to the route reflector. |
| Prefixes bound to route | Forwarding equivalent class (FEC) bound to this route. Applicable only to routes installed by LDP. |
| Communities | Community path attribute for the route. See the Output Field table in the show route detail command for all possible values for this field. |

Table 133: show route extensive Output Fields (*continued*)

| Field Name | Field Description |
|------------------------------|--|
| Layer2-info: encaps | Layer 2 encapsulation (for example, VPLS). |
| control flags | Control flags: none or Site Down. |
| mtu | Maximum transmission unit (MTU) information. |
| Label-Base, range | First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device. |
| status vector | Layer 2 VPN and VPLS network layer reachability information (NLRI). |
| Localpref | Local preference value included in the route. |
| Router ID | BGP router ID as advertised by the neighbor in the open message. |
| Primary Routing Table | In a routing table group, the name of the primary routing table in which the route resides. |
| Secondary Tables | In a routing table group, the name of one or more secondary tables in which the route resides. |
| Originating RIB | Name of the routing table whose active route was used to determine the forwarding next-hop entry in the resolution database. For example, in the case of inet.0 resolving through inet.0 and inet.3, this field indicates which routing table, inet.0 or inet.3, provided the best path for a particular prefix. |
| Node path count | Number of nodes in the path. |
| Forwarding nexthops | Number of forwarding next hops. The forwarding next hop is the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it. |

Sample Output

show route extensive

```

user@host> show route extensive
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
    *Static Preference: 5
        Next-hop reference count: 29
        Next hop: 192.168.71.254 via fxp0.0, selected
        State: <Active NoReadvrt Int Ext>
        Local AS: 69
        Age: 1:34:06
        Task: RT
        Announcement bits (2): 0-KRT 3-Resolve tree 2
        AS path: I

10.31.1.0/30 (2 entries, 1 announced)
    *Direct Preference: 0
        Next hop type: Interface
        Next-hop reference count: 2
        Next hop: via so-0/3/0.0, selected
        State: <Active Int>
        Local AS: 69
        Age: 1:32:40
        Task: IF
        Announcement bits (1): 3-Resolve tree 2
        AS path: I
    OSPF Preference: 10
        Next-hop reference count: 1
        Next hop: via so-0/3/0.0, selected
        State: <Int>
        Inactive reason: Route Preference
        Local AS: 69
        Age: 1:32:40 Metric: 1
        Area: 0.0.0.0
        Task: OSPF
        AS path: I

10.31.1.1/32 (1 entry, 1 announced)
    *Local Preference: 0
        Next hop type: Local
        Next-hop reference count: 7
        Interface: so-0/3/0.0
        State: <Active NoReadvrt Int>
        Local AS: 69
        Age: 1:32:43
        Task: IF
        Announcement bits (1): 3-Resolve tree 2
        AS path: I

...

10.31.2.0/30 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.31.2.0/30 -> {10.31.1.6}
    *OSPF Preference: 10
        Next-hop reference count: 9
        Next hop: via so-0/3/0.0
        Next hop: 10.31.1.6 via ge-3/1/0.0, selected

```

```

        State: <Active Int>
        Local AS: 69
        Age: 1:32:19    Metric: 2
        Area: 0.0.0.0
        Task: OSPF
        Announcement bits (2): 0-KRT 3-Resolve tree 2
        AS path: I

...

224.0.0.2/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 224.0.0.2/32 -> {}
    *PIM    Preference: 0
            Next-hop reference count: 18
            State: <Active NoReadvrt Int>
            Local AS: 69
            Age: 1:34:08
            Task: PIM Recv
            Announcement bits (2): 0-KRT 3-Resolve tree 2
            AS path: I

...

224.0.0.22/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 224.0.0.22/32 -> {}
    *IGMP   Preference: 0
            Next-hop reference count: 18
            State: <Active NoReadvrt Int>
            Local AS: 69
            Age: 1:34:06
            Task: IGMP
            Announcement bits (2): 0-KRT 3-Resolve tree 2
            AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.255.70.103/32 (1 entry, 1 announced)
    State: <FlashAll>
    *RSVP   Preference: 7
            Next-hop reference count: 6
            Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
            Label-switched-path green-r1-r3
            Label operation: Push 100096
            State: <Active Int>
            Local AS: 69
            Age: 1:28:12    Metric: 2
            Task: RSVP
            Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
            AS path: I

10.255.71.238/32 (1 entry, 1 announced)
    State: <FlashAll>
    *RSVP   Preference: 7
            Next-hop reference count: 6
            Next hop: via so-0/3/0.0 weight 0x1, selected
            Label-switched-path green-r1-r2
            State: <Active Int>
            Local AS: 69
            Age: 1:28:12    Metric: 1

```



```

Task: RSVP
Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
AS path: I

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

...

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

47.0005.80ff.f800.0000.0108.0001.0102.5507.1052/152 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
    State: <Active Int>
    Local AS: 69
    Age: 1:34:07
    Task: IF
    AS path: I

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

0 (1 entry, 1 announced)
TSI:
KRT in-kernel 0 /36 -> {}
  *MPLS Preference: 0
    Next hop type: Receive
    Next-hop reference count: 6
    State: <Active Int>
    Local AS: 69
    Age: 1:34:08 Metric: 1
    Task: MPLS
    Announcement bits (1): 0-KRT
    AS path: I

...

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
299776 (1 entry, 1 announced)
TSI:
KRT in-kernel 299776 /52 -> {Flood}
  *RSVP Preference: 7
    Next hop type: Flood
    Next-hop reference count: 130
    Flood nexthop branches exceed maximum
    Address: 0x8ea65d0

...

800010 (1 entry, 1 announced)

TSI:
KRT in-kernel 800010 /36 -> {vt-3/2/0.32769}
  *VPLS Preference: 7
    Next-hop reference count: 2
    Next hop: via vt-3/2/0.32769, selected
    Label operation: Pop
    State: <Active Int>
    Age: 1:31:53
    Task: Common L2 VC
    Announcement bits (1): 0-KRT

```

```

AS path: I

vt-3/2/0.32769 (1 entry, 1 announced)
TSI:
KRT in-kernel vt-3/2/0.32769.0 /16 -> {indirect(1048574)}
  *VPLS Preference: 7
    Next-hop reference count: 2
    Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
    Label-switched-path green-r1-r3
    Label operation: Push 800012, Push 100096(top)
    Protocol next hop: 10.255.70.103
    Push 800012
    Indirect next hop: 87272e4 1048574
    State: <Active Int>
    Age: 1:31:53 Metric2: 2
    Task: Common L2 VC
    Announcement bits (2): 0-KRT 1-Common L2 VC
    AS path: I
    Communities: target:11111:1 Layer2-info: encaps:VPLS,
    control flags:, mtu: 0
    Indirect next hops: 1
      Protocol next hop: 10.255.70.103 Metric: 2
      Push 800012
      Indirect next hop: 87272e4 1048574
      Indirect path forwarding next hops: 1
        Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1
        10.255.70.103/32 Originating RIB: inet.3
        Metric: 2 Node path count: 1
        Forwarding nexthops: 1
          Nexthop: 10.31.1.6 via ge-3/1/0.0

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

abcd::10:255:71:52/128 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
    State: <Active Int>
    Local AS: 69
    Age: 1:34:07
    Task: IF
    AS path: I

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:34:07
    Task: IF
    AS path: I

ff02::2/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::2/128 -> {}
  *PIM Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>

```

```

        Local AS:    69
        Age: 1:34:08
        Task: PIM Recv6
        Announcement bits (1): 0-KRT
        AS path: I

ff02::d/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::d/128 -> {}
    *PIM      Preference: 0
              Next-hop reference count: 18
              State: <Active NoReadvrt Int>
              Local AS:    69
              Age: 1:34:08
              Task: PIM Recv6
              Announcement bits (1): 0-KRT
              AS path: I

ff02::16/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::16/128 -> {}
    *MLD      Preference: 0
              Next-hop reference count: 18
              State: <Active NoReadvrt Int>
              Local AS:    69
              Age: 1:34:06
              Task: MLD
              Announcement bits (1): 0-KRT
              AS path: I

private.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
    *Direct Preference: 0
              Next hop type: Interface
              Next-hop reference count: 1
              Next hop: via lo0.16385, selected
              State: <Active NoReadvrt Int>
              Age: 1:34:07
              Task: IF
              AS path: I

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

10.255.70.103:1:3:1/96 (1 entry, 1 announced)
    *BGP      Preference: 170/-101
              Route Distinguisher: 10.255.70.103:1
              Next-hop reference count: 7
              Source: 10.255.70.103
              Protocol next hop: 10.255.70.103
              Indirect next hop: 2 no-forward
              State: <Secondary Active Int Ext>
              Local AS:    69 Peer AS:    69
              Age: 1:28:12   Metric2: 1
              Task: BGP_69.10.255.70.103+179
              Announcement bits (1): 0-green-l2vpn
              AS path: I
              Communities: target:11111:1 Layer2-info: encaps:VPLS,
              control flags:, mtu: 0
              Label-base: 800008, range: 8
              Localpref: 100

```

```

Router ID: 10.255.70.103
Primary Routing Table bgp.l2vpn.0

10.255.71.52:1:1:1/96 (1 entry, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 8699540
    *L2VPN Preference: 170/-1
        Next-hop reference count: 5
        Protocol next hop: 10.255.71.52
        Indirect next hop: 0 -
        State: <Active Int Ext>
        Age: 1:34:03 Metric2: 1
        Task: green-l2vpn
        Announcement bits (1): 1-BGP.0.0.0+179
        AS path: I
        Communities: Layer2-info: encaps:VPLS, control flags:Site-Down,
        mtu: 0
        Label-base: 800016, range: 8, status-vector: 0x9F

10.255.71.52:1:5:1/96 (1 entry, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 8699528
    *L2VPN Preference: 170/-101
        Next-hop reference count: 5
        Protocol next hop: 10.255.71.52
        Indirect next hop: 0 -
        State: <Active Int Ext>
        Age: 1:34:03 Metric2: 1
        Task: green-l2vpn
        Announcement bits (1): 1-BGP.0.0.0+179
        AS path: I
        Communities: Layer2-info: encaps:VPLS, control flags:, mtu: 0
        Label-base: 800008, range: 8, status-vector: 0x9F

...

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
TSI:

10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
    *L2CKT Preference: 7
        Next hop: via so-1/1/2.0 weight 1, selected
        Label-switched-path my-lsp
        Label operation: Push 100000[0]
        Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
        State: <Active Int>
        Local AS: 99
        Age: 10:21
        Task: l2 circuit
        Announcement bits (1): 0-LDP
        AS path: I
        VC Label 100000, MTU 1500, VLAN ID 512

55.0.0.0/24 (1 entry, 1 announced)
TSI:
KRT queued (pending) add
    55.0.0.0/24 -> {Push 300112}
        *BGP Preference: 170/-101
            Next hop type: Router

```

```

Address: 0x925c208
Next-hop reference count: 2
Source: 10.0.0.9
Next hop: 10.0.0.9 via ge-1/2/0.15, selected
Label operation: Push 300112
Label TTL action: prop-ttl
State: <Active Ext>
Local AS: 7019 Peer AS: 13979
Age: 1w0d 23:06:56
AIGP: 25
Task: BGP_13979.10.0.0.9+56732
Announcement bits (1): 0-KRT
AS path: 13979 7018 I
Accepted
Route Label: 300112
Localpref: 100
Router ID: 10.9.9.1

```

show route extensive (Access Route)

```

user@host> show route 13.160.0.102 extensive
inet.0: 39256 destinations, 39258 routes (39255 active, 0 holddown, 1 hidden)
13.160.0.102/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 13.160.0.102/32 -> {13.160.0.2}
OSPF area : 0.0.0.0, LSA ID : 13.160.0.102, LSA type : Extern
  *Access Preference: 13
    Next-hop reference count: 78472
    Next hop: 13.160.0.2 via fe-0/0/0.0, selected
    State: <Active Int>
  Age: 12
    Task: RPD Unix Domain Server./var/run/rpd_serv.local
    Announcement bits (2): 0-KRT 1-OSPFv2
    AS path: I

```

show route extensive (Route Reflector)

```

user@host> show route extensive
1.0.0.0/8 (1 entry, 1 announced)

TSI:
KRT in-kernel 1.0.0.0/8 -> {indirect(40)}
  *BGP Preference: 170/-101
    Source: 192.168.4.214
    Protocol next hop: 207.17.136.192 Indirect next hop: 84ac908 40
    State: <Active Int Ext>
    Local AS: 10458 Peer AS: 10458
    Age: 3:09 Metric: 0 Metric2: 0
    Task: BGP_10458.192.168.4.214+1033
    Announcement bits (2): 0-KRT 4-Resolve inet.0
    AS path: 3944 7777 I <Originator>
    Cluster list: 1.1.1.1
    Originator ID: 10.255.245.88
    Communities: 7777:7777
    Localpref: 100
    Router ID: 4.4.4.4
    Indirect next hops: 1
      Protocol next hop: 207.17.136.192 Metric: 0
      Indirect next hop: 84ac908 40
      Indirect path forwarding next hops: 0
      Next hop type: Discard

```

```

user@host> show route 20:31:2:0 extensive

```

```

show route extensive      inet.0: 46 destinations, 49 routes (45 active, 0 holddown, 1 hidden)
(FRR and LFA)           20.31.2.0/24 (2 entries, 1 announced)
                        State: FlashAll

TSI:
KRT in-kernel 20.31.2.0/24 -> {Push 299776, Push 299792}
    *RSVP Preference: 7/1
        Next hop type: Router, Next hop index: 1048574
        Address: 0xbbbc010
        Next-hop reference count: 5
        Next hop: 10.31.1.2 via ge-2/1/8.0 weight 0x1, selected
        Label-switched-path europa-d-to-europa-e
        Label operation: Push 299776
        Label TTL action: prop-ttl
        Session Id: 0x201
        Next hop: 10.31.2.2 via ge-2/1/4.0 weight 0x4001
        Label-switched-path europa-d-to-europa-e
        Label operation: Push 299792
        Label TTL action: prop-ttl
        Session Id: 0x202
        State: Active Int
        Local AS: 100
        Age: 5:31 Metric: 2
        Task: RSVP
        Announcement bits (1): 0-KRT
        AS path: I
    OSPF Preference: 10
        Next hop type: Router, Next hop index: 615
        Address: 0xb9d78c4
        Next-hop reference count: 7
        Next hop: 10.31.1.2 via ge-2/1/8.0, selected
        Session Id: 0x201
        State: Int
        Inactive reason: Route Preference
        Local AS: 100
        Age: 5:35 Metric: 3
        Area: 0.0.0.0
        Task: OSPF
        AS path: I

```

```

show route extensive      user@host> show route 20.31.2.0 extensive
(FRR and LFA)           inet.0: 46 destinations, 49 routes (45 active, 0 holddown, 1 hidden)
                        20.31.2.0/24 (2 entries, 1 announced)
                        State: FlashAll

TSI:
KRT in-kernel 20.31.2.0/24 -> {Push 299776, Push 299792}
    *RSVP Preference: 7/1
        Next hop type: Router, Next hop index: 1048574
        Address: 0xbbbc010
        Next-hop reference count: 5
        Next hop: 10.31.1.2 via ge-2/1/8.0 weight 0x1, selected
        Label-switched-path europa-d-to-europa-e
        Label operation: Push 299776
        Label TTL action: prop-ttl
        Session Id: 0x201
        Next hop: 10.31.2.2 via ge-2/1/4.0 weight 0x4001
        Label-switched-path europa-d-to-europa-e
        Label operation: Push 299792
        Label TTL action: prop-ttl
        Session Id: 0x202
        State: Active Int
        Local AS: 100

```

```
Age: 5:31 Metric: 2
Task: RSVP
Announcement bits (1): 0-KRT
AS path: I
OSPF Preference: 10
Next hop type: Router, Next hop index: 615
Address: 0xb9d78c4
Next-hop reference count: 7
Next hop: 10.31.1.2 via ge-2/1/8.0, selected
Session Id: 0x201
State: Int
Inactive reason: Route Preference
Local AS: 100
Age: 5:35 Metric: 3
Area: 0.0.0.0
Task: OSPF
AS path: I
```

show route flow validation

| | |
|------------------------------------|--|
| Syntax | show route flow validation <brief detail> <ip-prefix> <table table-name> <logical-system (all logical-system-name)> |
| Syntax (EX Series Switches) | show route flow validation <brief detail> <ip-prefix> <table table-name> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. |
| Description | Display flow route information. |
| Options | <p>none—Display flow route information.</p> <p>brief detail—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>ip-prefix—(Optional) IP address for the flow route.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>table table-name—(Optional) Display flow route information for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the show route flow validation inet command).</p> |
| Required Privilege Level | view |
| List of Sample Output | show route flow validation on page 559 |
| Output Fields | Table 134 on page 558 lists the output fields for the show route flow validation command. Output fields are listed in the approximate order in which they appear. |

Table 134: show route flow validation Output Fields

| Field Name | Field Description | Level of Output |
|------------------------------------|--|-----------------|
| <i>routing-table-name</i> | Name of the routing table (for example, inet.0). | All levels |
| <i>prefix</i> | Route address. | All levels |
| Active unicast route | Active route in the routing table. | All levels |
| Dependent flow destinations | Number of flows for which there are routes in the routing table. | All levels |

Table 134: show route flow validation Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------------|---|-----------------|
| Origin | Source of the route flow. | All levels |
| Neighbor AS | Autonomous system identifier of the neighbor. | All levels |
| Flow destination | Number of entries and number of destinations that match the route flow. | All levels |
| Unicast best match | Destination that is the best match for the route flow. | All levels |
| Flags | Information about the route flow. | All levels |

Sample Output

**show route flow
validation**

```
user@host> show route flow validation
inet.0:
10.0.5.0/24Active unicast route
Dependent flow destinations: 1
Origin: 192.168.224.218, Neighbor AS: 65001
Flow destination (3 entries, 1 match origin)
Unicast best match: 10.0.5.0/24
Flags: SubtreeApex Consistent
```

show route forwarding-table

| | |
|-----------------------------------|---|
| Syntax | <pre>show route forwarding-table <detail extensive summary> <all> <ccc interface-name> <destination destination-prefix> <family family matching matching> <interface-name interface-name> <label name> <matching matching> <multicast> <table (default logical-system-name/routing-instance-name routing-instance-name)> <vlan (all vlan-name)> <vpn vpn></pre> |
| Syntax (MX Series Routers) | <pre>show route forwarding-table <detail extensive summary> <all> <bridge-domain (all domain-name)> <ccc interface-name> <destination destination-prefix> <family family matching matching> <interface-name interface-name> <label name> <learning-vlan-id learning-vlan-id> <matching matching> <multicast> <table (default logical-system-name/routing-instance-name routing-instance-name)> <vlan (all vlan-name)> <vpn vpn></pre> |
| Syntax (Routing Matrix) | <pre>show route forwarding-table <detail extensive summary> <all> <ccc interface-name> <destination destination-prefix> <family family matching matching> <interface-name interface-name> <matching matching> <label name> <lcc number> <multicast> <table routing-instance-name> <vpn vpn></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Option bridge-domain introduced in Junos OS Release 7.5</p> <p>Option learning-vlan-id introduced in Junos OS Release 8.4</p> <p>Options all and vlan introduced in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> |

Description Display the Routing Engine's forwarding table, including the network-layer prefixes and their next hops. This command is used to help verify that the routing protocol process has relayed the correction information to the forwarding table. The Routing Engine constructs and maintains one or more routing tables. From the routing tables, the Routing Engine derives a table of active routes, called the forwarding table.



NOTE: The Routing Engine copies the forwarding table to the Packet Forwarding Engine, the part of the router that is responsible for forwarding packets. To display the entries in the Packet Forwarding Engine's forwarding table, use the **show pfe route** command.

Options **none**—Display the routes in the forwarding tables. By default, the **show route forwarding-table** command does not display information about private, or internal, forwarding tables.

detail | extensive | summary—(Optional) Display the specified level of output.

all—(Optional) Display routing table entries for all forwarding tables, including private, or internal, tables.

bridge-domain (all | bridge-domain-name)—(MX Series routers only) (Optional) Display route entries for all bridge domains or the specified bridge domain.

ccc interface-name—(Optional) Display route entries for the specified circuit cross-connect interface.

destination destination-prefix—(Optional) Destination prefix.

family family—(Optional) Display routing table entries for the specified family: **fibre-channel**, **fmembers**, **inet**, **inet6**, **iso**, **mpls**, **tnp**, **unix**, **vpls**, or **vlan-classification**.

interface-name interface-name—(Optional) Display routing table entries for the specified interface.

label name—(Optional) Display route entries for the specified label.

lcc number—(Routing Matrix only) (Optional) On a routing matrix composed of a TX Matrix Plus router and T640 routers configured in the routing matrix, display information for the specified T640 router (or line-card chassis) connected to the TX Matrix router. On a routing matrix composed of the TX Matrix Plus router and T1600 routers configured in the routing matrix, display information for the specified T1600 router (or line-card chassis) connected to the TX Matrix Plus router. Replace **number** with a value from 0 through 3.

learning-vlan-id learning-vlan-id—(MX Series routers only) (Optional) Display learned information for all VLANs or for the specified VLAN.

matching matching—(Optional) Display routing table entries matching the specified prefix or prefix length.

multicast—(Optional) Display routing table entries for multicast routes.

table (**default** | *logical-system-name/routing-instance-name* |

routing-instance-name)—(Optional) Display route entries for all the routing tables in the main routing instance or for the specified routing instance. If your device supports logical systems, you can also display route entries for the specified logical system and routing instance. To view the routing instances on your device, use the [show route instance](#) command.

vlan (**all** | *vlan-name*)—(Optional) Display information for all VLANs or for the specified VLAN.

vpn *vpn*—(Optional) Display routing table entries for a specified VPN.

Required Privilege Level

view

List of Sample Output

[show route forwarding-table on page 565](#)
[show route forwarding-table detail on page 565](#)
[show route forwarding-table destination extensive \(Weights and Balances\) on page 566](#)
[show route forwarding-table extensive on page 567](#)
[show route forwarding-table extensive \(RPF\) on page 568](#)
[show route forwarding-table family mpls on page 570](#)
[show route forwarding-table family vpls on page 570](#)
[show route forwarding-table family vpls extensive on page 570](#)
[show route forwarding-table table default on page 571](#)
[show route forwarding-table table](#)
[logical-system-name/routing-instance-name on page 572](#)
[show route forwarding-table vpn on page 573](#)

Output Fields

[Table 135 on page 562](#) lists the output fields for the **show route forwarding-table** command. Output fields are listed in the approximate order in which they appear. Field names might be abbreviated (as shown in parentheses) when no level of output is specified, or when the **detail** keyword is used instead of the **extensive** keyword.

Table 135: show route forwarding-table Output Fields

| Field Name | Field Description | Level of Output |
|----------------|--|--------------------------------|
| Logical system | Name of the logical system. This field is displayed if you specify the table <i>logical-system-name/routing-instance-name</i> option on a device that is configured for and supports logical systems. | All levels |
| Routing table | Name of the routing table (for example, inet, inet6, mpls). | All levels |
| Address family | Address family (for example, IP, IPv6, ISO, MPLS, and VPLS). | All levels |
| Destination | Destination of the route. | detail extensive |

Table 135: show route forwarding-table Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------------------|--|-------------------------|
| Route Type (Type) | How the route was placed into the forwarding table. When the detail keyword is used, the route type might be abbreviated (as shown in parentheses): <ul style="list-style-type: none"> • cloned (clon)—(TCP or multicast only) Cloned route. • destination (dest)—Remote addresses directly reachable through an interface. • destination down (iddn)—Destination route for which the interface is unreachable. • interface cloned (ifcl)—Cloned route for which the interface is unreachable. • route down (ifdn)—Interface route for which the interface is unreachable. • ignore (ignr)—Ignore this route. • interface (intf)—Installed as a result of configuring an interface. • permanent (perm)—Routes installed by the kernel when the routing table is initialized. • user—Routes installed by the routing protocol process or as a result of the configuration. | All levels |
| Route Reference (RtRef) | Number of routes to reference. | detail extensive |
| Flags | Route type flags: <ul style="list-style-type: none"> • none—No flags are enabled. • accounting—Route has accounting enabled. • cached—Cache route. • incoming-iface interface-number—Check against incoming interface. • prefix load balance—Load balancing is enabled for this prefix. • rt nh decoupled—Route has been decoupled from the next hop to the destination. • sent to PFE—Route has been sent to the Packet Forwarding Engine. • static—Static route. | extensive |
| Next hop | IP address of the next hop to the destination. | detail extensive |

Table 135: show route forwarding-table Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-----------------------------------|--|------------------------------|
| Next hop Type (Type) | <p>Next-hop type. When the detail keyword is used, the next-hop type might be abbreviated (as indicated in parentheses):</p> <ul style="list-style-type: none"> • broadcast (bcst)—Broadcast. • deny—Deny. • discard (dscd)—Discard. • hold—Next hop is waiting to be resolved into a unicast or multicast type. • indexed (idxd)—Indexed next hop. • indirect (indr)—Indirect next hop. • local (locl)—Local address on an interface. • routed multicast (mcrst)—Regular multicast next hop. • multicast (mcst)—Wire multicast next hop (limited to the LAN). • multicast discard (mdsc)—Multicast discard. • multicast group (mgrp)—Multicast group member. • receive (rcv)—Receive. • reject (rjct)—Discard. An ICMP unreachable message was sent. • resolve (rslv)—Resolving the next hop. • unicast (ucst)—Unicast. • unilist (ulst)—List of unicast next hops. A packet sent to this next hop goes to any next hop in the list. | detail extensive |
| Index | Software index of the next hop that is used to route the traffic for a given prefix. | detail extensive none |
| Route interface-index | Logical interface index from which the route is learned. For example, for interface routes, this is the logical interface index of the route itself. For static routes, this field is zero. For routes learned through routing protocols, this is the logical interface index from which the route is learned. | extensive |
| Reference (NhRef) | Number of routes that refer to this next hop. | detail extensive none |
| Next-hop interface (Netif) | Interface used to reach the next hop. | detail extensive none |
| Weight | Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible (see the Balance field description). | extensive |
| Balance | Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a router is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing. | extensive |
| RPF interface | List of interfaces from which the prefix can be accepted. Reverse path forwarding (RPF) information is displayed only when rpf-check is configured on the interface. | extensive |

Sample Output

```
show route
forwarding-table
```

```
user@host> show route forwarding-table
```

```
Routing table: default.inet
```

```
Internet:
```

| Destination | Type | RtRef | Next hop | Type | Index | NhRef | Netif |
|------------------|------|-------|-------------------|------|-------|-------|------------|
| default | perm | 0 | | rjct | 46 | 4 | |
| 0.0.0.0/32 | perm | 0 | | dscd | 44 | 1 | |
| 1.1.1.0/24 | ifdn | 0 | | rslv | 608 | 1 | ge-2/0/1.0 |
| 1.1.1.0/32 | iddn | 0 | 1.1.1.0 | recv | 606 | 1 | ge-2/0/1.0 |
| 1.1.1.1/32 | user | 0 | | rjct | 46 | 4 | |
| 1.1.1.1/32 | intf | 0 | 1.1.1.1 | loc1 | 607 | 2 | |
| 1.1.1.1/32 | iddn | 0 | 1.1.1.1 | loc1 | 607 | 2 | |
| 1.1.1.255/32 | iddn | 0 | ff:ff:ff:ff:ff:ff | bcst | 605 | 1 | ge-2/0/1.0 |
| 10.0.0.0/24 | intf | 0 | | rslv | 616 | 1 | ge-2/0/0.0 |
| 10.0.0.0/32 | dest | 0 | 10.0.0.0 | recv | 614 | 1 | ge-2/0/0.0 |
| 10.0.0.1/32 | intf | 0 | 10.0.0.1 | loc1 | 615 | 2 | |
| 10.0.0.1/32 | dest | 0 | 10.0.0.1 | loc1 | 615 | 2 | |
| 10.0.0.255/32 | dest | 0 | 10.0.0.255 | bcst | 613 | 1 | ge-2/0/0.0 |
| 10.1.1.0/24 | ifdn | 0 | | rslv | 612 | 1 | ge-2/0/1.0 |
| 10.1.1.0/32 | iddn | 0 | 10.1.1.0 | recv | 610 | 1 | ge-2/0/1.0 |
| 10.1.1.1/32 | user | 0 | | rjct | 46 | 4 | |
| 10.1.1.1/32 | intf | 0 | 10.1.1.1 | loc1 | 611 | 2 | |
| 10.1.1.1/32 | iddn | 0 | 10.1.1.1 | loc1 | 611 | 2 | |
| 10.1.1.255/32 | iddn | 0 | ff:ff:ff:ff:ff:ff | bcst | 609 | 1 | ge-2/0/1.0 |
| 10.209.0.0/16 | user | 0 | 10.209.63.254 | ucst | 419 | 20 | fxp0.0 |
| 10.209.0.0/16 | user | 1 | 0:12:1e:ca:98:0 | ucst | 419 | 20 | fxp0.0 |
| 10.209.0.0/18 | intf | 0 | | rslv | 418 | 1 | fxp0.0 |
| 10.209.0.0/32 | dest | 0 | 10.209.0.0 | recv | 416 | 1 | fxp0.0 |
| 10.209.2.131/32 | intf | 0 | 10.209.2.131 | loc1 | 417 | 2 | |
| 10.209.2.131/32 | dest | 0 | 10.209.2.131 | loc1 | 417 | 2 | |
| 10.209.17.55/32 | dest | 0 | 0:30:48:5b:78:d2 | ucst | 435 | 1 | fxp0.0 |
| 10.209.63.42/32 | dest | 0 | 0:23:7d:58:92:ca | ucst | 434 | 1 | fxp0.0 |
| 10.209.63.254/32 | dest | 0 | 0:12:1e:ca:98:0 | ucst | 419 | 20 | fxp0.0 |
| 10.209.63.255/32 | dest | 0 | 10.209.63.255 | bcst | 415 | 1 | fxp0.0 |
| 10.227.0.0/16 | user | 0 | 10.209.63.254 | ucst | 419 | 20 | fxp0.0 |

```
...
```

```
Routing table: iso
```

```
ISO:
```

| Destination | Type | RtRef | Next hop | Type | Index | NhRef | Netif |
|--|------|-------|----------|------|-------|-------|-------|
| default | perm | 0 | | rjct | 27 | 1 | |
| 47.0005.80ff.f800.0000.0108.0003.0102.5524.5220.00 | intf | 0 | loc1 28 | | | 1 | |

```
Routing table: inet6
```

```
Internet6:
```

| Destination | Type | RtRef | Next hop | Type | Index | NhRef | Netif |
|-------------|------|-------|----------|------|-------|-------|-------|
| default | perm | 0 | | rjct | 6 | 1 | |
| ff00::/8 | perm | 0 | | mdsc | 4 | 1 | |
| ff02::1/128 | perm | 0 | ff02::1 | mcst | 3 | 1 | |

```
Routing table: ccc
```

```
MPLS:
```

| Interface.Label | Type | RtRef | Next hop | Type | Index | NhRef | Netif |
|-----------------------|------|-------|----------|------|-------|-------|-------|
| default | perm | 0 | | rjct | 16 | 1 | |
| 100004(top)fe-0/0/1.0 | | | | | | | |

show route forwarding-table detail

```
user@host> show route forwarding-table detail
```

```
Routing table: inet
```

```
Internet:
```

| Destination | Type | RtRef | Next hop | Type | Index | NhRef | Netif |
|-----------------|------|-------|------------------|------|-------|-------|------------|
| default | user | 2 | 0:90:69:8e:b1:1b | ucst | 132 | 4 | fxp0.0 |
| default | perm | 0 | | rjct | 14 | 1 | |
| 10.1.1.0/24 | intf | 0 | ff.3.0.21 | ucst | 322 | 1 | so-5/3/0.0 |
| 10.1.1.0/32 | dest | 0 | 10.1.1.0 | recv | 324 | 1 | so-5/3/0.0 |
| 10.1.1.1/32 | intf | 0 | 10.1.1.1 | loc1 | 321 | 1 | |
| 10.1.1.255/32 | dest | 0 | 10.1.1.255 | bcst | 323 | 1 | so-5/3/0.0 |
| 10.21.21.0/24 | intf | 0 | ff.3.0.21 | ucst | 326 | 1 | so-5/3/0.0 |
| 10.21.21.0/32 | dest | 0 | 10.21.21.0 | recv | 328 | 1 | so-5/3/0.0 |
| 10.21.21.1/32 | intf | 0 | 10.21.21.1 | loc1 | 325 | 1 | |
| 10.21.21.255/32 | dest | 0 | 10.21.21.255 | bcst | 327 | 1 | so-5/3/0.0 |
| 127.0.0.1/32 | intf | 0 | 127.0.0.1 | loc1 | 320 | 1 | |
| 172.17.28.19/32 | clon | 1 | 192.168.4.254 | ucst | 132 | 4 | fxp0.0 |
| 172.17.28.44/32 | clon | 1 | 192.168.4.254 | ucst | 132 | 4 | fxp0.0 |

```
...
```

```
Routing table: private1__inet
```

```
Internet:
```

| Destination | Type | RtRef | Next hop | Type | Index | NhRef | Netif |
|-------------|------|-------|----------|------|-------|-------|--------|
| default | perm | 0 | | rjct | 46 | 1 | |
| 10.0.0.0/8 | intf | 0 | | rslv | 136 | 1 | fxp1.0 |
| 10.0.0.0/32 | dest | 0 | 10.0.0.0 | recv | 134 | 1 | fxp1.0 |
| 10.0.0.4/32 | intf | 0 | 10.0.0.4 | loc1 | 135 | 2 | |
| 10.0.0.4/32 | dest | 0 | 10.0.0.4 | loc1 | 135 | 2 | |

```
...
```

```
Routing table: iso
```

```
ISO:
```

| Destination | Type | RtRef | Next hop | Type | Index | NhRef | Netif |
|-------------|------|-------|----------|------|-------|-------|-------|
| default | perm | 0 | | rjct | 38 | 1 | |

```
Routing table: inet6
```

```
Internet6:
```

| Destination | Type | RtRef | Next hop | Type | Index | NhRef | Netif |
|-------------|------|-------|----------|------|-------|-------|-------|
| default | perm | 0 | | rjct | 22 | 1 | |
| ff00::/8 | perm | 0 | | mdsc | 21 | 1 | |
| ff02::1/128 | perm | 0 | ff02::1 | mcst | 17 | 1 | |

```
...
```

```
Routing table: mpls
```

```
MPLS:
```

| Destination | Type | RtRef | Next hop | Type | Index | NhRef | Netif |
|-------------|------|-------|----------|------|-------|-------|-------|
| default | perm | 0 | | rjct | 28 | 1 | |

show route forwarding-table destination extensive

```
user@host> show route forwarding-table destination 3.4.2.1 extensive
```

```
Routing table: inet [Index 0]
```

```
Internet:
```


(Weights and Balances)

```

Destination: 3.4.2.1/32
Route type: user
Route reference: 0
Flags: sent to PFE
Next-hop type: unicast
Nexthop: 4.4.4.4
Index: 262143 Reference: 1
Next-hop type: unicast
Next-hop interface: so-1/1/0.0
Index: 335 Reference: 2
Weight: 22 Balance: 3
Nexthop: 145.12.1.2
Next-hop type: unicast
Index: 337 Reference: 2
Next-hop interface: so-0/1/2.0
Weight: 33 Balance: 33

```

show route forwarding-table extensive

```

user@host> show route forwarding-table extensive
Routing table: inet [Index 0]
Internet:

Destination: default
Route type: user
Route reference: 2
Flags: sent to PFE
Nexthop: 0:90:69:8e:b1:1b
Index: 132 Reference: 4
Next-hop type: unicast
Next-hop interface: fxp0.0

Destination: default
Route type: permanent
Route reference: 0
Flags: none
Next-hop type: reject
Index: 14 Reference: 1

Destination: 127.0.0.1/32
Route type: interface
Route reference: 0
Flags: sent to PFE
Nexthop: 127.0.0.1
Index: 320 Reference: 1
Next-hop type: local

...

Routing table: private1__inet [Index 1]
Internet:

Destination: default
Route type: permanent
Route reference: 0
Flags: sent to PFE
Next-hop type: reject
Index: 46 Reference: 1

Destination: 10.0.0.0/8
Route type: interface
Route reference: 0
Flags: sent to PFE
Next-hop type: resolve
Index: 136 Reference: 1
Next-hop interface: fxp1.0

...

Routing table: iso [Index 0]
ISO:

Destination: default

```

```

Route type: permanent
Route reference: 0
Flags: sent to PFE
Next-hop type: reject
Route interface-index: 0
Index: 38      Reference: 1

Routing table: inet6 [Index 0]
Internet6:

Destination: default
Route type: permanent
Route reference: 0
Flags: sent to PFE
Next-hop type: reject
Route interface-index: 0
Index: 22      Reference: 1

Destination: ff00::/8
Route type: permanent
Route reference: 0
Flags: sent to PFE
Next-hop type: multicast discard
Route interface-index: 0
Index: 21      Reference: 1

...

Routing table: private1__inet6 [Index 1]
Internet6:

Destination: default
Route type: permanent
Route reference: 0
Flags: sent to PFE
Next-hop type: reject
Route interface-index: 0
Index: 54      Reference: 1

Destination: fe80::2a0:a5ff:fe3d:375/128
Route type: interface
Route reference: 0
Flags: sent to PFE
Nexthop: fe80::2a0:a5ff:fe3d:375
Next-hop type: local
Route interface-index: 0
Index: 75      Reference: 1

...

```

show route forwarding-table extensive (RPF)

The next example is based on the following configuration, which enables an RPF check on all routes that are learned from this interface, including the interface route:

```

so-1/1/0 {
  unit 0 {
    family inet {
      rpf-check;
      address 15.95.1.2/30;
    }
  }
}

```

```

user@host> show route forwarding-table extensive
Routing table: inet [Index 0]
Internet:
...
...
Destination: 15.95.1.3/32
Route type: destination
Route reference: 0
Route interface-index: 67

```

Flags: sent to PFE
Nexthop: 15.95.1.3
Next-hop type: broadcast Index: 328 Reference: 1
Next-hop interface: so-1/1/0.0
RPF interface: so-1/1/0.0

show route forwarding-table family mpls

```

user@host> show route forwarding-table family mpls
Routing table: mpls
MPLS:
Destination      Type RtRef Next hop      Type Index NhRef Netif
default          perm  0
0                user  0
1                user  0
2                user  0
100000           user  0 10.31.1.6      swap 100001      fe-1/1/0.0
800002           user  0                Pop                vt-0/3/0.32770

vt-0/3/0.32770 (VPLS)
                    user  0                indr  351      4
                    Push 800000, Push 100002(top)

so-0/0/0.0

```

show route forwarding-table family vpls

```

user@host> show route forwarding-table family vpls
Routing table: green.vpls
VPLS:
Destination      Type RtRef Next hop      Type Index NhRef Netif
default          dynm  0
default          perm  0
fe-0/1/0.0       dynm  0
00:90:69:0c:20:1f/48      <<<<<Remote CE

                    dynm  0                indr  351      4
                    Push 800000, Push 100002(top)

so-0/0/0.0
00:90:69:85:b0:1f/48      <<<<<Local CE

                    dynm  0                ucst  354      2 fe-0/1/0.0

```

show route forwarding-table family vpls extensive

```

user@host> show route forwarding-table family vpls extensive
Routing table: green.vpls [Index 2]
VPLS:

Destination: default
Route type: dynamic
Route reference: 0
Flags: sent to PFE
Next-hop type: flood
Next-hop type: unicast
Next-hop interface: fe-0/1/3.0
Next-hop type: unicast
Next-hop interface: fe-0/1/2.0
Route interface-index: 72
Index: 289 Reference: 1
Index: 291 Reference: 3
Index: 290 Reference: 3

Destination: default
Route type: permanent
Route reference: 0
Flags: none
Next-hop type: discard
Route interface-index: 0
Index: 341 Reference: 1

Destination: fe-0/1/2.0
Route type: dynamic
Route reference: 0
Flags: sent to PFE
Next-hop type: flood
Next-hop type: indirect
Next-hop type: Push 800016
Route interface-index: 69
Index: 293 Reference: 1
Index: 363 Reference: 4

```

```

Next-hop interface: at-1/0/1.0
Next-hop type: indirect          Index: 301      Reference: 5
Next hop: 10.31.3.2
Next-hop type: Push 800000
Next-hop interface: fe-0/1/1.0
Next-hop type: unicast          Index: 291      Reference: 3
Next-hop interface: fe-0/1/3.0

Destination: fe-0/1/3.0
Route type: dynamic
Route reference: 0               Route interface-index: 70
Flags: sent to PFE
Next-hop type: flood            Index: 292      Reference: 1
Next-hop type: indirect         Index: 363      Reference: 4
Next-hop type: Push 800016
Next-hop interface: at-1/0/1.0
Next-hop type: indirect         Index: 301      Reference: 5
Next hop: 10.31.3.2
Next-hop type: Push 800000
Next-hop interface: fe-0/1/1.0
Next-hop type: unicast          Index: 290      Reference: 3
Next-hop interface: fe-0/1/2.0

Destination: 10:00:00:01:01:01/48
Route type: dynamic
Route reference: 0               Route interface-index: 70
Flags: sent to PFE, prefix load balance
Next-hop type: unicast          Index: 291      Reference: 3
Next-hop interface: fe-0/1/3.0
Route used as destination:
  Packet count:      6640   Byte count:      675786
Route used as source
  Packet count:      6894   Byte count:      696424

Destination: 10:00:00:01:01:01:04/48
Route type: dynamic
Route reference: 0               Route interface-index: 69
Flags: sent to PFE, prefix load balance
Next-hop type: unicast          Index: 290      Reference: 3
Next-hop interface: fe-0/1/2.0
Route used as destination:
  Packet count:      96     Byte count:      8079
Route used as source:
  Packet count:      296    Byte count:      24955

Destination: 10:00:00:01:03:05/48
Route type: dynamic
Route reference: 0               Route interface-index: 74
Flags: sent to PFE, prefix load balance
Next-hop type: indirect         Index: 301      Reference: 5
Next hop: 10.31.3.2
Next-hop type: Push 800000
Next-hop interface: fe-0/1/1.0

```

show route forwarding-table table default

```

user@host> show route forwarding-table table default
Routing table: default.inet

```

```

Internet:
Destination      Type RtRef Next hop      Type Index NhRef Netif
default          perm  0
0.0.0.0/32       perm  0
10.0.60.0/30     user  0 10.0.60.13  ucst  713  5 fe-0/1/3.0

```

```

10.0.60.12/30      intf      0                rslv    688      1 fe-0/1/3.0
10.0.60.12/32      dest      0 10.0.60.12      recv    686      1 fe-0/1/3.0
10.0.60.13/32      dest      0 0:5:85:8b:bc:22 ucst    713      5 fe-0/1/3.0
10.0.60.14/32      intf      0 10.0.60.14      locl    687      2
10.0.60.14/32      dest      0 10.0.60.14      locl    687      2
10.0.60.15/32      dest      0 10.0.60.15      bcst    685      1 fe-0/1/3.0
10.0.67.12/30      user      0 10.0.60.13      ucst    713      5 fe-0/1/3.0
10.0.80.0/30       ifdn      0 ff.3.0.21       ucst    676      1 so-0/0/1.0
10.0.80.0/32       dest      0 10.0.80.0       recv    678      1 so-0/0/1.0
10.0.80.2/32       user      0                rjct     36      2
10.0.80.2/32       intf      0 10.0.80.2       locl    675      1
10.0.80.3/32       dest      0 10.0.80.3       bcst    677      1 so-0/0/1.0
10.0.90.12/30      intf      0                rslv    684      1 fe-0/1/0.0
10.0.90.12/32      dest      0 10.0.90.12      recv    682      1 fe-0/1/0.0
10.0.90.14/32      intf      0 10.0.90.14      locl    683      2
10.0.90.14/32      dest      0 10.0.90.14      locl    683      2
10.0.90.15/32      dest      0 10.0.90.15      bcst    681      1 fe-0/1/0.0
10.5.0.0/16        user      0 192.168.187.126 ucst    324     15 fxp0.0
10.10.0.0/16        user      0 192.168.187.126 ucst    324     15 fxp0.0
10.13.10.0/23       user      0 192.168.187.126 ucst    324     15 fxp0.0
10.84.0.0/16        user      0 192.168.187.126 ucst    324     15 fxp0.0
10.150.0.0/16       user      0 192.168.187.126 ucst    324     15 fxp0.0
10.157.64.0/19      user      0 192.168.187.126 ucst    324     15 fxp0.0
10.209.0.0/16       user      0 192.168.187.126 ucst    324     15 fxp0.0

```

...

Routing table: default.iso

ISO:

| Destination | Type | RtRef | Next hop | Type | Index | NhRef | Netif |
|-------------|------|-------|----------|------|-------|-------|-------|
| default | perm | 0 | | rjct | 60 | 1 | |

Routing table: default.inet6

Internet6:

| Destination | Type | RtRef | Next hop | Type | Index | NhRef | Netif |
|-------------|------|-------|----------|------|-------|-------|-------|
| default | perm | 0 | | rjct | 44 | 1 | |
| ::/128 | perm | 0 | | dscd | 42 | 1 | |
| ff00::/8 | perm | 0 | | mdsc | 43 | 1 | |
| ff02::1/128 | perm | 0 | ff02::1 | mcst | 39 | 1 | |

Routing table: default.mpls

MPLS:

| Destination | Type | RtRef | Next hop | Type | Index | NhRef | Netif |
|-------------|------|-------|----------|------|-------|-------|-------|
| default | perm | 0 | | dscd | 50 | 1 | |

show route
forwarding-table table
logical-system-name

user@host> show route forwarding-table table R4/vpn-red

Logical system: R4

Routing table: vpn-red.inet

Internet:

| Destination | Type | RtRef | Next hop | Type | Index | NhRef | Netif |
|--------------|------|-------|--|------|-------|-------|------------|
| default | perm | 0 | | rjct | 563 | 1 | |
| 0.0.0.0/32 | perm | 0 | | dscd | 561 | 2 | |
| 1.0.0.1/32 | user | 0 | | dscd | 561 | 2 | |
| 2.0.2.0/24 | intf | 0 | | rslv | 771 | 1 | ge-1/2/0.3 |
| 2.0.2.0/32 | dest | 0 | 2.0.2.0 | recv | 769 | 1 | ge-1/2/0.3 |
| 2.0.2.1/32 | intf | 0 | 2.0.2.1 | locl | 770 | 2 | |
| 2.0.2.1/32 | dest | 0 | 2.0.2.1 | locl | 770 | 2 | |
| 2.0.2.2/32 | dest | 0 | 0.4.80.3.0.1b.c0.d5.e4.bd.0.1b.c0.d5.e4.bc.8.0 | ucst | 789 | 1 | ge-1/2/0.3 |
| 2.0.2.255/32 | dest | 0 | 2.0.2.255 | bcst | 768 | 1 | ge-1/2/0.3 |
| 224.0.0.0/4 | perm | 1 | | mdsc | 562 | 1 | |

```

224.0.0.1/32      perm      0 224.0.0.1      mcst    558      1
255.255.255.255/32 perm      0                bcst    559      1

```

```

Logical system: R4
Routing table: vpn-red.iso
ISO:

```

| Destination | Type | RtRef | Next hop | Type | Index | NhRef | Netif |
|-------------|------|-------|----------|------|-------|-------|-------|
| default | perm | 0 | | rjct | 608 | 1 | |

```

Logical system: R4
Routing table: vpn-red.inet6
Internet6:

```

| Destination | Type | RtRef | Next hop | Type | Index | NhRef | Netif |
|-------------|------|-------|----------|------|-------|-------|-------|
| default | perm | 0 | | rjct | 708 | 1 | |
| ::/128 | perm | 0 | | dscd | 706 | 1 | |
| ff00::/8 | perm | 0 | | mdsc | 707 | 1 | |
| ff02::1/128 | perm | 0 | ff02::1 | mcst | 704 | 1 | |

```

Logical system: R4
Routing table: vpn-red.mpls
MPLS:

```

| Destination | Type | RtRef | Next hop | Type | Index | NhRef | Netif |
|-------------|------|-------|----------|------|-------|-------|-------|
| default | perm | 0 | | dscd | 638 | | |

**show route
forwarding-table vpn**

```

user@host> show route forwarding-table vpn VPN-A
Routing table:: VPN-A.inet

```

```

Internet:
Destination      Type RtRef Nexthop      Type Index NhRef Netif
default          perm  0                rjct   4     4
10.39.10.20/30   intf  0 ff.3.0.21      ucst   40    1
so-0/0/0.0
10.39.10.21/32   intf  0 10.39.10.21    locl   36    1
10.255.14.172/32 user  0                ucst   69    2
so-0/0/0.0
10.255.14.175/32 user  0                indr   81    3
Push 100004, Push
100004(top) so-1/0/0.0
224.0.0.0/4      perm  2                mdsc   5     3
224.0.0.1/32     perm  0 224.0.0.1      mcst   1     8
224.0.0.5/32     user  1 224.0.0.5      mcst   1     8
255.255.255.255/32 perm  0                bcst   2     3

```

show route forwarding-table interface-name

| | |
|---------------------------------|--|
| Syntax | show route forwarding-table interface-name <i>interface-name</i> <detail extensive> <all> |
| Release Information | Command introduced in Junos OS Release 9.6. |
| Description | Display the interfaces in the Routing Engine's forwarding table. |
| Options | <p>none—Display information for the specified interface.</p> <p>detail extensive—(Optional) Display the specified level of output.</p> <p>all—(Optional) Display all interfaces in the routing table.</p> |
| Required Privilege Level | view |
| List of Sample Output | show route forwarding-table interface-name fe-0/1/1 on page 575 show route forwarding-table interface-name all on page 576 show route forwarding-table interface-name all detail on page 576 |
| Output Fields | Table 136 on page 574 lists the output fields for the show route forwarding-table interface-name command. Output fields are listed in the approximate order in which they appear. |

Table 136: show route forwarding-table interface-name Output Fields

| Field Name | Field Description | Level of Output |
|----------------|---|-------------------------|
| Name | Name of the interface (for example fe-0/1/1 , lo0 , ae0 , and so on). | All levels |
| MTU | Interface's maximum transmission unit (MTU). | All levels |
| Afam | Configured address family (for example inet , tnp , inet6 , and so on). | detail extensive |
| Network | Network information: <ul style="list-style-type: none"> • <Link>—Physical interface, not a logical interface. • <PtoP>—Point-to-point network. • ipaddress—Network address. | All levels |
| Address | Address of the interface. The address can be a MAC address, IPv4 address, IPv6 address, and so on. | All levels |
| IPkts | Number of packets received on the interface. | All levels |
| Ierr | Number of packets received on the interface with errors. | All levels |
| Opkts | Number of packets transmitted or sent from the interface. | All levels |

Table 136: show route forwarding-table interface-name Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|------------|---|-----------------|
| Oerr | Number of packets transmitted or sent from the interface with errors. | All levels |
| Coll | Number of packets that experienced collisions on the interface. | All levels |

Sample Output

show route
forwarding-table

```
user@host> show route forwarding-table interface-name fe-0/1/1
Name      Mtu Network      Address      Ipkts Ierr    Opkts Oerr    Coll
fe-0/1/1  1514 <Link>      00.05.85.88.cc.20  0    0        0    0        0
```

interface-name
fe-0/1/1

show route
forwarding-table
interface-name all

```
user@host> show route forwarding-table interface-name all
Name      Mtu Network      Address      IpKts Ierr      Opkts Oerr      Coll
fxp0      1514 <Link>      00.a0.a5.56.03.83  180965  0      39907  0      0

  unit 0      1500 192.168.187.0/ 192.168.187.10
fxp1      1514 <Link>      02.00.00.00.00.04  33010497  0 30110800  0      0

  unit 0      1500 10.0.0.0/8      10.0.0.1
           10.0.0.0/8      10.0.0.4
           128.0.0.0/2      128.0.0.1
           128.0.0.0/2      128.0.0.4
           1500 fe80::/64      fe80::200:ff:fe0
           fec0::/64      fec0::a:0:0:4
           1500      4
lsi       1496 <Link>
dsc       max <Link>      0 0      0 0      0
lo0       max <Link>      8980 0      8980 0      0

  unit 0      max 127.0.0.1/8      127.0.0.1
           192.168.0.1/8 192.168.0.1
unit 16384 max 127.0.0.1/8      127.0.0.1
unit 16385 max
gre       max <Link>
ipip      max <Link>
tap       max <Link>
pime      max <Link>
pimd      max <Link>
mtun      max <Link>
so-0/0/0  4474 <Link>      1679900 0 1068611 0      0

  unit 0      4470 <PtoP>      10.0.60.2      0 0      0 0
0
so-0/0/1  4474 <Link>      0 0      0 0      0

  unit 0      4470 <PtoP>      10.0.80.2      0 0      0 0
0
so-0/0/2  4474 <Link>      0 0      0 0      0
so-0/0/3  4474 <Link>      0 0      0 0      0
fe-0/1/0  1514 <Link>      00.05.85.88.cc.1f  523120 0 623044 0      0

  unit 0      1500 10.0.90.12/30 10.0.90.14      0 0      0 0
0
fe-0/1/1  1514 <Link>      00.05.85.88.cc.20  0 0      0 0      0
fe-0/1/2  1514 <Link>      00.05.85.88.cc.21  0 0      0 0      0

...
```

show route
forwarding-table

```
user@host> show route forwarding-table interface-name all detail
Name      Mtu AFam      Network      Address      IpKts Ierr      Opkts
Oerr      Coll
```

interface-name all
detail

```

fxp0      1514      <Link>      00.a0.a5.56.03.83      181005      0      39948
0      0
unit 0      1500 inet      192.168.187.0/ 192.168.187.10
fxp1      1514      <Link>      02.00.00.00.00.04      33012676      0      30112468
0      0
unit 0      1500 inet      10.0.0.0/8      10.0.0.1
                                10.0.0.0/8      10.0.0.4
                                128.0.0.0/2      128.0.0.1
                                128.0.0.0/2      128.0.0.4
                                1500 inet6      fe80::/64      fe80::200:ff:fe0
                                fec0::/64      fec0::a:0:0:4
                                1500 tnp
                                4
lsi      1496      <Link>
dsc      max      <Link>
0      0
lo0      max      <Link>
0      0
unit 0      max inet      127.0.0.1/8      127.0.0.1
                                192.168.0.1/8      192.168.0.1
unit 16384 max inet      127.0.0.1/8      127.0.0.1
unit 16385 max inet
gre      max      <Link>
ipip     max      <Link>
tap      max      <Link>
pime     max      <Link>
pimd     max      <Link>
mtun     max      <Link>
so-0/0/0 4474      <Link>
0      0
unit 0      4470 inet      <PtoP>      10.0.60.2
0      0
...

```

show route hidden

| | |
|---------------------------------|--|
| Syntax | <code>show route hidden</code> <code><brief detail extensive terse></code> <code><logical-system (all <i>logical-system-name</i>)></code> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display only hidden route information. A hidden route is unusable, even if it is the best path. |
| Options | brief detail extensive terse —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief . logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | view |
| List of Sample Output | show route hidden on page 579 show route hidden detail on page 579 show route hidden extensive on page 580 show route hidden terse on page 580 |
| Output Fields | For information about output fields, see the output field table for the show route command, the show route detail command, the show route extensive command, or the show route terse command. |

Sample Output

show route hidden

```

user@host> show route hidden
inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
127.0.0.1/32      [Direct/0] 04:26:38
                  > via lo0.0

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.5.5.5/32      [BGP/170] 03:44:10, localpref 100, from 10.4.4.4
                  AS path: 100 I
                  Unusable
10.12.1.0/24     [BGP/170] 03:44:10, localpref 100, from 10.4.4.4
                  AS path: 100 I
                  Unusable
10.12.80.4/30    [BGP/170] 03:44:10, localpref 100, from 10.4.4.4
                  AS path: I
                  Unusable
...

```

show route hidden detail

```

user@host> show route hidden detail

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
127.0.0.1/32 (1 entry, 0 announced)
    Direct Preference: 0
        Next hop type: Interface
        Next-hop reference count: 1
        Next hop: via lo0.0, selected
        State: <Hidden Martian Int>
        Local AS: 1
        Age: 4:27:37
        Task: IF
        AS path: I

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
10.5.5.5/32 (1 entry, 0 announced)
    BGP Preference: 170/-101
        Route Distinguisher: 10.4.4.4:4
        Next hop type: Unusable
        Next-hop reference count: 6
        State: <Secondary Hidden Int Ext>
        Local AS: 1 Peer AS: 1
        Age: 3:45:09
        Task: BGP_1.10.4.4.4+2493
        AS path: 100 I
        Communities: target:1:999
        VPN Label: 100064
        Localpref: 100
        Router ID: 10.4.4.4

```

Primary Routing Table bgp.13vpn.0

...

show route hidden extensive

The output for the **show route hidden extensive** command is identical to that of the **show route hidden detail** command. For sample output, see [show route hidden detail on page 579](#).

show route hidden terse

```
user@host> show route hidden terse
```

```
inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1   Metric 2   Next hop      AS path
  127.0.0.1/32      D   0           >100.0

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1   Metric 2   Next hop      AS path
  10.5.5.5/32       B 170      100        Unusable     100 I
  10.12.1.0/24      B 170      100        Unusable     100 I
  10.12.80.4/30     B 170      100        Unusable      I

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

bgp.13vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1   Metric 2   Next hop      AS path
  10.4.4.4:4:10.5.5.5/32
                        B 170      100        Unusable     100 I
  10.4.4.4:4:10.12.1.0/24
                        B 170      100        Unusable     100 I
  10.4.4.4:4:10.12.80.4/30
                        B 170      100        Unusable      I

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route inactive-path

| | |
|------------------------------------|--|
| Syntax | show route inactive-path <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switches) | show route inactive-path <brief detail extensive terse> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. |
| Description | Display routes for destinations that have no active route. An inactive route is a route that was not selected as the best path. |
| Options | <p>none—Display all inactive routes.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show route inactive-path on page 582 show route inactive-path detail on page 582 show route inactive-path extensive on page 583 show route inactive-path terse on page 583 |
| Output Fields | For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command. |

Sample Output

show route inactive-path

```

user@host> show route inactive-path

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.12.100.12/30      [OSPF/10] 03:57:28, metric 1
> via so-0/3/0.0

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.0/8          [Direct/0] 04:39:56
> via fxp1.0

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.12.80.0/30       [BGP/170] 04:38:17, localpref 100
AS path: 100 I
> to 10.12.80.1 via ge-6/3/2.0

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

bgp.l3vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

show route inactive-path detail

```

user@host> show route inactive-path detail

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete

10.12.100.12/30 (2 entries, 1 announced)
  OSPF   Preference: 10
        Next-hop reference count: 1
        Next hop: via so-0/3/0.0, selected
        State: <Int>
        Inactive reason: Route Preference
        Local AS:      1
        Age: 3:58:24    Metric: 1
        Area: 0.0.0.0
        Task: OSPF
        AS path: I

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

10.0.0.0/8 (2 entries, 0 announced)

```



```

Direct Preference: 0
  Next hop type: Interface
  Next-hop reference count: 1
  Next hop: via fxp1.0, selected
  State: <NotBest Int>
  Inactive reason: No difference
  Age: 4:40:52
  Task: IF
  AS path: I

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete

10.12.80.0/30 (2 entries, 1 announced)
  BGP Preference: 170/-101
  Next-hop reference count: 6
  Source: 10.12.80.1
  Next hop: 10.12.80.1 via ge-6/3/2.0, selected
  State: <Ext>
  Inactive reason: Route Preference
  Peer AS: 100
  Age: 4:39:13
  Task: BGP_100.10.12.80.1+179
  AS path: 100 I
  Localpref: 100
  Router ID: 10.0.0.0

```

**show route
inactive-path
extensive**

The output for the **show route inactive-path extensive** command is identical to that of the **show route inactive-path detail** command. For sample output, see [show route inactive-path detail on page 582](#).

**show route
inactive-path terse**

```

user@host> show route inactive-path terse

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
  10.12.100.12/30   0  10           1           >so-0/3/0.0

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
  10.0.0.0/8        D   0           0           >fxp1.0

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
  10.12.80.0/30     B 170          100           >10.12.80.1    100 I

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

```

```
bgp.l3vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)
Restart Complete
```

```
inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
```

```
private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route inactive-prefix

| | |
|------------------------------------|--|
| Syntax | show route inactive-prefix <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switches) | show route inactive-prefix <brief detail extensive terse> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. |
| Description | Display inactive route destinations in each routing table. |
| Options | <p>none—Display all inactive route destination.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show route inactive-prefix on page 586 show route inactive-prefix detail on page 586 show route inactive-prefix extensive on page 586 show route inactive-prefix terse on page 586 |
| Output Fields | For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command. |

Sample Output

show route inactive-prefix

```
user@host> show route inactive-prefix

inet.0: 14 destinations, 14 routes (13 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

127.0.0.1/32          [Direct/0] 00:04:54
> via lo0.0
```

show route inactive-prefix detail

```
user@host> show route inactive-prefix detail

inet.0: 14 destinations, 14 routes (13 active, 0 holddown, 1 hidden)
127.0.0.1/32 (1 entry, 0 announced)
  Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
    State: <Hidden Martian Int>
    Age: 4:51
    Task: IF
    AS path: I00:04:54
    > via lo0.0
```

show route inactive-prefix extensive

The output for the **show route inactive-prefix extensive** command is identical to that of the **show route inactive-path detail** command. For sample output, see [show route inactive-prefix detail on page 586](#).

show route inactive-prefix terse

```
user@host> show route inactive-prefix terse

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination          P Prf  Metric 1  Metric 2  Next hop          AS path
127.0.0.1/32          D   0                >lo0.0
```

show route instance

| | |
|---|---|
| Syntax | show route instance <brief detail summary> <instance-name> <logical-system (all <i>logical-system-name</i>)> <operational> |
| Syntax (EX Series Switches and QFX Series) | show route instance <brief detail summary> <instance-name> <operational> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | Display routing instance information. |
| Options | <p>none—(Same as brief) Display standard information about all routing instances.</p> <p>brief detail summary—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief. (These options are not available with the operational keyword.)</p> <p>instance-name—(Optional) Display information for all routing instances whose name begins with this string (for example, cust1, cust11, and cust111 are all displayed when you run the show route instance cust1 command).</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>operational—(Optional) Display operational routing instances.</p> |
| Required Privilege Level | view |
| List of Sample Output | show route instance on page 589 show route instance detail (Graceful Restart Complete) on page 589 show route instance detail (Graceful Restart Incomplete) on page 591 show route instance detail (VPLS Routing Instance) on page 592 show route instance operational on page 593 show route instance summary on page 593 |
| Output Fields | Table 137 on page 587 lists the output fields for the show route instance command. Output fields are listed in the approximate order in which they appear. |

Table 137: show route instance Output Fields

| Field Name | Field Description | Level of Output |
|----------------------------------|-------------------------------|-----------------|
| Instance or <i>instance-name</i> | Name of the routing instance. | All levels |

Table 137: show route instance Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------------------------|--|---------------------------|
| Operational Routing Instances | (operational keyword only) Names of all operational routing instances. | — |
| Type | Type of routing instance: forwarding , l2vpn , no-forwarding , vpls , virtual-router , or vrf . | All levels |
| State | State of the routing instance: active or inactive . | brief detail none |
| Interfaces | Name of interfaces belonging to this routing instance. | brief detail none |
| Restart State | Status of graceful restart for this instance: Pending or Complete . | detail |
| Path selection timeout | Maximum amount of time, in seconds, remaining until graceful restart is declared complete. The default is 300 . | detail |
| Tables | Tables (and number of routes) associated with this routing instance. | brief detail none |
| Route-distinguisher | Unique route distinguisher associated with this routing instance. | detail |
| Vrf-import | VPN routing and forwarding instance import policy name. | detail |
| Vrf-export | VPN routing and forwarding instance export policy name. | detail |
| Vrf-import-target | VPN routing and forwarding instance import target community name. | detail |
| Vrf-export-target | VPN routing and forwarding instance export target community name. | detail |
| Fast-reroute-priority | Fast reroute priority setting for a VPLS routing instance: high , medium , or low . The default is low . | detail |
| Restart State | Restart state: <ul style="list-style-type: none"> • Pending:protocol-name—List of protocols that have not yet completed graceful restart for this routing table. • Complete—All protocols have restarted for this routing table. | detail |
| Primary rib | Primary table for this routing instance. | brief none summary |
| Active/holddown/hidden | Number of active, hold-down, and hidden routes. | All levels |

Sample Output

show route instance

```

user@host> show route instance
Instance              Type
      Primary RIB
master                forwarding
      inet.0              16/0/1
      iso.0               1/0/0
      mpls.0              0/0/0
      inet6.0             2/0/0
      l2circuit.0         0/0/0
__juniper_private1__ forwarding
  __juniper_private1__.inet.0 12/0/0
  __juniper_private1__.inet6.0 1/0/0

```

show route instance detail (Graceful Restart Complete)

```

user@host> show route instance detail
master:
  Router ID: 10.255.14.176
  Type: forwarding      State: Active
  Restart State: Complete Path selection timeout: 300
  Tables:
    inet.0              : 17 routes (15 active, 0 holddown, 1 hidden)
    Restart Complete
    inet.3              : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    iso.0               : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
    mpls.0              : 19 routes (19 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l3vpn.0         : 10 routes (10 active, 0 holddown, 0 hidden)
    Restart Complete
    inet6.0             : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l2vpn.0         : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
  BGP-INET:
    Router ID: 10.69.103.1
    Type: vrf            State: Active
    Restart State: Complete Path selection timeout: 300
    Interfaces:
      t3-0/0/0.103
    Route-distinguisher: 10.255.14.176:103
    Vrf-import: [ BGP-INET-import ]
    Vrf-export: [ BGP-INET-export ]
    Tables:
      BGP-INET.inet.0    : 4 routes (4 active, 0 holddown, 0 hidden)
      Restart Complete
  BGP-L:
    Router ID: 10.69.104.1
    Type: vrf            State: Active
    Restart State: Complete Path selection timeout: 300
    Interfaces:
      t3-0/0/0.104
    Route-distinguisher: 10.255.14.176:104
    Vrf-import: [ BGP-L-import ]
    Vrf-export: [ BGP-L-export ]
    Tables:
      BGP-L.inet.0       : 4 routes (4 active, 0 holddown, 0 hidden)
      Restart Complete

```

```
BGP-L.mpls.0          : 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete
L2VPN:
Router ID: 0.0.0.0
Type: l2vpn           State: Active
Restart State: Complete Path selection timeout: 300
Interfaces:
  t3-0/0/0.512
Route-distinguisher: 10.255.14.176:512
Vrf-import: [ L2VPN-import ]
Vrf-export: [ L2VPN-export ]
Tables:
  L2VPN.l2vpn.0       : 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
LDP:
Router ID: 10.69.105.1
Type: vrf             State: Active
Restart State: Complete Path selection timeout: 300
Interfaces:
  t3-0/0/0.105
Route-distinguisher: 10.255.14.176:105
Vrf-import: [ LDP-import ]
Vrf-export: [ LDP-export ]
Tables:
  LDP.inet.0          : 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
OSPF:
Router ID: 10.69.101.1
Type: vrf             State: Active
Restart State: Complete Path selection timeout: 300
Interfaces:
  t3-0/0/0.101
Route-distinguisher: 10.255.14.176:101
Vrf-import: [ OSPF-import ]
Vrf-export: [ OSPF-export ]
Vrf-import-target: [ target:11111
Tables:
  OSPF.inet.0         : 8 routes (7 active, 0 holddown, 0 hidden)
Restart Complete
RIP:
Router ID: 10.69.102.1
Type: vrf             State: Active
Restart State: Complete Path selection timeout: 300
Interfaces:
  t3-0/0/0.102
Route-distinguisher: 10.255.14.176:102
Vrf-import: [ RIP-import ]
Vrf-export: [ RIP-export ]
Tables:
  RIP.inet.0          : 6 routes (6 active, 0 holddown, 0 hidden)
Restart Complete
STATIC:
Router ID: 10.69.100.1
Type: vrf             State: Active
Restart State: Complete Path selection timeout: 300
Interfaces:
  t3-0/0/0.100
Route-distinguisher: 10.255.14.176:100
Vrf-import: [ STATIC-import ]
Vrf-export: [ STATIC-export ]
Tables:
```



```

STATIC.inet.0          : 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

```

**show route instance
detail (Graceful
Restart Incomplete)**

```

user@host> show route instance detail
master:
  Router ID: 10.255.14.176
  Type: forwarding      State: Active
  Restart State: Pending Path selection timeout: 300
  Tables:
    inet.0              : 17 routes (15 active, 1 holddown, 1 hidden)
    Restart Pending: OSPF LDP
    inet.3              : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Pending: OSPF LDP
    iso.0               : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
    mpls.0              : 23 routes (23 active, 0 holddown, 0 hidden)
    Restart Pending: LDP VPN
    bgp.l3vpn.0         : 10 routes (10 active, 0 holddown, 0 hidden)
    Restart Pending: BGP VPN
    inet6.0             : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l2vpn.0         : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Pending: BGP VPN
BGP-INET:
  Router ID: 10.69.103.1
  Type: vrf             State: Active
  Restart State: Pending Path selection timeout: 300
  Interfaces:
    t3-0/0/0.103
  Route-distinguisher: 10.255.14.176:103
  Vrf-import: [ BGP-INET-import ]
  Vrf-export: [ BGP-INET-export ]
  Tables:
    BGP-INET.inet.0     : 6 routes (5 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
BGP-L:
  Router ID: 10.69.104.1
  Type: vrf             State: Active
  Restart State: Pending Path selection timeout: 300
  Interfaces:
    t3-0/0/0.104
  Route-distinguisher: 10.255.14.176:104
  Vrf-import: [ BGP-L-import ]
  Vrf-export: [ BGP-L-export ]
  Tables:
    BGP-L.inet.0        : 6 routes (5 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
    BGP-L.mpls.0        : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
L2VPN:
  Router ID: 0.0.0.0
  Type: l2vpn           State: Active
  Restart State: Pending Path selection timeout: 300
  Interfaces:
    t3-0/0/0.512
  Route-distinguisher: 10.255.14.176:512
  Vrf-import: [ L2VPN-import ]
  Vrf-export: [ L2VPN-export ]
  Tables:
    L2VPN.l2vpn.0       : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Pending: VPN L2VPN

```

```

LDP:
  Router ID: 10.69.105.1
  Type: vrf                      State: Active
  Restart State: Pending Path selection timeout: 300
  Interfaces:
    t3-0/0/0.105
  Route-distinguisher: 10.255.14.176:105
  Vrf-import: [ LDP-import ]
  Vrf-export: [ LDP-export ]
  Tables:
    LDP.inet.0                  : 5 routes (4 active, 1 holddown, 0 hidden)
    Restart Pending: OSPF LDP VPN

OSPF:
  Router ID: 10.69.101.1
  Type: vrf                      State: Active
  Restart State: Pending Path selection timeout: 300
  Interfaces:
    t3-0/0/0.101
  Route-distinguisher: 10.255.14.176:101
  Vrf-import: [ OSPF-import ]
  Vrf-export: [ OSPF-export ]
  Tables:
    OSPF.inet.0                : 8 routes (7 active, 1 holddown, 0 hidden)
    Restart Pending: OSPF VPN

RIP:
  Router ID: 10.69.102.1
  Type: vrf                      State: Active
  Restart State: Pending Path selection timeout: 300
  Interfaces:
    t3-0/0/0.102
  Route-distinguisher: 10.255.14.176:102
  Vrf-import: [ RIP-import ]
  Vrf-export: [ RIP-export ]
  Tables:
    RIP.inet.0                 : 8 routes (6 active, 2 holddown, 0 hidden)
    Restart Pending: RIP VPN

STATIC:
  Router ID: 10.69.100.1
  Type: vrf                      State: Active
  Restart State: Pending Path selection timeout: 300
  Interfaces:
    t3-0/0/0.100
  Route-distinguisher: 10.255.14.176:100
  Vrf-import: [ STATIC-import ]
  Vrf-export: [ STATIC-export ]
  Tables:
    STATIC.inet.0              : 4 routes (4 active, 0 holddown, 0 hidden)
    Restart Pending: VPN

```

show route instance detail (VPLS Routing Instance)

```

user@host> show route instance detail test-vpls
test-vpls:
  Router ID: 0.0.0.0
  Type: vpls                      State: Active
  Interfaces:
    lsi.1048833
    lsi.1048832
    fe-0/1/0.513
  Route-distinguisher: 10.255.37.65:1
  Vrf-import: [ __vrf-import-test-vpls-internal__ ]
  Vrf-export: [ __vrf-export-test-vpls-internal__ ]
  Vrf-import-target: [ target:300:1 ]

```

```

Vrf-export-target: [ target:300:1 ]
Fast-reroute-priority: high
Tables:
    test-vpls.l2vpn.0          : 3 routes (3 active, 0 holddown, 0 hidden)

```

show route instance operational

```

user@host> show route instance operational
Operational Routing Instances:

```

```

master
default

```

show route instance summary

```

user@host> show route instance summary

```

| Instance | Type | Primary rib | Active/holddown/hidden |
|----------|------------|------------------|------------------------|
| master | forwarding | inet.0 | 15/0/1 |
| | | iso.0 | 1/0/0 |
| | | mpls.0 | 35/0/0 |
| | | l3vpn.0 | 0/0/0 |
| | | inet6.0 | 2/0/0 |
| | | l2vpn.0 | 0/0/0 |
| | | l2circuit.0 | 0/0/0 |
| BGP-INET | vrf | BGP-INET.inet.0 | 5/0/0 |
| | | BGP-INET.iso.0 | 0/0/0 |
| | | BGP-INET.inet6.0 | 0/0/0 |
| BGP-L | vrf | BGP-L.inet.0 | 5/0/0 |
| | | BGP-L.iso.0 | 0/0/0 |
| | | BGP-L.mpls.0 | 4/0/0 |
| | | BGP-L.inet6.0 | 0/0/0 |
| L2VPN | l2vpn | L2VPN.inet.0 | 0/0/0 |
| | | L2VPN.iso.0 | 0/0/0 |
| | | L2VPN.inet6.0 | 0/0/0 |
| | | L2VPN.l2vpn.0 | 2/0/0 |
| LDP | vrf | LDP.inet.0 | 4/0/0 |
| | | LDP.iso.0 | 0/0/0 |
| | | LDP.mpls.0 | 0/0/0 |
| | | LDP.inet6.0 | 0/0/0 |
| | | LDP.l2circuit.0 | 0/0/0 |
| OSPF | vrf | OSPF.inet.0 | 7/0/0 |
| | | OSPF.iso.0 | 0/0/0 |
| | | OSPF.inet6.0 | 0/0/0 |
| RIP | vrf | RIP.inet.0 | 6/0/0 |
| | | RIP.iso.0 | 0/0/0 |
| | | RIP.inet6.0 | 0/0/0 |
| STATIC | vrf | STATIC.inet.0 | 4/0/0 |
| | | STATIC.iso.0 | 0/0/0 |
| | | STATIC.inet6.0 | 0/0/0 |

show route label

| | |
|------------------------------------|---|
| Syntax | <code>show route label <i>label</i></code> <code><brief detail extensive terse></code> <code><logical-system (all <i>logical-system-name</i>)></code> |
| Syntax (EX Series Switches) | <code>show route label <i>label</i></code> <code><brief detail extensive terse></code> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches. |
| Description | Display the routes based on a specified Multiprotocol Label Switching (MPLS) label value. |
| Options | <p><i>label</i>—Value of the MPLS label.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show route label on page 595 show route label detail on page 595 show route label extensive on page 595 show route label terse on page 595 |
| Output Fields | For information about output fields, see the output field table for the show route command, the show route detail command, the show route extensive command, or the show route terse command. |

Sample Output

show route label

```
user@host> show route label 100016

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
100016          *[VPN/170] 03:25:41
                > to 10.12.80.1 via ge-6/3/2.0, Pop
```

show route label detail

```
user@host> show route label 100016 detail

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
100016 (1 entry, 1 announced)
    *VPN      Preference: 170
              Next-hop reference count: 2
              Source: 10.12.80.1
              Next hop: 10.12.80.1 via ge-6/3/2.0, selected
              Label operation: Pop
              State: <Active Int Ext>
              Local AS:      1
              Age: 3:23:31
              Task: BGP.0.0.0.0+179
              Announcement bits (1): 0-KRT
              AS path: 100 I
              Ref Cnt: 2
```

show route label extensive

The output for the show route label extensive command is identical to that of the **show route label detail** command. For sample output, see [show route label detail on page 595](#).

show route label terse

```
user@host> show route label 100016 terse

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
* 100016           V 170                >10.12.80.1
```

show route label-switched-path

| | |
|------------------------------------|--|
| Syntax | <code>show route label-switched-path <i>path-name</i></code> <code><brief detail extensive terse></code> <code><logical-system (all <i>logical-system-name</i>)></code> |
| Syntax (EX Series Switches) | <code>show route label-switched-path <i>path-name</i></code> <code><brief detail extensive terse></code> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches. |
| Description | Display the routes used in an MPLS label-switched path (LSP). |
| Options | brief detail extensive terse —(Optional) Display the specified level of output. <i>path-name</i> —LSP tunnel name. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | view |
| List of Sample Output | show route label-switched-path on page 597 |
| Output Fields | For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command. |

Sample Output

show route
label-switched-path

```
user@host> show route label-switched-path sf-to-ny
inet.0: 29 destinations, 29 routes (29 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.1/32      [MPLS/7] 00:00:06, metric 0
> to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny
3.3.3.3/32      *[MPLS/7] 00:00:06, metric 0
> to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2.2.2.2/32      *[MPLS/7] 00:00:06, metric 0
> to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny
4.4.4.4/32      *[MPLS/7] 00:00:06, metric 0
to 111.222.1.9 via s0-0/0/0, label-switched-path abc
> to 111.222.1.9 via s0-0/0/0, label-switched-path xyz
to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny
111.222.1.9/32  [MPLS/7] 00:00:06, metric 0
> to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

mpls.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

show route localization

| | |
|---------------------------------|---|
| Syntax | show route localization |
| Release Information | Command introduced in Junos OS Release 11.4. |
| Description | (T320, T640, and T1600 routers only) Display route localization details. |
| Options | detail —Display detailed output. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> Example: Configuring Packet Forwarding Engine FIB Localization |
| Output Fields | Table 138 on page 598 lists the output fields for the show route localization command. Output fields are listed in the approximate order in which they appear. |

Table 138: show route localization Output Fields

| Field Name | Field Description |
|-------------------|--|
| FIB-local | FPCs configured as FIB-local. |
| FIB-remote | FPCs configured as FIB-remote. |
| Normal | FPCs neither configured as FIB-local or FIB-remote . |
| Protocols | IPv4 (inet) or IPv6 (inet6) traffic configured for route localization. |

Sample Output

```

user@R0> show route localization
FIB localization ready FPCs (and FIB-local Forwarding Engine addresses)
  FIB-local:  FPC2(4,5)
  FIB-remote: FPC0, FPC1
  Normal:     FPC3, FPC4, FPC5, FPC6, FPC7

user@R0> show route localization detail
FIB localization ready FPCs (and FIB-local Forwarding Engine addresses)
  FIB-local:  FPC2(4,5)
  FIB-remote: FPC0, FPC1
  Normal:     FPC3, FPC4, FPC5, FPC6, FPC7
FIB localization configuration
  Protocols:  inet, inet6
  FIB-local:  FPC2
  FIB-remote: FPC0, FPC1
Forwarding Engine addresses
  FPC0: 1
  FPC1: 2
  FPC2: 4, 5
  FPC3: 6

```


FPC4: 8
FPC5: 11
FPC6: 13
FPC7: 15

show route martians

| | |
|------------------------------------|---|
| Syntax | show route martians <logical-system (all <i>logical-system-name</i>)> <table <i>routing-table-name</i> > |
| Syntax (EX Series Switches) | show route martians <table <i>routing-table-name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. |
| Description | Display the martian (invalid and ignored) entries associated with each routing table. |
| Options | <p>none—Display standard information about route martians for all routing tables.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>table <i>routing-table-name</i>—(Optional) Display information about route martians for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the show route martians table inet command).</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> Example: Configuring Martian Addresses |
| List of Sample Output | show route martians on page 601 |
| Output Fields | Table 139 on page 600 lists the output fields for the show route martians command. Output fields are listed in the approximate order in which they appear |

Table 139: show route martians Output Fields

| Field Name | Field Description |
|---------------------------|---|
| <i>table-name</i> | Name of the route table in which the route martians reside. |
| <i>destination-prefix</i> | Route destination. |
| <i>match value</i> | Route match parameter. |
| <i>status</i> | Status of the route: allowed or disallowed . |

Sample Output

```

show route martians      user@host> show route martians

inet.0:
    0.0.0.0/0 exact -- allowed
    0.0.0.0/8 orlonger -- disallowed
    127.0.0.0/8 orlonger -- disallowed
    192.0.0.0/24 orlonger -- disallowed
    240.0.0.0/4 orlonger -- disallowed
    224.0.0.0/4 exact -- disallowed
    224.0.0.0/24 exact -- disallowed

inet.1:
    0.0.0.0/0 exact -- allowed
    0.0.0.0/8 orlonger -- disallowed
    127.0.0.0/8 orlonger -- disallowed
    192.0.0.0/24 orlonger -- disallowed
    240.0.0.0/4 orlonger -- disallowed

inet.2:
    0.0.0.0/0 exact -- allowed
    0.0.0.0/8 orlonger -- disallowed
    127.0.0.0/8 orlonger -- disallowed
    192.0.0.0/24 orlonger -- disallowed
    240.0.0.0/4 orlonger -- disallowed
    224.0.0.0/4 exact -- disallowed
    224.0.0.0/24 exact -- disallowed

inet.3:
    0.0.0.0/0 exact -- allowed
    0.0.0.0/8 orlonger -- disallowed
    127.0.0.0/8 orlonger -- disallowed
    192.0.0.0/24 orlonger -- disallowed
    240.0.0.0/4 orlonger -- disallowed
    224.0.0.0/4 exact -- disallowed
    224.0.0.0/24 exact -- disallowed
...

inet6.0:
    ::1/128 exact -- disallowed
    ff00::/8 exact -- disallowed
    ff02::/16 exact -- disallowed

inet6.1:
    ::1/128 exact -- disallowed

inet6.2:
    ::1/128 exact -- disallowed
    ff00::/8 exact -- disallowed
    ff02::/16 exact -- disallowed

inet6.3:
    ::1/128 exact -- disallowed
    ff00::/8 exact -- disallowed
    ff02::/16 exact -- disallowed
...

```

show route next-hop

| | |
|------------------------------------|--|
| Syntax | <code>show route next-hop <i>next-hop</i></code> <code><brief detail extensive terse></code> <code><logical-system (all <i>logical-system-name</i>)></code> |
| Syntax (EX Series Switches) | <code>show route next-hop <i>next-hop</i></code> <code><brief detail extensive terse></code> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. |
| Description | Display the entries in the routing table that are being sent to the specified next-hop address. |
| Options | brief detail extensive terse —(Optional) Display the specified level of output. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. <i>next-hop</i> —Next-hop address. |
| Required Privilege Level | view |
| List of Sample Output | show route next-hop on page 603 show route next-hop detail on page 603 show route next-hop extensive on page 605 show route next-hop terse on page 607 |
| Output Fields | For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command. |

Sample Output

show route next-hop

```

user@host> show route next-hop 192.168.71.254

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.10.0.0/16      *[Static/5] 06:26:25
                  > to 192.168.71.254 via fxp0.0
10.209.0.0/16    *[Static/5] 06:26:25
                  > to 192.168.71.254 via fxp0.0
172.16.0.0/12    *[Static/5] 06:26:25
                  > to 192.168.71.254 via fxp0.0
192.168.0.0/16   *[Static/5] 06:26:25
                  > to 192.168.71.254 via fxp0.0
192.168.102.0/23 *[Static/5] 06:26:25
                  > to 192.168.71.254 via fxp0.0
207.17.136.0/24  *[Static/5] 06:26:25
                  > to 192.168.71.254 via fxp0.0
207.17.136.192/32 *[Static/5] 06:26:25
                  > to 192.168.71.254 via fxp0.0

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

show route next-hop detail

```

user@host> show route next-hop 192.168.71.254 detail

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
Restart Complete
10.10.0.0/16 (1 entry, 1 announced)
    *Static Preference: 5
        Next-hop reference count: 36
        Next hop: 192.168.71.254 via fxp0.0, selected
        State: <Active NoReadvrt Int Ext>
        Local AS: 1
        Age: 6:27:41
        Task: RT
        Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
        AS path: I
10.209.0.0/16 (1 entry, 1 announced)
    *Static Preference: 5
        Next-hop reference count: 36
        Next hop: 192.168.71.254 via fxp0.0, selected
        State: <Active NoReadvrt Int Ext>

```

```
Local AS:      1
Age: 6:27:41
Task: RT
Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
AS path: I

172.16.0.0/12 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 36
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS:      1
    Age: 6:27:41
    Task: RT
    Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
    AS path: I

192.168.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 36
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS:      1
    Age: 6:27:41
    Task: RT
    Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
    AS path: I

192.168.102.0/23 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 36
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS:      1
    Age: 6:27:41
    Task: RT
    Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
    AS path: I

207.17.136.0/24 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 36
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS:      1
    Age: 6:27:41
    Task: RT
    Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
    AS path: I

207.17.136.192/32 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 36
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS:      1
    Age: 6:27:41
    Task: RT
    Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
    AS path: I
```

```

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

show route next-hop extensive

```

user@host> show route next-hop 192.168.71.254 extensive

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 2:02:28
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

10.209.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.209.0.0/16 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 2:02:28
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

172.16.0.0/12 (1 entry, 1 announced)
TSI:
KRT in-kernel 172.16.0.0/12 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 2:02:28
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

192.168.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 192.168.0.0/16 -> {192.168.71.254}

```

```
*Static Preference: 5
  Next-hop reference count: 22
  Next hop: 192.168.71.254 via fxp0.0, selected
  State: <Active NoReadvrt Int Ext>
  Local AS: 69
  Age: 2:02:28
  Task: RT
  Announcement bits (1): 0-KRT
  AS path: I

192.168.102.0/23 (1 entry, 1 announced)
TSI:
KRT in-kernel 192.168.102.0/23 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 2:02:28
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

207.17.136.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.0/24 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 2:02:28
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

207.17.136.192/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.192/32 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 2:02:28
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

green.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
```



```
red.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route next-hop terse

```
user@host> show route next-hop 192.168.71.254 terse
```

```
inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
```

```
Restart Complete
```

```
+ = Active Route, - = Last Active, * = Both
```

| A | Destination | P | Prf | Metric 1 | Metric 2 | Next hop | AS path |
|---|-------------------|---|-----|----------|----------|-----------------|---------|
| * | 10.10.0.0/16 | S | 5 | | | >192.168.71.254 | |
| * | 10.209.0.0/16 | S | 5 | | | >192.168.71.254 | |
| * | 172.16.0.0/12 | S | 5 | | | >192.168.71.254 | |
| * | 192.168.0.0/16 | S | 5 | | | >192.168.71.254 | |
| * | 192.168.102.0/23 | S | 5 | | | >192.168.71.254 | |
| * | 207.17.136.0/24 | S | 5 | | | >192.168.71.254 | |
| * | 207.17.136.192/32 | S | 5 | | | >192.168.71.254 | |

```
private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
```

```
red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
```

```
Restart Complete
```

```
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

```
Restart Complete
```

```
mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
```

```
Restart Complete
```

```
inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
```

```
Restart Complete
```

```
private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route no-community

| | |
|------------------------------------|---|
| Syntax | show route no-community <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switches) | show route no-community <brief detail extensive terse> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. |
| Description | Display the route entries in each routing table that are not associated with any community. |
| Options | none —(Same as brief) Display the route entries in each routing table that are not associated with any community. brief detail extensive terse —(Optional) Display the specified level of output. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | view |
| List of Sample Output | show route no-community on page 609 show route no-community detail on page 609 show route no-community extensive on page 609 show route no-community terse on page 610 |
| Output Fields | For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command. |

Sample Output

show route no-community

```
user@host> show route no-community
inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.0.0/16      *[Static/5] 00:36:27
                  > to 192.168.71.254 via fxp0.0
10.209.0.0/16    *[Static/5] 00:36:27
                  > to 192.168.71.254 via fxp0.0
10.255.71.52/32  *[Direct/0] 00:36:27
                  > via lo0.0
10.255.71.63/32  *[OSPF/10] 00:04:39, metric 1
                  > to 35.1.1.2 via ge-3/1/0.0
10.255.71.64/32  *[OSPF/10] 00:00:08, metric 2
                  > to 35.1.1.2 via ge-3/1/0.0
10.255.71.240/32 *[OSPF/10] 00:05:04, metric 2
                  via so-0/1/2.0
                  > via so-0/3/2.0
10.255.71.241/32 *[OSPF/10] 00:05:14, metric 1
                  > via so-0/1/2.0
10.255.71.242/32 *[OSPF/10] 00:05:19, metric 1
                  > via so-0/3/2.0
12.1.1.0/24      *[OSPF/10] 00:05:14, metric 2
                  > via so-0/3/2.0
14.1.1.0/24      *[OSPF/10] 00:00:08, metric 3
                  > to 35.1.1.2 via ge-3/1/0.0
                  via so-0/1/2.0
                  via so-0/3/2.0
16.1.1.0/24      *[OSPF/10] 00:05:14, metric 2
                  > via so-0/1/2.0
.....
```

show route no-community detail

```
user@host> show route no-community detail

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
    *Static Preference: 5
        Next-hop reference count: 22
        Next hop: 192.168.71.254 via fxp0.0, selected
        State: <Active NoReadvrt Int Ext>
        Age: 38:08
        Task: RT
        Announcement bits (1): 0-KRT
        AS path: I

10.209.0.0/16 (1 entry, 1 announced)
    *Static Preference: 5
        Next-hop reference count: 22
        Next hop: 192.168.71.254 via fxp0.0, selected
        State: <Active NoReadvrt Int Ext>
        Age: 38:08
        Task: RT
        Announcement bits (1): 0-KRT
        AS path: I
....

user@host> show route no-community extensive
```

show route
no-community
extensive

```
inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 2:03:33
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

10.209.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.209.0.0/16 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 2:03:33
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I
```

show route
no-community terse

```
user@host> show route no-community terse

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
```

| A | Destination | P | Prf | Metric 1 | Metric 2 | Next hop | AS path |
|---|------------------|---|-----|----------|----------|-----------------|---------|
| * | 10.10.0.0/16 | S | 5 | | | >192.168.71.254 | |
| * | 10.209.0.0/16 | S | 5 | | | >192.168.71.254 | |
| * | 10.255.71.52/32 | D | 0 | | | >lo0.0 | |
| * | 10.255.71.63/32 | 0 | 10 | 1 | | >35.1.1.2 | |
| * | 10.255.71.64/32 | 0 | 10 | 2 | | >35.1.1.2 | |
| * | 10.255.71.240/32 | 0 | 10 | 2 | | so-0/1/2.0 | |
| | | | | | | >so-0/3/2.0 | |
| * | 10.255.71.241/32 | 0 | 10 | 1 | | >so-0/1/2.0 | |
| * | 10.255.71.242/32 | 0 | 10 | 1 | | >so-0/3/2.0 | |
| * | 12.1.1.0/24 | 0 | 10 | 2 | | >so-0/3/2.0 | |
| * | 14.1.1.0/24 | 0 | 10 | 3 | | >35.1.1.2 | |
| | | | | | | so-0/1/2.0 | |
| | | | | | | so-0/3/2.0 | |
| * | 16.1.1.0/24 | 0 | 10 | 2 | | >so-0/1/2.0 | |
| | | | | | | | |

...

show route output

| | |
|------------------------------------|--|
| Syntax | show route output (address <i>ip-address</i> interface <i>interface-name</i>) <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switches) | show route output (address <i>ip-address</i> interface <i>interface-name</i>) <brief detail extensive terse> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. |
| Description | <p>Display the entries in the routing table learned through static routes and interior gateway protocols that are to be sent out the interface with either the specified IP address or specified name.</p> <p>To view routes advertised to a neighbor or received from a neighbor for the BGP protocol, use the show route advertising-protocol bgp and show route receive-protocol bgp commands instead.</p> |
| Options | <p>address <i>ip-address</i>—Display entries in the routing table that are to be sent out the interface with the specified IP address.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>interface <i>interface-name</i>—Display entries in the routing table that are to be sent out the interface with the specified name.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show route output address on page 612 show route output address detail on page 612 show route output address extensive on page 612 show route output address terse on page 612 show route output interface on page 613 show route output interface detail on page 613 show route output interface extensive on page 614 show route output interface terse on page 614 |
| Output Fields | For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command. |

Sample Output

show route output address

```
user@host> show route output address 36.1.1.1/24

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

36.1.1.0/24          *[Direct/0] 00:19:56
                    > via so-0/1/2.0
                    [OSPF/10] 00:19:55, metric 1
                    > via so-0/1/2.0

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route output address detail

```
user@host> show route output address 36.1.1.1 detail

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
36.1.1.0/24 (2 entries, 0 announced)
    *Direct Preference: 0
        Next hop type: Interface
        Next-hop reference count: 1
        Next hop: via so-0/1/2.0, selected
        State: <Active Int>
        Age: 23:00
        Task: IF
        AS path: I
    OSPF Preference: 10
        Next-hop reference count: 1
        Next hop: via so-0/1/2.0, selected
        State: <Int>
        Inactive reason: Route Preference
        Age: 22:59      Metric: 1
        Area: 0.0.0.0
        Task: OSPF
        AS path: I

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route output address extensive

The output for the **show route output address extensive** command is identical to that of the **show route output address detail** command. For sample output, see [show route output address detail on page 612](#).

**show route output
address terse**

```
user@host> show route output address 36.1.1.1 terse
```

```
inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1   Metric 2   Next hop      AS path
* 36.1.1.0/24      D   0           1           >so-0/1/2.0
                   0  10           1           >so-0/1/2.0

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

**show route output
interface**

```
user@host> show route output interface so-0/1/2.0
```

```
inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.71.240/32  * [OSPF/10] 00:13:00, metric 2
                   via so-0/1/2.0
                   > via so-0/3/2.0
10.255.71.241/32  * [OSPF/10] 00:13:10, metric 1
                   > via so-0/1/2.0
14.1.1.0/24       * [OSPF/10] 00:05:11, metric 3
                   to 35.1.1.2 via ge-3/1/0.0
                   > via so-0/1/2.0
                   via so-0/3/2.0
16.1.1.0/24       * [OSPF/10] 00:13:10, metric 2
                   > via so-0/1/2.0
36.1.1.0/24       * [Direct/0] 00:13:21
                   > via so-0/1/2.0
                   [OSPF/10] 00:13:20, metric 1
                   > via so-0/1/2.0

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

**show route output
interface detail**

```
user@host> show route output interface so-0/1/2.0 detail
```

```
inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
10.255.71.240/32 (1 entry, 1 announced)
  *OSPF   Preference: 10
           Next-hop reference count: 2
           Next hop: via so-0/1/2.0
           Next hop: via so-0/3/2.0, selected
           State: <Active Int>
           Age: 14:52      Metric: 2
```

```

Area: 0.0.0.0
Task: OSPF
Announcement bits (1): 0-KRT
AS path: I

10.255.71.241/32 (1 entry, 1 announced)
*OSPF Preference: 10
Next-hop reference count: 4
Next hop: via so-0/1/2.0, selected
State: <Active Int>
Age: 15:02 Metric: 1
Area: 0.0.0.0
Task: OSPF
Announcement bits (1): 0-KRT
AS path: I
...

```

show route output interface extensive

The output for the **show route output interface extensive** command is identical to that of the **show route output interface detail** command. For sample output, see [show route output interface detail on page 613](#).

show route output interface terse

```

user@host> show route output interface so-0/1/2.0 terse

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
* 10.255.71.240/32  0 10      2          so-0/1/2.0
                        >so-0/3/2.0
* 10.255.71.241/32  0 10      1          >so-0/1/2.0
* 14.1.1.0/24       0 10      3          35.1.1.2
                        >so-0/1/2.0
                        so-0/3/2.0
* 16.1.1.0/24       0 10      2          >so-0/1/2.0
* 36.1.1.0/24       D 0
                        0 10      1          >so-0/1/2.0
                        >so-0/1/2.0

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```


show route protocol

| | |
|------------------------------------|--|
| Syntax | <pre>show route protocol <i>protocol</i> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)></pre> |
| Syntax (EX Series Switches) | <pre>show route protocol <i>protocol</i> <brief detail extensive terse></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>ospf2 and ospf3 options introduced in Junos OS Release 9.2.</p> <p>ospf2 and ospf3 options introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>flow option introduced in Junos OS Release 10.0.</p> <p>flow option introduced in Junos OS Release 10.0 for EX Series switches.</p> |
| Description | Display the route entries in the routing table that were learned from a particular protocol. |
| Options | <p>brief detail extensive terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><i>protocol</i>—Protocol from which the route was learned:</p> <ul style="list-style-type: none"> • access—Access route for use by DHCP application • access-internal—Access-internal route for use by DHCP application • aggregate—Locally generated aggregate route • arp—Route learned through the Address Resolution Protocol • atmvpn—Asynchronous Transfer Mode virtual private network • bgp—Border Gateway Protocol • ccc—Circuit cross-connect • direct—Directly connected route • dvmrp—Distance Vector Multicast Routing Protocol • esis—End System-to-Intermediate System • flow—Locally defined flow-specification route • frr—Precomputed protection route or backup route used when a link goes down • isis—Intermediate System-to-Intermediate System • ldp—Label Distribution Protocol • l2circuit—Layer 2 circuit • l2vpn—Layer 2 virtual private network |

- **local**—Local address
- **mpls**—Multiprotocol Label Switching
- **msdp**—Multicast Source Discovery Protocol
- **ospf**—Open Shortest Path First versions 2 and 3
- **ospf2**—Open Shortest Path First versions 2 only
- **ospf3**—Open Shortest Path First version 3 only
- **pim**—Protocol Independent Multicast
- **rip**—Routing Information Protocol
- **ripng**—Routing Information Protocol next generation
- **rsvp**—Resource Reservation Protocol
- **rtarget**—Local route target virtual private network
- **static**—Statically defined route
- **tunnel**—Dynamic tunnel
- **vpn**—Virtual private network



NOTE: EX Series switches run a subset of these protocols. See the switch CLI for details.

| | |
|---------------------------------|---|
| Required Privilege Level | view |
| List of Sample Output | show route protocol access on page 618 show route protocol access-internal extensive on page 618 show route protocol arp on page 618 show route protocol bgp on page 619 show route protocol bgp detail on page 619 show route protocol bgp extensive on page 619 show route protocol bgp terse on page 620 show route protocol direct on page 620 show route protocol frr on page 620 show route protocol l2circuit detail on page 621 show route protocol l2vpn extensive on page 622 show route protocol ldp on page 622 show route protocol ldp extensive on page 623 show route protocol ospf (Layer 3 VPN) on page 624 show route protocol ospf detail on page 624 show route protocol rip on page 625 show route protocol rip detail on page 625 show route protocol ripng table inet6 on page 625 show route protocol static detail on page 625 |

Output Fields For information about output fields, see the output field tables for the [show route](#) command, the [show route detail](#) command, the [show route extensive](#) command, or the [show route terse](#) command.

Sample Output

show route protocol access

```
user@host> show route protocol access
inet.0: 30380 destinations, 30382 routes (30379 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

13.160.0.3/32      *[Access/13] 00:00:09
                  > to 13.160.0.2 via fe-0/0/0.0
13.160.0.4/32      *[Access/13] 00:00:09
                  > to 13.160.0.2 via fe-0/0/0.0
13.160.0.5/32      *[Access/13] 00:00:09
                  > to 13.160.0.2 via fe-0/0/0.0
```

show route protocol access-internal extensive

```
user@host> show route protocol access-internal 13.160.0.19 extensive
inet.0: 100020 destinations, 100022 routes (100019 active, 0 holddown, 1 hidden)
13.160.0.19/32 (1 entry, 1 announced)
TSI:
KRT in-kerne1 13.160.0.19/32 -> {13.160.0.2}
    *Access-internal Preference: 12
      Next-hop reference count: 200000
      Next hop: 13.160.0.2 via fe-0/0/0.0, selected
      State: <Active Int>
    Age: 36
      Task: RPD Unix Domain Server./var/run/rpd_serv.local
      Announcement bits (1): 0-KRT
      AS path: I
```

show route protocol arp

```
user@host> show route protocol arp
inet.0: 43 destinations, 43 routes (42 active, 0 holddown, 1 hidden)

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

cust1.inet.0: 1033 destinations, 2043 routes (1033 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.1.3/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.4/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.5/32      [ARP/4294967293] 00:04:32, from 20.20.1.1
                  Unusable
20.20.1.6/32      [ARP/4294967293] 00:04:34, from 20.20.1.1
                  Unusable
20.20.1.7/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.8/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.9/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.10/32     [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.11/32     [ARP/4294967293] 00:04:33, from 20.20.1.1
                  Unusable
20.20.1.12/32     [ARP/4294967293] 00:04:33, from 20.20.1.1
                  Unusable
20.20.1.13/32     [ARP/4294967293] 00:04:33, from 20.20.1.1
                  Unusable
...
```

**show route protocol
bgp**

```

user@host> show route protocol bgp 192.168.64.0/21
inet.0: 335832 destinations, 335833 routes (335383 active, 0 holddown, 450 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.64.0/21      *[BGP/170] 6d 10:41:16, localpref 100, from 192.168.69.71
                    AS path: 10458 14203 2914 4788 4788 I
                    > to 192.168.167.254 via fxp0.0

```

**show route protocol
bgp detail**

```

user@host> show route protocol bgp 66.117.63.0/24 detail
inet.0: 335805 destinations, 335806 routes (335356 active, 0 holddown, 450 hidden)
66.117.63.0/24      (1 entry, 1 announced)
    *BGP           Preference: 170/-101
                   Next hop type: Indirect
                   Next-hop reference count: 1006436
                   Source: 192.168.69.71
                   Next hop type: Router, Next hop index: 324
                   Next hop: 192.168.167.254 via fxp0.0, selected
                   Protocol next hop: 192.168.69.71
                   Indirect next hop: 8e166c0 342
                   State: <Active Ext>
                   Local AS: 69 Peer AS: 10458
                   Age: 6d 10:42:42 Metric2: 0
                   Task: BGP_10458.192.168.69.71+179
                   Announcement bits (3): 0-KRT 2-BGP RT Background 3-Resolve tree

1
    AS path: 10458 14203 2914 4788 4788 I
    Communities: 2914:410 2914:2403 2914:3400
    Accepted
    Localpref: 100
    Router ID: 207.17.136.192

```

**show route protocol
bgp extensive**

```

user@host> show route protocol bgp 192.168.64.0/21 extensive
inet.0: 335827 destinations, 335828 routes (335378 active, 0 holddown, 450 hidden)
192.168.64.0/21 (1 entry, 1 announced)
TSI:
KRT in-kernel 1.9.0.0/16 -> {indirect(342)}
Page 0 idx 1 Type 1 val db31a80
  Nexthop: Self
  AS path: [69] 10458 14203 2914 4788 4788 I
  Communities: 2914:410 2914:2403 2914:3400
Path 1.9.0.0 from 192.168.69.71 Vector len 4. Val: 1
    *BGP           Preference: 170/-101
                   Next hop type: Indirect
                   Next-hop reference count: 1006502
                   Source: 192.168.69.71
                   Next hop type: Router, Next hop index: 324
                   Next hop: 192.168.167.254 via fxp0.0, selected
                   Protocol next hop: 192.168.69.71
                   Indirect next hop: 8e166c0 342
                   State: <Active Ext>
                   Local AS: 69 Peer AS: 10458
                   Age: 6d 10:44:45 Metric2: 0
                   Task: BGP_10458.192.168.69.71+179
                   Announcement bits (3): 0-KRT 2-BGP RT Background 3-Resolve tree

1
    AS path: 10458 14203 2914 4788 4788 I
    Communities: 2914:410 2914:2403 2914:3400

```

```

Accepted
Localpref: 100
Router ID: 207.17.136.192
Indirect next hops: 1
  Protocol next hop: 192.168.69.71
  Indirect next hop: 8e166c0 342
  Indirect path forwarding next hops: 1
    Next hop type: Router
    Next hop: 192.168.167.254 via fxp0.0
  192.168.0.0/16 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
    Nexthop: 192.168.167.254 via fxp0.0

```

show route protocol bgp terse

```
user@host> show route protocol bgp 192.168.64.0/21 terse
```

```
inet.0: 24 destinations, 32 routes (23 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
```

| A Destination | P Prf | Metric 1 | Metric 2 | Next hop | AS path |
|-----------------|-------|----------|----------|------------|------------|
| 192.168.64.0/21 | B 170 | 100 | | >100.1.3.2 | 10023 21 I |

show route protocol direct

```
user@host> show route protocol direct
```

```
inet.0: 335843 destinations, 335844 routes (335394 active, 0 holddown, 450 hidden)
+ = Active Route, - = Last Active, * = Both
```

```

8.8.8.0/24          *[Direct/0] 17w0d 10:31:49
> via fe-1/3/1.0
10.255.165.1/32     *[Direct/0] 25w4d 04:13:18
> via lo0.0
30.30.30.0/24       *[Direct/0] 17w0d 23:06:26
> via fe-1/3/2.0
192.168.164.0/22    *[Direct/0] 25w4d 04:13:20
> via fxp0.0

```

```
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```

47.0005.80ff.f800.0000.0108.0001.0102.5516.5001/152
*[Direct/0] 25w4d 04:13:21
> via lo0.0

```

```
inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```

abcd::10:255:165:1/128
*[Direct/0] 25w4d 04:13:21
> via lo0.0
fe80::2a0:a5ff:fe12:ad7/128
*[Direct/0] 25w4d 04:13:21
> via lo0.0

```

show route protocol frr

```
user@host> show route protocol frr
```

```
inet.0: 43 destinations, 43 routes (42 active, 0 holddown, 1 hidden)
```

```
inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
```

```
cust1.inet.0: 1033 destinations, 2043 routes (1033 active, 0 holddown, 0 hidden)
```

+ = Active Route, - = Last Active, * = Both

```

20.20.1.3/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.3 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.4/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.4 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.5/32      *[FRR/200] 00:05:35, from 20.20.1.1
                  > to 20.20.1.5 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.6/32      *[FRR/200] 00:05:37, from 20.20.1.1
                  > to 20.20.1.6 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.7/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.7 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.8/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.8 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.9/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.9 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.10/32     *[FRR/200] 00:05:38, from 20.20.1.1
...

```

show route protocol l2circuit detail

user@host> show route protocol l2circuit detail

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
100000 (1 entry, 1 announced)

```

*L2CKT Preference: 7
  Next hop: via ge-2/0/0.0, selected
  Label operation: Pop      Offset: 4
  State: <Active Int>
  Local AS:    99
  Age: 9:52
  Task: Common L2 VC
  Announcement bits (1): 0-KRT
  AS path: I

```

ge-2/0/0.0 (1 entry, 1 announced)

```

*L2CKT Preference: 7
  Next hop: via so-1/1/2.0 weight 1, selected
  Label-switched-path my-lsp
  Label operation: Push 100000, Push 100000(top)[0] Offset: -4
  Protocol next hop: 10.245.255.63
  Push 100000 Offset: -4
  Indirect next hop: 86af0c0 298
  State: <Active Int>
  Local AS:    99
  Age: 9:52
  Task: Common L2 VC
  Announcement bits (2): 0-KRT 1-Common L2 VC
  AS path: I

```

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)

```

*L2CKT Preference: 7
  Next hop: via so-1/1/2.0 weight 1, selected
  Label-switched-path my-lsp

```

```

Label operation: Push 100000[0]
Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
State: <Active Int>
Local AS: 99
Age: 10:21
Task: 12 circuit
Announcement bits (1): 0-LDP
AS path: I
VC Label 100000, MTU 1500, VLAN ID 512

```

show route protocol l2vpn extensive

```

user@host> show route protocol l2vpn extensive

inet.0: 14 destinations, 15 routes (13 active, 0 holddown, 1 hidden)

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
800001 (1 entry, 1 announced)
TSI:
KRT in-kernel 800001 /36 -> {so-0/0/0.0}
    *L2VPN Preference: 7
      Next hop: via so-0/0/0.0 weight 49087 balance 97%, selected
      Label operation: Pop Offset: 4
      State: <Active Int>
      Local AS: 69
      Age: 7:48
      Task: Common L2 VC
      Announcement bits (1): 0-KRT
      AS path: I

so-0/0/0.0 (1 entry, 1 announced)
TSI:
KRT in-kernel so-0/0/0.0 /16 -> {indirect(288)}
    *L2VPN Preference: 7
      Next hop: via so-0/0/1.0, selected
      Label operation: Push 800000 Offset: -4
      Protocol next hop: 10.255.14.220
      Push 800000 Offset: -4
      Indirect next hop: 85142a0 288
      State: <Active Int>
      Local AS: 69
      Age: 7:48
      Task: Common L2 VC
      Announcement bits (2): 0-KRT 1-Common L2 VC
      AS path: I
      Communities: target:69:1 Layer2-info: encaps:PPP,
      control flags:2, mtu: 0

```

show route protocol ldp

```

user@host> show route protocol ldp

inet.0: 12 destinations, 13 routes (12 active, 0 holddown, 0 hidden)

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.16.1/32    *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Push 100000
192.168.17.1/32    *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0

```



```
private1___.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
```

```
mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
```

```
+ = Active Route, - = Last Active, * = Both
```

```
100064          *[LDP/9] 1d 23:03:35, metric 1
                 > via t1-4/0/0.0, Pop
100064(S=0)     *[LDP/9] 1d 23:03:35, metric 1
                 > via t1-4/0/0.0, Pop
100080          *[LDP/9] 1d 23:03:35, metric 1
                 > via t1-4/0/0.0, Swap 100000
```

show route protocol ldp extensive

```
user@host> show route protocol ldp extensive
192.168.16.1/32 (1 entry, 1 announced)
  State: <FlashAll>
  *LDP   Preference: 9
         Next-hop reference count: 3
         Next hop: via t1-4/0/0.0, selected
         Label operation: Push 100000
         State: <Active Int>
         Local AS: 65500
         Age: 1d 23:03:58           Metric: 1
         Task: LDP
         Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 2
         AS path: I

192.168.17.1/32 (1 entry, 1 announced)
  State: <FlashAll>
  *LDP   Preference: 9
         Next-hop reference count: 3
         Next hop: via t1-4/0/0.0, selected
         State: <Active Int>
         Local AS: 65500
         Age: 1d 23:03:58           Metric: 1
         Task: LDP
         Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 2
         AS path: I

private1___.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)

100064 (1 entry, 1 announced)
TSI:
KRT in-kerne1 100064 /36 -> {t1-4/0/0.0}
  *LDP   Preference: 9
         Next-hop reference count: 2
         Next hop: via t1-4/0/0.0, selected
         State: <Active Int>
         Local AS: 65500
         Age: 1d 23:03:58           Metric: 1
         Task: LDP
         Announcement bits (1): 0-KRT
         AS path: I
         Prefixes bound to route: 192.168.17.1/32

100064(S=0) (1 entry, 1 announced)
TSI:
KRT in-kerne1 100064 /40 -> {t1-4/0/0.0}
  *LDP   Preference: 9
```

```

Next-hop reference count: 2
Next hop: via t1-4/0/0.0, selected
Label operation: Pop
State: <Active Int>
Local AS: 65500
Age: 1d 23:03:58      Metric: 1
Task: LDP
Announcement bits (1): 0-KRT
AS path: I

100080 (1 entry, 1 announced)
TSI:
KRT in-kernel 100080 /36 -> {t1-4/0/0.0}
  *LDP      Preference: 9
            Next-hop reference count: 2
            Next hop: via t1-4/0/0.0, selected
            Label operation: Swap 100000
            State: <Active Int>
            Local AS: 65500
            Age: 1d 23:03:58      Metric: 1
            Task: LDP
            Announcement bits (1): 0-KRT
            AS path: I
            Prefixes bound to route: 192.168.16.1/32

```

show route protocol ospf (Layer 3 VPN)

```

user@host> show route protocol ospf
inet.0: 40 destinations, 40 routes (39 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.4/30      *[OSPF/10] 00:05:18, metric 4
                  > via t3-3/2/0.0
10.39.1.8/30      [OSPF/10] 00:05:18, metric 2
                  > via t3-3/2/0.0
10.255.14.171/32 *[OSPF/10] 00:05:18, metric 4
                  > via t3-3/2/0.0
10.255.14.179/32 *[OSPF/10] 00:05:18, metric 2
                  > via t3-3/2/0.0
224.0.0.5/32     *[OSPF/10] 20:25:55, metric 1

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.16/30     [OSPF/10] 00:05:43, metric 1
                  > via so-0/2/2.0
10.255.14.173/32 *[OSPF/10] 00:05:43, metric 1
                  > via so-0/2/2.0
224.0.0.5/32     *[OSPF/10] 20:26:20, metric 1

```

show route protocol ospf detail

```

user@host> show route protocol ospf detail
VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.16/30 (2 entries, 0 announced)
  OSPF      Preference: 10
            Nexthop: via so-0/2/2.0, selected
            State: <Int>
            Inactive reason: Route Preference
            Age: 6:25      Metric: 1
            Area: 0.0.0.0
            Task: VPN-AB-OSPF

```

```
AS path: I
Communities: Route-Type:0.0.0.0:1:0
```

```
...
```

show route protocol rip

```
user@host> show route protocol rip
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32    *[RIP/100] 20:24:34, metric 2
                  > to 10.39.1.22 via t3-0/2/2.0
224.0.0.9/32      *[RIP/100] 00:03:59, metric 1
```

show route protocol rip detail

```
user@host> show route protocol rip detail
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32 (1 entry, 1 announced)
    *RIP      Preference: 100
              Nexthop: 10.39.1.22 via t3-0/2/2.0, selected
              State: <Active Int>
              Age: 20:25:02   Metric: 2
              Task: VPN-AB-RIPv2
              Announcement bits (2): 0-KRT 2-BGP.0.0.0.0+179
              AS path: I
              Route learned from 10.39.1.22 expires in 96 seconds
```

show route protocol ripng table inet6

```
user@host> show route protocol ripng table inet6
inet6.0: 4215 destinations, 4215 routes (4214 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

1111::1/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::2/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::3/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::4/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::5/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::6/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
```

show route protocol static detail

```
user@host> show route protocol static detail
inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
10.5.0.0/16 (1 entry, 1 announced)
    *Static Preference: 5
      Next hop type: Router, Next hop index: 324
      Address: 0x9274010
      Next-hop reference count: 27
      Next hop: 192.168.187.126 via fxp0.0, selected
      Session Id: 0x0
      State: <Active NoReadvrt Int Ext>
```

Age: 7w3d 21:24:25
 Validation State: unverified
 Task: RT
 Announcement bits (1): 0-KRT
 AS path: I

10.10.0.0/16 (1 entry, 1 announced)

*Static Preference: 5
 Next hop type: Router, Next hop index: 324
 Address: 0x9274010
 Next-hop reference count: 27
 Next hop: 192.168.187.126 via fxp0.0, selected
 Session Id: 0x0
 State: <Active NoReadvrt Int Ext>
 Age: 7w3d 21:24:25
 Validation State: unverified
 Task: RT
 Announcement bits (1): 0-KRT
 AS path: I

10.13.10.0/23 (1 entry, 1 announced)

*Static Preference: 5
 Next hop type: Router, Next hop index: 324
 Address: 0x9274010
 Next-hop reference count: 27
 Next hop: 192.168.187.126 via fxp0.0, selected
 Session Id: 0x0
 State: <Active NoReadvrt Int Ext>
 Age: 7w3d 21:24:25
 Validation State: unverified
 Task: RT
 Announcement bits (1): 0-KRT
 AS path: I

show route range

| | |
|------------------------------------|---|
| Syntax | <pre>show route range <brief detail extensive terse> <destination-prefix> <logical-system (all <i>logical-system-name</i>)></pre> |
| Syntax (EX Series Switches) | <pre>show route range <brief detail extensive terse> <destination-prefix></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> |
| Description | Display routing table entries using a prefix range. |
| Options | <p>none—Display standard information about all routing table entries using a prefix range.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>destination-prefix—(Optional) Destination and prefix mask for the range.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | <p>show route range on page 628</p> <p>show route range destination-prefix on page 628</p> <p>show route range detail on page 628</p> <p>show route range extensive on page 629</p> <p>show route range terse on page 630</p> |
| Output Fields | For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command. |

Sample Output

show route range

```
user@host> show route range

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.0.0/16      *[Static/5] 00:30:01
                  > to 192.168.71.254 via fxp0.0
10.209.0.0/16    *[Static/5] 00:30:01
                  > to 192.168.71.254 via fxp0.0
10.255.71.14/32  *[Direct/0] 00:30:01
                  > via lo0.0
172.16.0.0/12    *[Static/5] 00:30:01
                  > to 192.168.71.254 via fxp0.0
192.168.0.0/16   *[Static/5] 00:30:01
                  > to 192.168.71.254 via fxp0.0
192.168.64.0/21  *[Direct/0] 00:30:01
                  > via fxp0.0
192.168.71.14/32 *[Local/0] 00:30:01
                  Local via fxp0.0
192.168.102.0/23 *[Static/5] 00:30:01
                  > to 192.168.71.254 via fxp0.0
...
```

show route range destination-prefix

```
user@host> show route range 192.168.0.0

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.0.0/16   *[Static/5] 00:31:14
                  > to 192.168.71.254 via fxp0.0
192.168.64.0/21  *[Direct/0] 00:31:14
                  > via fxp0.0
192.168.71.14/32 *[Local/0] 00:31:14
                  Local via fxp0.0
192.168.102.0/23 *[Static/5] 00:31:14
                  > to 192.168.71.254 via fxp0.0
```

show route range detail

```
user@host> show route range detail

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
    *Static Preference: 5
        Next-hop reference count: 22
        Next hop: 192.168.71.254 via fxp0.0, selected
        State: <Active NoReadvrt Int Ext>
        Age: 30:05
        Task: RT
        Announcement bits (1): 0-KRT
        AS path: I

10.209.0.0/16 (1 entry, 1 announced)
    *Static Preference: 5
        Next-hop reference count: 22
        Next hop: 192.168.71.254 via fxp0.0, selected
        State: <Active NoReadvrt Int Ext>
        Age: 30:05
```

```

Task: RT
Announcement bits (1): 0-KRT
AS path: I

10.255.71.14/32 (1 entry, 0 announced)
*Direct Preference: 0
  Next hop type: Interface
  Next-hop reference count: 1
  Next hop: via lo0.0, selected
  State: <Active Int>
  Age: 30:05
  Task: IF
  AS path: I

172.16.0.0/12 (1 entry, 1 announced)
*Static Preference: 5
  Next-hop reference count: 22
  Next hop: 192.168.71.254 via fxp0.0, selected
  State: <Active NoReadvrt Int Ext>
  Age: 30:05
  Task: RT
  Announcement bits (1): 0-KRT
  AS path: I

```

...

show route range extensive

```

user@host> show route range extensive

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kerne1 10.10.0.0/16 -> {192.168.71.254}
*Static Preference: 5
  Next-hop reference count: 22
  Next hop: 192.168.71.254 via fxp0.0, selected
  State: <Active NoReadvrt Int Ext>
  Age: 30:17
  Task: RT
  Announcement bits (1): 0-KRT
  AS path: I

10.209.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kerne1 10.209.0.0/16 -> {192.168.71.254}
*Static Preference: 5
  Next-hop reference count: 22
  Next hop: 192.168.71.254 via fxp0.0, selected
  State: <Active NoReadvrt Int Ext>
  Age: 30:17
  Task: RT
  Announcement bits (1): 0-KRT
  AS path: I

10.255.71.14/32 (1 entry, 0 announced)
*Direct Preference: 0
  Next hop type: Interface
  Next-hop reference count: 1
  Next hop: via lo0.0, selected
  State: <Active Int>
  Age: 30:17
  Task: IF

```

```

AS path: I

172.16.0.0/12 (1 entry, 1 announced)
TSI:
KRT in-kerne1 172.16.0.0/12 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 30:17
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

...

```

show route range terse user@host> **show route range terse**

```

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
* 10.10.0.0/16      S  5                >192.168.71.254
* 10.209.0.0/16     S  5                >192.168.71.254
* 10.255.71.14/32   D  0                >1o0.0
* 172.16.0.0/12     S  5                >192.168.71.254
* 192.168.0.0/16     S  5                >192.168.71.254
* 192.168.64.0/21   D  0                >fxp0.0
* 192.168.71.14/32  L  0                Local
* 192.168.102.0/23  S  5                >192.168.71.254
* 207.17.136.0/24   S  5                >192.168.71.254
* 207.17.136.192/32 S  5                >192.168.71.254

__juniper_private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
* 10.0.0.0/8        D  0                >fxp2.0
                   D  0                >fxp1.0
* 10.0.0.4/32       L  0                Local

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
47.0005.80ff.f800.0000.0108.0001.0102.5507.1014/152
*                   D  0                >1o0.0

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
abcd::10:255:71:14/128
*                   D  0                >1o0.0
fe80::280:42ff:fe11:226f/128
*                   D  0                >1o0.0

__juniper_private1__inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both

```


| A | Destination | P | Prf | Metric 1 | Metric 2 | Next hop | AS path |
|---|------------------------------|---|-----|----------|----------|------------|---------|
| | fe80::280:42ff:fe11:226f/128 | | | | | | |
| * | | D | 0 | | | >1o0.16385 | |

show route receive-protocol

| | | |
|------------------------------------|--|--|
| Syntax | show route receive-protocol <i>protocol neighbor-address</i> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)> | |
| Syntax (EX Series Switches) | show route receive-protocol <i>protocol neighbor-address</i> <brief detail extensive terse> | |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. | |
| Description | Display the routing information as it was received through a particular neighbor using a particular dynamic routing protocol. | |
| Options | brief detail extensive terse —(Optional) Display the specified level of output. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. <i>protocol neighbor-address</i> —Protocol transmitting the route (bgp , dvmrp , msdp , pim , rip , or ripng) and address of the neighboring router from which the route entry was received. | |
| Additional Information | The output displays the selected routes and the attributes with which they were received, but does not show the effects of import policy on the routing attributes. | |
| Required Privilege Level | view | |
| List of Sample Output | show route receive-protocol bgp on page 635 show route receive-protocol bgp extensive on page 635 show route receive-protocol bgp table extensive on page 635 show route receive-protocol bgp logical-system extensive on page 635 show route receive-protocol bgp detail (Layer 2 VPN) on page 636 show route receive-protocol bgp extensive (Layer 2 VPN) on page 637 show route receive-protocol bgp (Layer 3 VPN) on page 638 show route receive-protocol bgp detail (Layer 3 VPN) on page 638 show route receive-protocol bgp extensive (Layer 3 VPN) on page 639 | |
| Output Fields | Table 140 on page 632 describes the output fields for the show route receive-protocol command. Output fields are listed in the approximate order in which they appear. | |

Table 140: show route receive-protocol Output Fields

| Field Name | Field Description | Level of Output |
|----------------------------|---|-----------------|
| <i>routing-table-name</i> | Name of the routing table—for example, inet.0. | All levels |
| <i>number destinations</i> | Number of destinations for which there are routes in the routing table. | All levels |

Table 140: show route receive-protocol Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--|--|-------------------------|
| number routes | Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> • active • holddown (routes that are in pending state before being declared inactive) • hidden (routes that are not used because of a routing policy) | All levels |
| Prefix | Destination prefix. | none brief |
| MED | Multiple exit discriminator value included in the route. | none brief |
| destination-prefix (entry, announced) | Destination prefix. The entry value is the number of routes for this destination, and the announced value is the number of routes being announced for this destination. | detail extensive |
| Route Distinguisher | 64-bit prefix added to IP subnets to make them unique. | detail extensive |
| Label-Base, range | First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device. | detail extensive |
| VPN Label | Virtual private network (VPN) label. Packets are sent between CE and PE routing devices by advertising VPN labels. VPN labels transit over either an RSVP or an LDP label-switched path (LSP) tunnel. | detail extensive |
| Next hop | Next hop to the destination. An angle bracket (>) indicates that the route is the selected route. | All levels |
| Localpref or Lclpref | Local preference value included in the route. | All levels |

Table 140: show route receive-protocol Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|----------------------------|--|-------------------------|
| AS path | <p>Autonomous system (AS) path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> • I—IGP. • E—EGP. • ?—Incomplete; typically, the AS path was aggregated. <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> • []—Brackets enclose the number that precedes the AS path. This number represents the number of ASs present in the AS path, when calculated as defined in RFC 4271. This value is used the AS-path merge process, as defined in RFC 4893. • []—If more than one AS number is configured on the router, or if AS path prepending is configured, brackets enclose the local AS number associated with the AS path. • { }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • ()—Parentheses enclose a confederation. • ([])—Parentheses and brackets enclose a confederation set. <p>NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p> | All levels |
| Cluster list | (For route reflected output only) Cluster ID sent by the route reflector. | detail extensive |
| Originator ID | (For route reflected output only) Address of routing device that originally sent the route to the route reflector. | detail extensive |
| Communities | Community path attribute for the route. See the Output Field table in the show route detail command for all possible values for this field. | detail extensive |
| AIGP | Accumulated interior gateway protocol (AIGP) BGP attribute. | detail extensive |
| Attrset AS | Number, local preference, and path of the AS that originated the route. These values are stored in the Attrset attribute at the originating routing device. | detail extensive |
| Layer2-info: encaps | Layer 2 encapsulation (for example, VPLS). | detail extensive |
| control flags | Control flags: none or Site Down . | detail extensive |
| mtu | Maximum transmission unit (MTU) of the Layer 2 circuit. | detail extensive |

Sample Output

**show route
receive-protocol bgp**

```
user@host> show route receive-protocol bgp 10.255.245.215

inet.0: 28 destinations, 33 routes (27 active, 0 holddown, 1 hidden)
Prefix          Next hop          MED      Lclpref    AS path
10.22.1.0/24     10.255.245.215    0        100        I
10.22.2.0/24     10.255.245.215    0        100        I
```

**show route
receive-protocol bgp
extensive**

```
user@host> show route receive-protocol bgp 10.255.245.63 extensive

inet.0: 244 destinations, 244 routes (243 active, 0 holddown, 1 hidden)
Prefix          Next hop          MED      Lclpref    AS path
1.1.1.0/24 (1 entry, 1 announced)
  Next hop: 10.0.50.3
  Localpref: 100
  AS path: I <Originator>
  Cluster list: 10.2.3.1
  Originator ID: 10.255.245.45
165.3.0.0/16 (1 entry, 1 announced)
  Next hop: 111.222.5.254
  Localpref: 100
  AS path: I <Originator>
  Cluster list: 10.2.3.1
  Originator ID: 10.255.245.68
165.4.0.0/16 (1 entry, 1 announced)
  Next hop: 111.222.5.254
  Localpref: 100
  AS path: I <Originator>
  Cluster list: 10.2.3.1
  Originator ID: 10.255.245.45
195.1.2.0/24 (1 entry, 1 announced)
  Next hop: 111.222.5.254
  Localpref: 100
  AS path: I <Originator>
  Cluster list: 10.2.3.1
  Originator ID: 10.255.245.68
inet.2: 63 destinations, 63 routes (63 active, 0 holddown, 0 hidden)
Prefix          Next hop          MED      Lclpref    AS path
inet.3: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix          Next hop          MED      Lclpref    AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Next hop          MED      Lclpref    AS path
mpls.0: 48 destinations, 48 routes (48 active, 0 holddown, 0 hidden)
```

**show route
receive-protocol bgp
table extensive**

```
user@host> show route receive-protocol bgp 207.17.136.192 table inet.0 66.117.68.0/24 extensive

inet.0: 227315 destinations, 227316 routes (227302 active, 0 holddown, 13 hidden)
* 66.117.63.0/24 (1 entry, 1 announced)
  Nexthop: 207.17.136.29
  Localpref: 100
  AS path: AS2 PA[6]: 14203 2914 3356 29748 33437 AS_TRANS
  AS path: AS4 PA[2]: 33437 393219
  AS path: Merged[6]: 14203 2914 3356 29748 33437 393219 I
  Communities: 2914:420
```

**show route
receive-protocol bgp**

```
user@host> show route receive-protocol bgp 10.0.0.9 logical-system PE4 extensive

inet.0: 12 destinations, 13 routes (12 active, 0 holddown, 0 hidden)
* 10.0.0.0/30 (1 entry, 1 announced)
```

**logical-system
extensive**

```

Accepted
Route Label: 3
Nexthop: 10.0.0.9
AS path: 13979 I

* 10.0.0.4/30 (1 entry, 1 announced)
Accepted
Route Label: 3
Nexthop: 10.0.0.9
AS path: 13979 I

10.0.0.8/30 (2 entries, 1 announced)
Accepted
Route Label: 3
Nexthop: 10.0.0.9
AS path: 13979 I

* 10.9.9.1/32 (1 entry, 1 announced)
Accepted
Route Label: 3
Nexthop: 10.0.0.9
AS path: 13979 I

* 10.100.1.1/32 (1 entry, 1 announced)
Accepted
Route Label: 3
Nexthop: 10.0.0.9
AS path: 13979 I

* 44.0.0.0/24 (1 entry, 1 announced)
Accepted
Route Label: 300096
Nexthop: 10.0.0.9
AS path: 13979 I
AIGP: 203

* 55.0.0.0/24 (1 entry, 1 announced)
Accepted
Route Label: 300112
Nexthop: 10.0.0.9
AS path: 13979 7018 I
AIGP: 25

* 66.0.0.0/24 (1 entry, 1 announced)
Accepted
Route Label: 300144
Nexthop: 10.0.0.9
AS path: 13979 7018 I

* 99.0.0.0/24 (1 entry, 1 announced)
Accepted
Route Label: 300160
Nexthop: 10.0.0.9
AS path: 13979 7018 I

```

**show route
receive-protocol bgp
detail (Layer 2 VPN)**

```

user@host> show route receive-protocol bgp 10.255.14.171 detail
inet.0: 68 destinations, 68 routes (67 active, 0 holddown, 1 hidden)
Prefix          Nexthop          MED    Lclpref AS path
inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

```

Prefix          Nexthop          MED    Lclpref AS path
mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
frame-vpn.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0
hidden)
Prefix          Nexthop          MED    Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 1 announced)
  Route Distinguisher: 10.255.245.35:1
  Label-base : 800000, range : 4, status-vector : 0x0
  Nexthop: 10.255.245.35
  Localpref: 100
  AS path: I
  Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
control flags: 0, mtu: 0
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 0 announced)
  Route Distinguisher: 10.255.245.35:1
  Label-base : 800000, range : 4, status-vector : 0x0
  Nexthop: 10.255.245.35
  Localpref: 100
  AS path: I
  Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
control flags:0, mtu: 0

```

show route
receive-protocol bgp

```

user@host> show route receive-protocol bgp 10.255.14.171 extensive
inet.0: 68 destinations, 68 routes (67 active, 0 holddown, 1 hidden)
Prefix          Nexthop          MED    Lclpref AS path

```

extensive (Layer 2 VPN)

```

inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
frame-vpn.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 1 announced)
  Route Distinguisher: 10.255.245.35:1
  Label-base : 800000, range : 4, status-vector : 0x0
  Nexthop: 10.255.245.35
  Localpref: 100
  AS path: I
  Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
  control flags:0, mtu: 0
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 0 announced)
  Route Distinguisher: 10.255.245.35:1
  Label-base : 800000, range : 4, status-vector : 0x0
  Nexthop: 10.255.245.35
  Localpref: 100
  AS path: I
  Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
  control flags:0, mtu: 0

```

**show route
receive-protocol bgp
(Layer 3 VPN)**

```

user@host> show route receive-protocol bgp 10.255.14.171
inet.0: 33 destinations, 33 routes (32 active, 0 holddown, 1 hidden)
Prefix          Nexthop          MED      Lclpref AS path
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
VPN-A.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
10.255.14.175/32  10.255.14.171          100 2 I
10.255.14.179/32  10.255.14.171          2    100 I
VPN-B.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
10.255.14.175/32  10.255.14.171          100 2 I
10.255.14.177/32  10.255.14.171          100 I
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
mpls.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
bgp.l3vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
10.255.14.171:300:10.255.14.177/32
                  10.255.14.171          100 I
10.255.14.171:100:10.255.14.179/32
                  10.255.14.171          2    100 I
10.255.14.171:200:10.255.14.175/32
                  10.255.14.171          100 2 I

```

**show route
receive-protocol bgp
detail (Layer 3 VPN)**

```

user@host> show route receive-protocol bgp 10.255.14.174 detail
inet.0: 16 destinations, 17 routes (15 active, 0 holddown, 1 hidden)
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
vpna.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
* 10.49.0.0/30 (1 entry, 1 announced)
  Route Distinguisher: 10.255.14.176:2
  VPN Label: 101264

```



```

    Nexthop: 10.255.14.174
    Localpref: 100
    AS path: I
    Communities: target:200:100
    AttrSet AS: 100
      Localpref: 100
      AS path: I
* 10.255.14.172/32 (1 entry, 1 announced)
  Route Distinguisher: 10.255.14.176:2
  VPN Label: 101280
  Nexthop: 10.255.14.174
  Localpref: 100
  AS path: I
  Communities: target:200:100
  AttrSet AS: 100
    Localpref: 100
    AS path: I
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
* 10.255.14.174:2:10.49.0.0/30 (1 entry, 0 announced)
  Route Distinguisher: 10.255.14.174:2
  VPN Label: 101264
  Nexthop: 10.255.14.174
  Localpref: 100
  AS path: I
  Communities: target:200:100
  AttrSet AS: 100
    Localpref: 100
    AS path: I
* 10.255.14.174:2:10.255.14.172/32 (1 entry, 0 announced)
  Route Distinguisher: 10.255.14.174:2
  VPN Label: 101280
  Nexthop: 10.255.14.174
  Localpref: 100
  AS path: I
  Communities: target:200:100
  AttrSet AS: 100
    Localpref: 100
    AS path: I
inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

```

show route
receive-protocol bgp

```

user@host> show route receive-protocol bgp 10.255.245.63 extensive
inet.0: 244 destinations, 244 routes (243 active, 0 holddown, 1 hidden)
  Prefix                Nexthop                MED    Lclpref AS path

```

extensive (Layer 3
VPN)

```

1.1.1.0/24 (1 entry, 1 announced)
  Nexthop: 10.0.50.3
  Localpref: 100
  AS path: I <Originator>
  Cluster list: 10.2.3.1
  Originator ID: 10.255.245.45
165.3.0.0/16 (1 entry, 1 announced)
  Nexthop: 111.222.5.254
  Localpref: 100
  AS path: I <Originator>
  Cluster list: 10.2.3.1
  Originator ID: 10.255.245.68
165.4.0.0/16 (1 entry, 1 announced)
  Nexthop: 111.222.5.254
  Localpref: 100
  AS path: I <Originator>
  Cluster list: 10.2.3.1
  Originator ID: 10.255.245.45
195.1.2.0/24 (1 entry, 1 announced)
  Nexthop: 111.222.5.254
  Localpref: 100
  AS path: I <Originator>
  Cluster list: 10.2.3.1
  Originator ID: 10.255.245.68
inet.2: 63 destinations, 63 routes (63 active, 0 holddown, 0 hidden)
Prefix      Nexthop      MED    Lclpref AS path
inet.3: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix      Nexthop      MED    Lclpref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix      Nexthop      MED    Lclpref AS path
mpls.0: 48 destinations, 48 routes (48 active, 0 holddown, 0 hidden)

```

show route resolution

| | |
|------------------------------------|--|
| Syntax | <pre>show route resolution <brief detail extensive summary> <index <i>index</i>> <logical-system (all <i>logical-system-name</i>)> <prefix> <table <i>routing-table-name</i>> <unresolved></pre> |
| Syntax (EX Series Switches) | <pre>show route resolution <brief detail extensive summary> <index <i>index</i>> <prefix> <table <i>routing-table-name</i>> <unresolved></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> |
| Description | <p>Display the entries in the next-hop resolution database. This database provides for recursive resolution of next hops through other prefixes in the routing table.</p> |
| Options | <p>none—Display standard information about all entries in the next-hop resolution database.</p> <p>brief detail extensive summary—(Optional) Display the specified level of output.</p> <p>index <i>index</i>—(Optional) Show the index of the resolution tree.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>prefix <i>network/destination-prefix</i>—(Optional) Display database entries for the specified address.</p> <p>table <i>routing-table-name</i>—(Optional) Display information about a particular routing table (for example, inet.0) where policy-based export is currently enabled.</p> <p>unresolved—(Optional) Display routes that could not be resolved.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> Example: Configuring Route Resolution on PE Routers |
| List of Sample Output | <p>show route resolution detail on page 643</p> <p>show route resolution summary on page 643</p> <p>show route resolution unresolved on page 643</p> |

Output Fields Table 141 on page 642 describes the output fields for the **show route resolution** command. Output fields are listed in the approximate order in which they appear.

Table 141: show route resolution Output Fields

| Field Name | Field Description |
|------------------------------------|---|
| <i>routing-table-name</i> | Name of the routing table whose prefixes are resolved using the entries in the route resolution database. For routing table groups, this is the name of the primary routing table whose prefixes are resolved using the entries in the route resolution database. |
| Tree index | Tree index identifier. |
| Nodes | Number of nodes in the tree. |
| Reference count | Number of references made to the next hop. |
| Contributing routing tables | Routing tables used for next-hop resolution. |
| Originating RIB | Name of the routing table whose active route was used to determine the forwarding next-hop entry in the resolution database. For example, in the case of inet.0 resolving through inet.0 and inet.3 , this field indicates which routing table, inet.0 or inet.3 , provided the best path for a particular prefix. |
| Metric | Metric associated with the forwarding next hop. |
| Node path count | Number of nodes in the path. |
| Forwarding next hops | Number of forwarding next hops. The forwarding next hop is the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it. |

Sample Output

show route resolution detail

```
user@host> show route resolution detail
Tree Index: 1, Nodes 0, Reference Count 1
Contributing routing tables: inet.3
Tree Index: 2, Nodes 23, Reference Count 1
Contributing routing tables: inet.0 inet.3
10.10.0.0/16 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
10.31.1.0/30 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
10.31.1.1/32 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 0
10.31.1.4/30 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
10.31.1.5/32 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 0
10.31.2.0/30 Originating RIB: inet.0
  Metric: 2 Node path count: 1
  Forwarding nexthops: 2
10.31.11.0/24 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
```

show route resolution summary

```
user@host> show route resolution summary
Tree Index: 1, Nodes 24, Reference Count 1
Contributing routing tables: :voice.inet.0 :voice.inet.3
Tree Index: 2, Nodes 2, Reference Count 1
Contributing routing tables: inet.3
Tree Index: 3, Nodes 43, Reference Count 1
Contributing routing tables: inet.0 inet.3
```

show route resolution unresolved

```
user@host> show route resolution unresolved
Tree Index 1
vt-3/2/0.32769.0      /16
  Protocol Nexthop: 10.255.71.238 Push 800000
  Indirect nexthop: 0 -
vt-3/2/0.32772.0      /16
  Protocol Nexthop: 10.255.70.103 Push 800008
  Indirect nexthop: 0 -
Tree Index 2
```

show route snooping

| | |
|---------------------------------|---|
| Syntax | <code>show route snooping</code> <code><brief detail extensive terse></code> <code><all></code> <code><best address/prefix></code> <code><exact address></code> <code><range prefix-range></code> <code><summary></code> <code><table table-name></code> |
| Release Information | Command introduced in Junos OS Release 8.5. Command introduced in Junos OS Release 9.0 for EX Series switches. |
| Description | Display the entries in the routing table that were learned from snooping. |
| Options | <p>none—Display the entries in the routing table that were learned from snooping.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>all—(Optional) Display all entries, including hidden entries.</p> <p>best address/prefix—(Optional) Display the longest match for the provided address and optional prefix.</p> <p>exact address/prefix—(Optional) Display exact matches for the provided address and optional prefix.</p> <p>range prefix-range—(Optional) Display information for the provided address range.</p> <p>summary—(Optional) Display route snooping summary statistics.</p> <p>table table-name—(Optional) Display information for the named table.</p> |
| Required Privilege Level | view |
| List of Sample Output | show route snooping detail on page 645 |
| Output Fields | For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command. |

Sample Output

show route snooping
detail

```
user@host> show route snooping detail
__+domainAll__.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
  224.0.0.2/32 (1 entry, 1 announced)
    *IGMP Preference: 0
      Next hop type: MultiRecv
      Next-hop reference count: 4
      State: <Active NoReadvrt Int>
      Age: 2:24
      Task: IGMP
      Announcement bits (1): 0-KRT
      AS path: I

  224.0.0.22/32 (1 entry, 1 announced)
    *IGMP Preference: 0
      Next hop type: MultiRecv
      Next-hop reference count: 4
      State: <Active NoReadvrt Int>
      Age: 2:24
      Task: IGMP
      Announcement bits (1): 0-KRT
      AS path: I

__+domainAll__.inet.1: 36 destinations, 36 routes (36 active, 0 holddown, 0 hidden)
  224.0.0.0.0.0.0.0.0/24 (1 entry, 1 announced)
    *Multicast Preference: 180
      Next hop type: Multicast (IPv4), Next hop index: 1048584
      Next-hop reference count: 4
      State: <Active Int>
      Age: 2:24
      Task: MC
      Announcement bits (1): 0-KRT
      AS path: I

  225.0.0.2.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
    *Multicast Preference: 180
      Next hop type: Multicast (IPv4)
      Next-hop reference count: 113
      State: <Active Int>
      Age: 2:13
      Task: MC
      Announcement bits (1): 0-KRT
      AS path: I

  225.0.0.3.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
    *Multicast Preference: 180
      Next hop type: Multicast (IPv4)
      Next-hop reference count: 113
      State: <Active Int>
      Age: 2:15
      Task: MC
      Announcement bits (1): 0-KRT
      AS path: I

  225.0.0.4.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
    *Multicast Preference: 180
      Next hop type: Multicast (IPv4)
      Next-hop reference count: 113
```

```
State: <Active Int>
Age: 2:17
Task: MC
Announcement bits (1): 0-KRT
AS path: I

225.0.0.5.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 1:58
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.6.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:14
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.7.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:12
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.9.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:13
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.10.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

226.0.0.1.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
```



```
Age: 2:09
Task: MC
Announcement bits (1): 0-KRT
AS path: I

226.0.0.2.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 8
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

226.0.0.4.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:10
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

226.0.0.8.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:12
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

226.0.0.10.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 1:56
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.1.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:10
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.2.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:13
```

```
Task: MC
Announcement bits (1): 0-KRT
AS path: I

227.0.0.3.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:16
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.4.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.5.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 1:57
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.7.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 1:57
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.8.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:10
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.10.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
```

```
Announcement bits (1): 0-KRT
AS path: I

228.0.0.1.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:09
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.2.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:18
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.7.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:11
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.8.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:17
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.9.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 8
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.10.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:12
    Task: MC
    Announcement bits (1): 0-KRT
```

```
AS path: I

229.0.0.3.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:09
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

229.0.0.4.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:12
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

229.0.0.5.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 9
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

229.0.0.6.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

229.0.0.7.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

229.0.0.8.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I
```

229.0.0.9.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)

*Multicast Preference: 180
Next hop type: Multicast (IPv4)
Next-hop reference count: 113
State: <Active Int>
Age: 2:14
Task: MC
Announcement bits (1): 0-KRT
AS path: I

229.0.0.10.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)

*Multicast Preference: 180
Next hop type: Multicast (IPv4)
Next-hop reference count: 113
State: <Active Int>
Age: 2:13
Task: MC
Announcement bits (1): 0-KRT
AS path: I

show route source-gateway

| | |
|------------------------------------|--|
| Syntax | <code>show route source-gateway <i>address</i></code> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switches) | <code>show route source-gateway <i>address</i></code> <brief detail extensive terse> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. |
| Description | Display the entries in the routing table that were learned from a particular address. The Source field in the <code>show route detail</code> command output lists the source for each route, if known. |
| Options | brief detail extensive terse —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief . <i>address</i> —IP address of the system. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | view |
| List of Sample Output | show route source-gateway on page 653 show route source-gateway detail on page 653 show route source-gateway extensive on page 655 |
| Output Fields | For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command. |

Sample Output

**show route
source-gateway**

```
user@host> show route source-gateway 10.255.70.103
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.255.70.103:1:3:1/96
    *[BGP/170] 12:12:24, localpref 100, from 10.255.70.103
    AS path: I
    > via so-0/3/0.0, label-switched-path green-r1-r3

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.255.70.103:2:3:1/96
    *[BGP/170] 12:12:24, localpref 0, from 10.255.70.103
    AS path: I
    > via so-0/3/0.0, label-switched-path green-r1-r3

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.255.70.103:1:3:1/96
    *[BGP/170] 12:12:24, localpref 100, from 10.255.70.103
    AS path: I
    > via so-0/3/0.0, label-switched-path green-r1-r3

10.255.70.103:2:3:1/96
    *[BGP/170] 12:12:24, localpref 0, from 10.255.70.103
    AS path: I
    > via so-0/3/0.0, label-switched-path green-r1-r3
```

**show route
source-gateway detail**

```
user@host> show route source-gateway 10.255.70.103 detail
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
```

Restart Complete

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

Restart Complete

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)

Restart Complete

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

Restart Complete

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

Restart Complete

10.255.70.103:1:3:1/96 (1 entry, 1 announced)

```
*BGP      Preference: 170/-101
          Route Distinguisher: 10.255.70.103:1
          Next-hop reference count: 7
          Source: 10.255.70.103
          Protocol next hop: 10.255.70.103
          Indirect next hop: 2 no-forward
          State: <Secondary Active Int Ext>
          Local AS: 69 Peer AS: 69
          Age: 12:14:00 Metric2: 1
          Task: BGP_69.10.255.70.103+179
          Announcement bits (1): 0-green-l2vpn
          AS path: I
          Communities: target:11111:1 Layer2-info: encaps:VPLS,
          control flags:, mtu: 0
          Label-base: 800008, range: 8
          Localpref: 100
          Router ID: 10.255.70.103
          Primary Routing Table bgp.l2vpn.0
```

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

Restart Complete

10.255.70.103:2:3:1/96 (1 entry, 1 announced)

```
*BGP      Preference: 170/-1
          Route Distinguisher: 10.255.70.103:2
          Next-hop reference count: 7
          Source: 10.255.70.103
          Protocol next hop: 10.255.70.103
          Indirect next hop: 2 no-forward
          State: <Secondary Active Int Ext>
          Local AS: 69 Peer AS: 69
          Age: 12:14:00 Metric2: 1
          Task: BGP_69.10.255.70.103+179
          Announcement bits (1): 0-red-l2vpn
          AS path: I
          Communities: target:11111:2 Layer2-info: encaps:VPLS,
          control flags:Site-Down, mtu: 0
          Label-base: 800016, range: 8
          Localpref: 0
          Router ID: 10.255.70.103
          Primary Routing Table bgp.l2vpn.0
```

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

Restart Complete


```

10.255.70.103:1:3:1/96 (1 entry, 0 announced)
  *BGP    Preference: 170/-101
          Route Distinguisher: 10.255.70.103:1
          Next-hop reference count: 7
          Source: 10.255.70.103
          Protocol next hop: 10.255.70.103
          Indirect next hop: 2 no-forward
          State: <Active Int Ext>
          Local AS: 69 Peer AS: 69
          Age: 12:14:00 Metric2: 1
          Task: BGP_69.10.255.70.103+179
          AS path: I
          Communities: target:11111:1 Layer2-info: encaps:VPLS, control
flags:, mtu: 0
  Label-base: 800008, range: 8
  Localpref: 100
  Router ID: 10.255.70.103
  Secondary Tables: green.l2vpn.0
10.255.70.103:2:3:1/96 (1 entry, 0 announced)
  *BGP    Preference: 170/-1
          Route Distinguisher: 10.255.70.103:2
          Next-hop reference count: 7
          Source: 10.255.70.103
          Protocol next hop: 10.255.70.103
          Indirect next hop: 2 no-forward
          State: <Active Int Ext>
          Local AS: 69 Peer AS: 69
          Age: 12:14:00 Metric2: 1
          Task: BGP_69.10.255.70.103+179
          AS path: I
          Communities: target:11111:2 Layer2-info: encaps:VPLS,
control flags:Site-Down,
mtu: 0
  Label-base: 800016, range: 8
  Localpref: 0
  Router ID: 10.255.70.103
  Secondary Tables: red.l2vpn.0

```

**show route
source-gateway
extensive**

```

user@host> show route source-gateway 10.255.70.103 extensive
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
10.255.70.103:1:3:1/96 (1 entry, 1 announced)
  *BGP    Preference: 170/-101
          Route Distinguisher: 10.255.70.103:1

```

```

Next-hop reference count: 7
Source: 10.255.70.103
Protocol next hop: 10.255.70.103
Indirect next hop: 2 no-forward
State: <Secondary Active Int Ext>
Local AS: 69 Peer AS: 69
Age: 12:15:24 Metric2: 1
Task: BGP_69.10.255.70.103+179
Announcement bits (1): 0-green-12vpn
AS path: I
Communities: target:11111:1 Layer2-info: encaps:VPLS,
control flags:, mtu: 0
Label-base: 800008, range: 8
Localpref: 100
Router ID: 10.255.70.103
Primary Routing Table bgp.12vpn.0

red.12vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:2:3:1/96 (1 entry, 1 announced)
*BGP Preference: 170/-1
Route Distinguisher: 10.255.70.103:2
Next-hop reference count: 7
Source: 10.255.70.103
Protocol next hop: 10.255.70.103
Indirect next hop: 2 no-forward
State: <Secondary Active Int Ext>
Local AS: 69 Peer AS: 69
Age: 12:15:24 Metric2: 1
Task: BGP_69.10.255.70.103+179
Announcement bits (1): 0-red-12vpn
AS path: I
Communities: target:11111:2 Layer2-info: encaps:VPLS,
control flags:Site-Down, mtu: 0
Label-base: 800016, range: 8
Localpref: 0
Router ID: 10.255.70.103
Primary Routing Table bgp.12vpn.0

bgp.12vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:1:3:1/96 (1 entry, 0 announced)
*BGP Preference: 170/-101
Route Distinguisher: 10.255.70.103:1
Next-hop reference count: 7
Source: 10.255.70.103
Protocol next hop: 10.255.70.103
Indirect next hop: 2 no-forward
State: <Active Int Ext>
Local AS: 69 Peer AS: 69
Age: 12:15:24 Metric2: 1
Task: BGP_69.10.255.70.103+179
AS path: I
Communities: target:11111:1 Layer2-info: encaps:VPLS,
control flags:, mtu: 0
Label-base: 800008, range: 8
Localpref: 100
Router ID: 10.255.70.103
Secondary Tables: green.12vpn.0

```

```

Indirect next hops: 1
  Protocol next hop: 10.255.70.103 Metric: 2
  Indirect next hop: 2 no-forward
  Indirect path forwarding next hops: 1
Next hop:      via so-0/3/0.0 weight 0x1
             10.255.70.103/32 Originating RIB: inet.3
             Metric: 2                      Node path count: 1
             Forwarding nexthops: 1
             Nexthop: via so-0/3/0.0

10.255.70.103:2:3:1/96 (1 entry, 0 announced)
  *BGP Preference: 170/-1
  Route Distinguisher: 10.255.70.103:2
  Next-hop reference count: 7
  Source: 10.255.70.103
  Protocol next hop: 10.255.70.103
  Indirect next hop: 2 no-forward
  State: <Active Int Ext>
  Local AS: 69 Peer AS: 69
  Age: 12:15:24 Metric2: 1
  Task: BGP_69.10.255.70.103+179
  AS path: I
  Communities: target:11111:2 Layer2-info: encaps:VPLS,
  control flags:Site-Down,
  mtu: 0
  Label-base: 800016, range: 8
  Localpref: 0
  Router ID: 10.255.70.103
  Secondary Tables: red.12vpn.0
  Indirect next hops: 1
    Protocol next hop: 10.255.70.103 Metric: 2
    Indirect next hop: 2 no-forward
    Indirect path forwarding next hops: 1
  Next hop:      via so-0/3/0.0 weight 0x1
                10.255.70.103/32 Originating RIB: inet.3
                Metric: 2                      Node path count: 1
                Forwarding nexthops: 1
                Nexthop: via so-0/3/0.0

```

show route summary

| | |
|------------------------------------|--|
| Syntax | <pre>show route summary <logical-system (all <i>logical-system-name</i>)> <table <i>routing-table-name</i>></pre> |
| Syntax (EX Series Switches) | show route summary |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> |
| Description | <p>Display summary statistics about the entries in the routing table.</p> <p>CPU utilization might increase while the device learns routes. We recommend that you use the show route summary command after the device learns and enters the routes into the routing table. Depending on the size of your network, this might take several minutes. If you receive a “timeout communicating with routing daemon” error when using the show route summary command, wait several minutes before attempting to use the command again. This is not a critical system error, but you might experience a delay in using the command-line interface (CLI).</p> |
| Options | <p>none—Display summary statistics about the entries in the routing table.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>table <i>routing-table-name</i>—(Optional) Display summary statistics for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the show route summary table inet command). If you only want to display statistics for a specific routing table, make sure to enter the exact name of that routing table.</p> |
| Required Privilege Level | view |
| List of Sample Output | <p>show route summary on page 660</p> <p>show route summary table on page 660</p> <p>show route summary table (with Route Limits Configured for the Routing Table) on page 661</p> |
| Output Fields | <p>Table 142 on page 658 lists the output fields for the show route summary command. Output fields are listed in the approximate order in which they appear.</p> |

Table 142: show route summary Output Fields

| Field Name | Field Description |
|----------------------------------|--|
| Router ID | Address of the local routing device. |
| <i>routing-table-name</i> | Name of the routing table (for example, inet.0). |

Table 142: show route summary Output Fields (*continued*)

| Field Name | Field Description |
|------------------------|---|
| destinations | Number of destinations for which there are routes in the routing table. |
| routes | Number of routes in the routing table: <ul style="list-style-type: none"> • active—Number of routes that are active. • holddown—Number of routes that are in the hold-down state before being declared inactive. • hidden—Number of routes that are not used because of routing policy. |
| Limit/Threshold | Displays the configured route limits for the routing table set with the maximum-prefixes and the maximum-paths statements. If you do not configure route limits for the routing table, the show output does not display this information. <ul style="list-style-type: none"> • destinations—The first number represents the maximum number of route prefixes installed in the routing table. The second number represents the number of route prefixes that trigger a warning message. • routes—The first number represents the maximum number of routes. The second number represents the number of routes that trigger a warning message. |
| Direct | Routes on the directly connected network. |
| Local | Local routes. |
| protocol-name | Name of the protocol from which the route was learned. For example, OSPF , RSVP , and Static . |

Sample Output

```

show route summary  user@host> show route summary
Autonomous system number: 69
Router ID: 10.255.71.52
Maximum-ECMP: 32
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
    Direct:      6 routes,      5 active
    Local:      4 routes,      4 active
    OSPF:       5 routes,      4 active
    Static:     7 routes,      7 active
    IGMP:       1 routes,      1 active
    PIM:        2 routes,      2 active

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
    RSVP:       2 routes,      2 active

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete
    Direct:      1 routes,      1 active

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
    MPLS:       3 routes,      3 active
    VPLS:       4 routes,      2 active

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete
    Direct:      2 routes,      2 active
    PIM:         2 routes,      2 active
    MLD:         1 routes,      1 active

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
    BGP:        2 routes,      2 active
    L2VPN:      2 routes,      2 active

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete
    BGP:        2 routes,      2 active
    L2VPN:      1 routes,      1 active

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
    BGP:        4 routes,      4 active

```

```

show route summary  user@host> show route summary table inet
table               Router ID: 192.168.0.1

```

```

inet.0: 32 destinations, 34 routes (31 active, 0 holddown, 1 hidden)
    Direct:      6 routes,      5 active
    Local:      9 routes,      9 active
    OSPF:       3 routes,      1 active
    Static:    13 routes,     13 active
    IGMP:       1 routes,      1 active
    PIM:        2 routes,      2 active

```

```

inet.1: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
      Multicast:      1 routes,      1 active

inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
      Local:      1 routes,      1 active
      PIM:      2 routes,      2 active

inet6.1: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
      Multicast:      1 routes,      1 active

```

**show route summary
table (with Route
Limits Configured for
the Routing Table)**

```

user@host> show route summary table VPN-A.inet.0
Autonomous system number: 100
Router ID: 10.255.182.142

VPN-A.inet.0: 13 destinations, 14 routes (13 active, 0 holddown, 0 hidden)
Limit/Threshold: 2000/200 destinations 20/12 routes
      Direct:      2 routes,      2 active
      Local:      1 routes,      1 active
      OSPF:      4 routes,      3 active
      BGP:      4 routes,      4 active
      IGMP:      1 routes,      1 active
      PIM:      2 routes,      2 active

```

show route table

| | |
|------------------------------------|---|
| Syntax | <pre>show route table <i>routing-table-name</i> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)></pre> |
| Syntax (EX Series Switches) | <pre>show route table <i>routing-table-name</i> <brief detail extensive terse></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> |
| Description | Display the route entries in a particular routing table. |
| Options | <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><i>routing-table-name</i>—Display route entries for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the show route table inet command).</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • show route summary on page 658 |
| List of Sample Output | <p>show route table bgp.l2.vpn on page 664</p> <p>show route table bgp.l3vpn.0 on page 664</p> <p>show route table bgp.l3vpn.0 detail on page 664</p> <p>show route table bgp.rtarget.0 (When Proxy BGP Route Target Filtering Is Configured) on page 665</p> <p>show route table inet.0 on page 666</p> <p>show route table inet6.0 on page 666</p> <p>show route table inet6.3 on page 666</p> <p>show route table inetflow detail on page 666</p> <p>show route table l2circuit.0 on page 667</p> <p>show route table mpls on page 667</p> <p>show route table mpls extensive on page 668</p> <p>show route table mpls.0 on page 668</p> <p>show route table mpls.0 (RSVP Route—Transit LSP) on page 668</p> <p>show route table vpls_1 detail on page 669</p> <p>show route table vpn-a on page 669</p> <p>show route table vpn-a.mdt.0 on page 669</p> <p>show route table VPN-A detail on page 670</p> <p>show route table VPN-AB.inet.0 on page 670</p> <p>show route table VPN_blue.mvpn-inet6.0 on page 670</p> <p>show route table VPN-A detail on page 671</p> |

[show route table inetflow detail on page 671](#)

Output Fields For information about output fields, see the output field tables for the [show route](#) command, the [show route detail](#) command, the [show route extensive](#) command, or the [show route terse](#) command.

Sample Output

show route table bgp.l2vpn

```
user@host> show route table bgp.l2vpn
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.24.1:1:4:1/96
    *[BGP/170] 01:08:58, localpref 100, from 192.168.24.1
    AS path: I
    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
```

show route table bgp.l3vpn.0

```
user@host> show route table bgp.l3vpn.0
bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.71.15:100:10.255.71.17/32
    *[BGP/170] 00:03:59, MED 1, localpref 100, from
10.255.71.15
    AS path: I
    > via so-2/1/0.0, Push 100020, Push 100011(top)
10.255.71.15:200:10.255.71.18/32
    *[BGP/170] 00:03:59, MED 1, localpref 100, from
10.255.71.15
    AS path: I
    > via so-2/1/0.0, Push 100021, Push 100011(top)
```

show route table bgp.l3vpn.0 detail

```
user@host> show route table bgp.l3vpn.0 detail
bgp.l3vpn.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)

10.255.245.12:1:4.0.0.0/8 (1 entry, 1 announced)
  *BGP Preference: 170/-101
    Route Distinguisher: 10.255.245.12:1
    Source: 10.255.245.12
    Next hop: 192.168.208.66 via fe-0/0/0.0, selected
    Label operation: Push 182449
    Protocol next hop: 10.255.245.12
    Push 182449
    Indirect next hop: 863a630 297
    State: <Active Int Ext>
    Local AS: 35 Peer AS: 35
    Age: 12:19 Metric2: 1
    Task: BGP_35.10.255.245.12+179
    Announcement bits (1): 0-BGP.0.0.0.0+179
    AS path: 30 10458 14203 2914 3356 I (Atomic) Aggregator: 3356 4.68.0.11

    Communities: 2914:420 target:11111:1 origin:56:78
    VPN Label: 182449
    Localpref: 100
    Router ID: 10.255.245.12

10.255.245.12:1:4.17.225.0/24 (1 entry, 1 announced)
  *BGP Preference: 170/-101
    Route Distinguisher: 10.255.245.12:1
    Source: 10.255.245.12
    Next hop: 192.168.208.66 via fe-0/0/0.0, selected
    Label operation: Push 182465
    Protocol next hop: 10.255.245.12
    Push 182465
```

```

Indirect next hop: 863a8f0 305
State: <Active Int Ext>
Local AS: 35 Peer AS: 35
Age: 12:19 Metric2: 1
Task: BGP_35.10.255.245.12+179
Announcement bits (1): 0-BGP.0.0.0.0+179
AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496 I
Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
VPN Label: 182465
Localpref: 100
Router ID: 10.255.245.12

10.255.245.12:1:4.17.226.0/23 (1 entry, 1 announced)
*BGP Preference: 170/-101
Route Distinguisher: 10.255.245.12:1
Source: 10.255.245.12
Next hop: 192.168.208.66 via fe-0/0/0.0, selected
Label operation: Push 182465
Protocol next hop: 10.255.245.12
Push 182465
Indirect next hop: 86bd210 330
State: <Active Int Ext>
Local AS: 35 Peer AS: 35
Age: 12:19 Metric2: 1
Task: BGP_35.10.255.245.12+179
Announcement bits (1): 0-BGP.0.0.0.0+179
AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496
6496 I
Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
VPN Label: 182465
Localpref: 100
Router ID: 10.255.245.12

10.255.245.12:1:4.17.251.0/24 (1 entry, 1 announced)
*BGP Preference: 170/-101
Route Distinguisher: 10.255.245.12:1
Source: 10.255.245.12
Next hop: 192.168.208.66 via fe-0/0/0.0, selected
Label operation: Push 182465
Protocol next hop: 10.255.245.12
Push 182465
Indirect next hop: 86bd210 330
State: <Active Int Ext>
Local AS: 35 Peer AS: 35
Age: 12:19 Metric2: 1
Task: BGP_35.10.255.245.12+179
Announcement bits (1): 0-BGP.0.0.0.0+179
AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496
6496 I
Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
VPN Label: 182465
Localpref: 100

```

**show route table
bgp.rtarget.0 (When
Proxy BGP Route**

```

user@host> show route table bgp.rtarget.0
bgp.rtarget.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

```

Target Filtering Is Configured)

```
100:100:100/96
*[RTarget/5] 00:03:14
  Type Proxy
    for 10.255.165.103
    for 10.255.166.124
  Local
```

show route table inet.0

```
user@host> show route table inet.0
inet.0: 12 destinations, 12 routes (11 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 00:51:57
                   > to 111.222.5.254 via fxp0.0
1.0.0.1/32         *[Direct/0] 00:51:58
                   > via at-5/3/0.0
1.0.0.2/32         *[Local/0] 00:51:58
                   Local
12.12.12.21/32     *[Local/0] 00:51:57
                   Reject
13.13.13.13/32     *[Direct/0] 00:51:58
                   > via t3-5/2/1.0
13.13.13.14/32     *[Local/0] 00:51:58
                   Local
13.13.13.21/32     *[Local/0] 00:51:58
                   Local
13.13.13.22/32     *[Direct/0] 00:33:59
                   > via t3-5/2/0.0
127.0.0.1/32       [Direct/0] 00:51:58
                   > via lo0.0
111.222.5.0/24     *[Direct/0] 00:51:58
                   > via fxp0.0
111.222.5.81/32    *[Local/0] 00:51:58
                   Local
```

show route table inet6.0

```
user@host> show route table inet6.0
inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Route, * = Both

fec0:0:0:3::/64   *[Direct/0] 00:01:34
>via fe-0/1/0.0

fec0:0:0:3::/128  *[Local/0] 00:01:34
>Local

fec0:0:0:4::/64   *[Static/5] 00:01:34
>to fec0:0:0:3::ffff via fe-0/1/0.0
```

show route table inet6.3

```
user@router> show route table inet6.3
inet6.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

::10.255.245.195/128
                   *[LDP/9] 00:00:22, metric 1
                   > via so-1/0/0.0
::10.255.245.196/128
                   *[LDP/9] 00:00:08, metric 1
                   > via so-1/0/0.0, Push 100008
```

show route table inetflow detail

```

user@host> show route table inetflow detail
inetflow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.12.44.1,*/48 (1 entry, 1 announced)
    *BGP    Preference: 170/-101
            Next-hop reference count: 2
            State: **Active Ext>
            Local AS: 65002 Peer AS: 65000
            Age: 4
            Task: BGP_65000.10.12.99.5+3792
            Announcement bits (1): 0-Flow
            AS path: 65000 I
            Communities: traffic-rate:0:0
            Validation state: Accept, Originator: 10.12.99.5
            Via: 10.12.44.0/24, Active
            Localpref: 100
            Router ID: 10.255.71.161

10.12.56.1,*/48 (1 entry, 1 announced)
    *Flow    Preference: 5
            Next-hop reference count: 2
            State: **Active>
            Local AS: 65002
            Age: 6:30
            Task: RT Flow
            Announcement bits (2): 0-Flow 1-BGP.0.0.0.0+179
            AS path: I
            Communities: 1:1

```

show route table l2circuit.0

```

user@host> show route table l2circuit.0
l2circuit.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.1.1.195:NoCtrlWord:1:1:Local/96
    * [L2CKT/7] 00:50:47
    > via so-0/1/2.0, Push 100049
    via so-0/1/3.0, Push 100049
10.1.1.195:NoCtrlWord:1:1:Remote/96
    * [LDP/9] 00:50:14
    Discard
10.1.1.195:CtrlWord:1:2:Local/96
    * [L2CKT/7] 00:50:47
    > via so-0/1/2.0, Push 100049
    via so-0/1/3.0, Push 100049
10.1.1.195:CtrlWord:1:2:Remote/96
    * [LDP/9] 00:50:14
    Discard

```

show route table mpls

```

user@host> show route table mpls
mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0          * [MPLS/0] 00:13:55, metric 1
           Receive
1          * [MPLS/0] 00:13:55, metric 1
           Receive
2          * [MPLS/0] 00:13:55, metric 1
           Receive
1024       * [VPN/0] 00:04:18
           to table red.inet.0, Pop

```

show route table mpls extensive

```

user@host> show route table mpls extensive
100000 (1 entry, 1 announced)
TSI:
KRT in-kerne1 100000 /36 -> {so-1/0/0.0}
    *LDP      Preference: 9
              Next hop: via so-1/0/0.0, selected
              Pop
              State: <Active Int>
              Age: 29:50      Metric: 1
              Task: LDP
              Announcement bits (1): 0-KRT
              AS path: I
              Prefixes bound to route: 10.0.0.194/32

```

show route table mpls.0

```

user@host> show route table mpls.0
mpls.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0          *[MPLS/0] 00:45:09, metric 1
           Receive
1          *[MPLS/0] 00:45:09, metric 1
           Receive
2          *[MPLS/0] 00:45:09, metric 1
           Receive
100000     *[L2VPN/7] 00:43:04
           > via so-0/1/0.1, Pop
100001     *[L2VPN/7] 00:43:03
           > via so-0/1/0.2, Pop      Offset: 4
100002     *[LDP/9] 00:43:22, metric 1
           via so-0/1/2.0, Pop
           > via so-0/1/3.0, Pop
100002(S=0) *[LDP/9] 00:43:22, metric 1
           via so-0/1/2.0, Pop
           > via so-0/1/3.0, Pop
100003     *[LDP/9] 00:43:22, metric 1
           > via so-0/1/2.0, Swap 100002
           via so-0/1/3.0, Swap 100002
100004     *[LDP/9] 00:43:16, metric 1
           via so-0/1/2.0, Swap 100049
           > via so-0/1/3.0, Swap 100049
so-0/1/0.1 *[L2VPN/7] 00:43:04
           > via so-0/1/2.0, Push 100001, Push 100049(top)
           via so-0/1/3.0, Push 100001, Push 100049(top)
so-0/1/0.2 *[L2VPN/7] 00:43:03
           via so-0/1/2.0, Push 100000, Push 100049(top) Offset: -4
           > via so-0/1/3.0, Push 100000, Push 100049(top) Offset: -4

```

show route table mpls.0 (RSVP Route—Transit LSP)

```

user@host> show route table mpls.0
mpls.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0          *[MPLS/0] 00:37:31, metric 1
           Receive
1          *[MPLS/0] 00:37:31, metric 1
           Receive
2          *[MPLS/0] 00:37:31, metric 1
           Receive

```

```

13          *[MPLS/0] 00:37:31, metric 1
            Receive
300352      *[RSVP/7/1] 00:08:00, metric 1
            > to 8.64.0.106 via ge-1/0/1.0, label-switched-path lsp1_p2p
300352(S=0) *[RSVP/7/1] 00:08:00, metric 1
            > to 8.64.0.106 via ge-1/0/1.0, label-switched-path lsp1_p2p
300384      *[RSVP/7/2] 00:05:20, metric 1
            > to 8.64.1.106 via ge-1/0/0.0, Pop
300384(S=0) *[RSVP/7/2] 00:05:20, metric 1
            > to 8.64.1.106 via ge-1/0/0.0, Pop

```

show route table vpls_1 detail

```

user@host> show route table vpls_1 detail
vpls_1.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

1.1.1.11:1000:1:1/96 (1 entry, 1 announced)
*L2VPN Preference: 170/-1
Receive table: vpls_1.l2vpn.0
Next-hop reference count: 2
State: <Active Int Ext>
Age: 4:29:47 Metric2: 1
Task: vpls_1-l2vpn
Announcement bits (1): 1-BGP.0.0.0.0+179
AS path: I
Communities: Layer2-info: encaps:VPLS, control flags:Site-Down
Label-base: 800000, range: 8, status-vector: 0xFF

```

show route table vpn-a

```

user@host> show route table vpn-a
vpn-a.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

+ = Active Route, - = Last Active, * = Both
192.168.16.1:1:1:1/96
    *[VPN/7] 05:48:27
    Discard
192.168.24.1:1:2:1/96
    *[BGP/170] 00:02:53, localpref 100, from 192.168.24.1
    AS path: I
    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
192.168.24.1:1:3:1/96
    *[BGP/170] 00:02:53, localpref 100, from 192.168.24.1
    AS path: I
    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am

```

show route table vpn-a.mdt.0

```

user@host> show route table vpn-a.mdt.0
vpn-a.mdt.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:1:0:10.255.14.216:232.1.1.1/144
    *[MVPN/70] 01:23:05, metric2 1
    Indirect
1:1:1:10.255.14.218:232.1.1.1/144
    *[BGP/170] 00:57:49, localpref 100, from 10.255.14.218
    AS path: I
    > via so-0/0/0.0, label-switched-path r0e-to-r1
1:1:2:10.255.14.217:232.1.1.1/144
    *[BGP/170] 00:57:49, localpref 100, from 10.255.14.217
    AS path: I
    > via so-0/0/1.0, label-switched-path r0-to-r2

```

show route table VPN-A detail

```

user@host> show route table VPN-A detail
VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
10.255.179.9/32 (1 entry, 1 announced)
    *BGP      Preference: 170/-101
                Route Distinguisher: 10.255.179.13:200
                Next hop type: Indirect
                Next-hop reference count: 5
                Source: 10.255.179.13
                Next hop type: Router, Next hop index: 732
                Next hop: 10.39.1.14 via fe-0/3/0.0, selected
                Label operation: Push 299824, Push 299824(top)
                Protocol next hop: 10.255.179.13
                Push 299824
                Indirect next hop: 8f275a0 1048574
                State: (Secondary Active Int Ext)
                Local AS: 1 Peer AS: 1
                Age: 3:41:06 Metric: 1 Metric2: 1
                Task: BGP_1.10.255.179.13+64309
                Announcement bits (2): 0-KRT 1-BGP RT Background
                AS path: I
                Communities: target:1:200 rte-type:0.0.0.0:1:0
                Import Accepted
                VPN Label: 299824 TTL Action: vrf-ttl-propagate
                Localpref: 100
                Router ID: 10.255.179.13
                Primary Routing Table bgp.13vpn.0

```

show route table VPN-AB.inet.0

```

user@host> show route table VPN-AB.inet.0
VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.0/30      *[OSPF/10] 00:07:24, metric 1
                  > via so-7/3/1.0
10.39.1.4/30      *[Direct/0] 00:08:42
                  > via so-5/1/0.0
10.39.1.6/32      *[Local/0] 00:08:46
                  Local
10.255.71.16/32   *[Static/5] 00:07:24
                  > via so-2/0/0.0
10.255.71.17/32   *[BGP/170] 00:07:24, MED 1, localpref 100, from
10.255.71.15
                  AS path: I
                  > via so-2/1/0.0, Push 100020, Push 100011(top)
10.255.71.18/32   *[BGP/170] 00:07:24, MED 1, localpref 100, from
10.255.71.15
                  AS path: I
                  > via so-2/1/0.0, Push 100021, Push 100011(top)
10.255.245.245/32 *[BGP/170] 00:08:35, localpref 100
                  AS path: 2 I
                  > to 10.39.1.5 via so-5/1/0.0
10.255.245.246/32 *[OSPF/10] 00:07:24, metric 1
                  > via so-7/3/1.0

```

show route table VPN_blue.mvpn-inet6.0

```

user@host> show route table VPN_blue.mvpn-inet6.0
vpn_blue.mvpn-inet6.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:10.255.2.202:65535:10.255.2.202/432

```



```

          *[BGP/170] 00:02:37, localpref 100, from 10.255.2.202
          AS path: I
          > via so-0/1/3.0
1:10.255.2.203:65535:10.255.2.203/432
          *[BGP/170] 00:02:37, localpref 100, from 10.255.2.203
          AS path: I
          > via so-0/1/0.0
1:10.255.2.204:65535:10.255.2.204/432
          *[MVPN/70] 00:57:23, metric2 1
          Indirect
5:10.255.2.202:65535:128::192.168.90.2:128:ffff::1/432
          *[BGP/170] 00:02:37, localpref 100, from 10.255.2.202
          AS path: I
          > via so-0/1/3.0
6:10.255.2.203:65535:65000:128::10.12.53.12:128:ffff::1/432
          *[PIM/105] 00:02:37
          Multicast (IPv6)
7:10.255.2.202:65535:65000:128::192.168.90.2:128:ffff::1/432
          *[MVPN/70] 00:02:37, metric2 1
          Indirect

```

show route table VPN-A detail

```

user@host> show route table VPN-A detail
VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
10.255.179.9/32 (1 entry, 1 announced)
  *BGP    Preference: 170/-101
          Route Distinguisher: 10.255.179.13:200
          Next hop type: Indirect
          Next-hop reference count: 5
          Source: 10.255.179.13
          Next hop type: Router, Next hop index: 732
          Next hop: 10.39.1.14 via fe-0/3/0.0, selected
          Label operation: Push 299824, Push 299824(top)
          Protocol next hop: 10.255.179.13
          Push 299824
          Indirect next hop: 8f275a0 1048574
          State: (Secondary Active Int Ext)
          Local AS: 1 Peer AS: 1
          Age: 3:41:06 Metric: 1 Metric2: 1
          Task: BGP_1.10.255.179.13+64309
          Announcement bits (2): 0-KRT 1-BGP RT Background
          AS path: I
          Communities: target:1:200 rte-type:0.0.0.0:1:0
          Import Accepted
          VPN Label: 299824 TTL Action: vrf-ttl-propagate
          Localpref: 100
          Router ID: 10.255.179.13
          Primary Routing Table bgp.13vpn.0

```

show route table inetflow detail

```

user@host> show route table inetflow detail
inetflow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.12.44.1,*/48 (1 entry, 1 announced)
  *BGP    Preference: 170/-101
          Next-hop reference count: 2
          State: **Active Ext>
          Local AS: 65002 Peer AS: 65000
          Age: 4
          Task: BGP_65000.10.12.99.5+3792
          Announcement bits (1): 0-Flow
          AS path: 65000 I
          Communities: traffic-rate:0:0

```

```

Validation state: Accept, Originator: 10.12.99.5
Via: 10.12.44.0/24, Active
Localpref: 100
Router ID: 10.255.71.161

10.12.56.1,*/48 (1 entry, 1 announced)
  *Flow Preference: 5
    Next-hop reference count: 2
    State: **Active>
    Local AS: 65002
    Age: 6:30
    Task: RT Flow
    Announcement bits (2): 0-Flow 1-BGP.0.0.0.0+179
    AS path: I
    Communities: 1:1

user@PE1> show route table green.l2vpn.0 (VPLS Multihoming with FEC 129)
green.l2vpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.2:100:1.1.1.2/96 AD
    *[VPLS/170] 1d 03:11:03, metric2 1
    Indirect
1.1.1.4:100:1.1.1.4/96 AD
    *[BGP/170] 1d 03:11:02, localpref 100, from 1.1.1.4
    AS path: I, validation-state: unverified
    > via ge-1/2/1.5
1.1.1.2:100:1:0/96 MH
    *[VPLS/170] 1d 03:11:03, metric2 1
    Indirect
1.1.1.4:100:1:0/96 MH
    *[BGP/170] 1d 03:11:02, localpref 100, from 1.1.1.4
    AS path: I, validation-state: unverified
    > via ge-1/2/1.5
1.1.1.4:NoCtrlWord:5:100:100:1.1.1.2:1.1.1.4/176
    *[VPLS/7] 1d 03:11:02, metric2 1
    > via ge-1/2/1.5
1.1.1.4:NoCtrlWord:5:100:100:1.1.1.4:1.1.1.2/176
    *[LDP/9] 1d 03:11:02
    Discard

```

show route terse


| | |
|------------------------------------|---|
| Syntax | show route terse <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switches) | show route terse |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. |
| Description | Display a high-level summary of the routes in the routing table. |
| | <div>  <p>NOTE: For BGP routes, the show route terse command displays the local preference attribute and MED instead of the metric1 and metric2 values. This is mostly due to historical reasons.</p> <p>To display the metric1 and metric2 value of a BGP route, use the show route extensive command.</p> </div> |
| Options | <p>none—Display a high-level summary of the routes in the routing table.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show route terse on page 675 |
| Output Fields | Table 143 on page 673 describes the output fields for the show route terse command. Output fields are listed in the approximate order in which they appear. |

Table 143: show route terse Output Fields

| Field Name | Field Description |
|----------------------------|---|
| <i>routing-table-name</i> | Name of the routing table (for example, inet.0). |
| <i>number destinations</i> | Number of destinations for which there are routes in the routing table. |
| <i>number routes</i> | Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> active (routes that are active) holddown (routes that are in the pending state before being declared inactive) hidden (routes that are not used because of a routing policy) |

Table 143: show route terse Output Fields (*continued*)

| Field Name | Field Description |
|--------------------|---|
| route key | <p>Key for the state of the route:</p> <ul style="list-style-type: none"> • +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • -—A hyphen indicates the last active route. • *—An asterisk indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route. |
| A | Active route. An asterisk (*) indicates this is the active route. |
| V | <p>Validation status of the route:</p> <ul style="list-style-type: none"> • ?—Not evaluated. Indicates that the route was not learned through BGP. • I—Invalid. Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • N—Unknown. Indicates that the prefix is not among the prefixes or prefix ranges in the database. • V—Valid. Indicates that the prefix and autonomous system pair are found in the database. |
| Destination | Destination of the route. |
| P | <p>Protocol through which the route was learned:</p> <ul style="list-style-type: none"> • A—Aggregate • B—BGP • C—CCC • D—Direct • G—GMPLS • I—IS-IS • L—L2CKT, L2VPN, LDP, Local • K—Kernel • M—MPLS, MSDP • O—OSPF • P—PIM • R—RIP, RIPng • S—Static • T—Tunnel |
| Prf | <p>Preference value of the route. In every routing metric except for the BGP LocalPref attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the LocalPref value in the Preference2 field. For example, if the LocalPref value for Route 1 is 100, the Preference2 value is -101. If the LocalPref value for Route 2 is 155, the Preference2 value is -156. Route 2 is preferred because it has a higher LocalPref value and a lower Preference2 value.</p> |
| Metric 1 | First metric value in the route. For routes learned from BGP, this is the MED metric. |
| Metric 2 | Second metric value in the route. For routes learned from BGP, this is the IGP metric. |

Table 143: show route terse Output Fields (*continued*)

| Field Name | Field Description |
|-----------------|---|
| Next hop | Next hop to the destination. An angle bracket (>) indicates that the route is the selected route. |
| AS path | <p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> • I—IGP. • E—EGP. • ?—Incomplete; typically, the AS path was aggregated. |

Sample Output

show route terse

```

user@host> show route terse
inet.0: 10 destinations, 12 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A V Destination      P Prf  Metric 1   Metric 2  Next hop      AS path
* ? 1.0.1.1/32        O  10      1           >10.0.0.2      I
?                               B 170      100           I
  unverified          >10.0.0.2
* ? 1.1.1.1/32        D   0           >1o0.2
* V 2.2.0.2/32        B 170      110          200 I
  valid              >10.0.0.2
* ? 10.0.0.0/30       D   0           >1t-1/2/0.1
?                               B 170      100           I
  unverified          >10.0.0.2
* ? 10.0.0.1/32       L   0           Local
* ? 10.0.0.4/30       B 170      100           I
  unverified          >10.0.0.2
* ? 10.0.0.8/30       B 170      100           I
  unverified          >10.0.0.2
* I 172.16.1.1/32     B 170      90          200 I
  invalid             >10.0.0.2
* N 192.168.2.3/32    B 170      100          200 I
  unknown             >10.0.0.2
* ? 224.0.0.5/32      O  10      1           MultiRecv

```


PTP Operational Mode Commands

Table 144 on page 677 summarizes the command line interface (CLI) commands you can use to monitor and troubleshoot Precision Time Protocol (PTP) sessions. Commands are listed in alphabetical order.

Table 144: PTP Operational Mode Commands

| Task | Command |
|---|--|
| Display the details of the clock configured on the node. | <code>show ptp clock</code> |
| Display PTP-related global information. | <code>show ptp global-information</code> |
| Display the current configuration and current operation mode of the slave. | <code>show ptp hybrid</code> |
| Display information about the lock status of the slave. | <code>show ptp lock-status</code> |
| Display information about the configured master and the status of the master. | <code>show ptp master</code> |
| Display information about the number of ports created according to the configuration. | <code>show ptp port</code> |
| Display information about the configured slave and the status of the slave. | <code>show ptp slave</code> |
| Display information about PTP statistics. | <code>show ptp statistics</code> |

show ptp clock

| | |
|---------------------------------|---|
| Syntax | show ptp clock |
| Release Information | Command introduced in Junos OS Release 12.2. Command introduced in Junos OS Release 12.3 for ACX Series Routers. |
| Description | (ACX Series, MX80, MX240, MX480, and MX960 routers) Display the details of the clock configured on the node. |
| Options | This command has no options. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> IEEE 1588v2 PTP Boundary Clock Overview IEEE 1588v2 Precision Timing Protocol (PTP) on ACX Series Universal Access Routers |
| List of Sample Output | show ptp clock on page 680 show ptp clock (ACX Series Routers) on page 680 |
| Output Fields | Table 145 on page 678 lists the output fields for the show ptp clock command. Output fields are listed in the approximate order in which they appear. |

Table 145: show ptp clock Output Fields

| Field Name | Field Description |
|-----------------------|--|
| Slot Number | Number of the FPC or MIC slot. |
| Two-step Clock | Whether the clock provides time information which is a combination of an event message and a subsequent general message: True or False . |
| Clock Identity | Clock identity of the slave or client as defined in IEEE 1588. |
| Total Ports on Device | Total number of PTP ports on the router. |
| Clock Class | Attribute of an ordinary or boundary clock that denotes the traceability of the time or frequency distributed by the grandmaster clock. |
| Clock Accuracy | Indicates the expected accuracy of a clock when it is the grandmaster, or in the event it becomes the grandmaster. |
| Log Variance | Represents an estimate of the variations of the local clock when it is not synchronized via PTP to another clock. |
| Clock Priority1 | Priority value of the clock. Lower value takes precedence. |
| Clock Priority2 | Prioritize the masters to avoid confusion when the Clock Priority1 value is the same for different masters in a network. |

Table 145: show ptp clock Output Fields (*continued*)

| Field Name | Field Description |
|-------------------------------|---|
| UTC Offset | Offset between International Atomic Time (TAI) and Coordinated Universal Time (UTC) times. The value is 34 seconds as of January 2012. |
| Leap59 | When TRUE , the last minute of the current UTC day has only 59 seconds (instead of the 60 SI seconds). |
| Leap61 | When TRUE , the last minute of the current UTC day has 61 seconds (instead of the 60 SI seconds). |
| Time Traceable | When TRUE , the timescale and the UTC offset are traceable to a primary reference. |
| Frequency Traceable | When TRUE , frequency determining the timescale is traceable to a primary reference. |
| Time Source | Time source external to the Precision Time Protocol (PTP), which provides time and/or frequency as appropriate. The time source is traceable to the international standards laboratories maintaining clocks that form the basis for the International Atomic Time (TAI) and Universal Coordinated Time (UTC) timescales. Examples of these are Global Positioning System (GPS), NTP, and National Institute of Standards and Technology (NIST) timeservers. |
| Delay Req Sending Time | Interval in seconds between the delay-request messages sent by the slave to the master. |
| Steps Removed | Number of boundary clocks between the local clock and the foreign master clock. |
| Slave-only | Set to TRUE , when the system is used in ordinary slave clock mode; otherwise, FALSE . |
| Parent Id | EUI-64 clock identifier of the immediate upstream master clock. |
| GMC Id | EUI-64 clock identifier of the grandmaster clock. |
| GMC Class | Denotes the grandmaster clock's traceability of the distributed time or frequency. |
| GMC Accuracy | Indicates the expected accuracy of the grandmaster clock. |
| GMC Variance | Represents an estimate of the variations of the grandmaster clock. |
| GMC Priority1 | Priority1-value of the grandmaster clock. |
| GMC Priority2 | Priority2-value of the grandmaster clock. |

Sample Output

show ptp clock

```

user@host> run show ptp clock
Clock Details:

Slot Number           : 7
Default Data:
Two-step Clock        : FALSE
00:05:85:ff:fe:73:ef:d0
Total Ports on Device : 0
Clock Accuracy        : 49
Clock Priority1       : 128
UTC Offset            : 33
Leap61                : FALSE
Frequency Traceable   : FALSE
Delay Req Sending Time: 0
Slave-only            : NA
Parent Data:
Parent Id             : 00:18:0b:ff:ff:20:01:62
GMC Id               : 00:18:0b:ff:ff:20:01:62
GMC Accuracy         : 254
GMC Priority1        : 0
Global Data:
UTC Offset           : 34
Leap-61             : FALSE
Freq Traceable      : FALSE
Time master         : 160

Clock Identity :
Clock Class      : 255
Log Variance    : -12944
Clock Priority2  : 128
Leap59          : FALSE
Time Traceable  : FALSE
Time master     : 0
Steps Removed   : 1

GMC Class       : 52
GMC Variance    : 11952
GMC Priority2    : 0

Leap-59        : FALSE
Time traceable : FALSE
Time Scale     : FALSE

```

show ptp clock (ACX Series Routers)

```

user@host> run show ptp clock
Clock Details:

Slot Number           : 0
Default Data:
Two-step Clock        : FALSE
84:18:88:ff:fe:c0:7a:00
Total Ports on Device : 0
Clock Accuracy        : 34
Clock Priority1       : 128
UTC Offset            : 0
Leap61                : FALSE
Frequency Traceable   : FALSE
Delay Req Sending Time: 0
Slave-only            : NA
Parent Data:
Parent Id             : 00:00:64:ff:fe:01:01:02
GMC Id               : 00:00:64:ff:fe:01:01:02
GMC Accuracy         : 35
GMC Priority1        : 128
Global Data:
UTC Offset           : 0
Leap-61             : FALSE
Freq Traceable      : FALSE
Time source         : 16

Clock Identity :
Clock Class      : 255
Log Variance    : 15353
Clock Priority2  : 128
Leap59          : FALSE
Time Traceable  : FALSE
Time Source     : 0
Steps Removed   : 0

GMC Class       : 80
GMC Variance    : 0
GMC Priority2    : 128

Leap-59        : FALSE
Time tracable  : FALSE
Time Scale     : FALSE

```

show ptp hybrid

| | |
|---------------------------------|---|
| Syntax | show ptp hybrid <config status> |
| Release Information | Command introduced in Junos OS Release 12.2R2. |
| Description | Display the current configuration and current operation mode of the slave. |
| Options | config —Display the PTP source to Synchronous Ethernet interface mappings. status —Display the current hybrid mode operational status. |
| Required Privilege Level | View |
| Output Fields | Table 146 on page 681 lists the output fields for the show ptp hybrid command. Output fields are listed in the approximate order in which they appear. |

Table 146: show ptp hybrid Output Fields

| Field Name | Field Description |
|---------------------------------------|---|
| ptp source | Displays the IP address of the PTP source. |
| sync source | Displays the interface name of the Synchronous Ethernet source through which the PTP source is traceable. |
| Configured Mode | Displays the current configured mode of the router as Hybrid . |
| Operating Mode | Displays the current operation mode: Hybrid or None . |
| PTP Reference | Displays the IP address and the interface name of the PTP reference clock. |
| Synchronous Ethernet Reference | Displays the interface name of the Synchronous Ethernet reference clock. |
| Lock state | Displays the current lock state of the router: Locked , Initializing , or Acquiring . |
| Lock state description | Displays the description for the current lock state of the router: <ul style="list-style-type: none"> • Initializing—Hybrid mode is being initialized. • Acquiring Frequency—Synchronous Ethernet source identified for frequency synchronization, acquiring frequency-related data from master clock. • Frequency Locked, Acquiring Phase—Frequency locked from the Synchronous Ethernet source, acquiring phase-related data from master clock. • Frequency and Phase Locked—Slave clock is frequency and phase synchronized with master clock. |

Sample Output

show ptp hybrid config user@host> show ptp hybrid config
ptp source sync source
100.1.1.2 ge-1/1/2

show ptp hybrid status user@host> show ptp hybrid status
Hybrid Mode Status:
Configured Mode : Hybrid
Operating Mode : Hybrid
PTP Reference : 100.1.1.2, ge-1/0/0.0
Synchronous Ethernet Reference : ge-1/1/2
Lock state : Locked
Lock state description : Frequency and Phase Locked

show ptp lock-status

| | |
|---------------------------------|--|
| Syntax | show ptp lock-status |
| Release Information | Command introduced in Junos OS Release 12.2. |
| Description | (ACX Series, MX80, MX240, MX480, and MX960 routers) Display information about the lock status of the slave. The output verifies whether the slave is aligned to the grandmaster (master clock) or not. |
| Options | detail —Display detailed information about the lock status of the slave. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> IEEE 1588v2 PTP Boundary Clock Overview IEEE 1588v2 Precision Timing Protocol (PTP) on ACX Series Universal Access Routers |
| List of Sample Output | show ptp lock-status on page 684 show ptp lock-status (ACX Series) on page 684 show ptp lock-status detail (ACX Series) on page 684 |
| Output Fields | Table 147 on page 683 lists the output fields for the show ptp lock-status command. Output fields are listed in the approximate order in which they appear. |

Table 147: show ptp lock-status Output Fields

| Field Name | Field Description |
|--------------------------------|--|
| Lock State | <p>State of the slave clock with respect to its master clock:</p> <ul style="list-style-type: none"> Freerun Holdover Phase Aligned Acquiring Initializing Freq locked |
| Phase offset | Time offset information of a slave clock with respect to its master clock. Precision of this time offset is 1 nanosecond. |
| Selected Master Details | <p>Details include the following:</p> <ul style="list-style-type: none"> Upstream Master address—The address of the remote master from which the slave acquires the clock. Slave interface—The slave interface on this router corresponding to the Master above. |

Sample Output

```
show ptp lock-status      user@host> run show ptp lock-status
Lock Status:

Lock State      : 5 (PHASE ALIGNED)
Phase offset    : 0.000000001 sec

show ptp lock-status      user@host> show ptp lock-status
(ACX Series)             Lock Status:

Lock State      : 1 (FREERUN)
Phase offset    : 0.000000869 sec

show ptp lock-status      user@host> show ptp lock-status detail
detail (ACX Series)       Lock Status:

Lock State      : 5 (PHASE ALIGNED)
Phase offset    : 0.000000030 sec

Selected Master Details:
Upstream Master address : 13.13.13.1
Slave interface        : ge-0/1/5.0
```

show ptp master

| | |
|---------------------------------|---|
| Syntax | show ptp master <brief detail interface> |
| Release Information | Command introduced in Junos OS Release 12.2. |
| Description | (MX80, MX240, MX480, and MX960 routers) Display information about the configured master and the status of the master. |
| Options | brief —Display information about the master in brief. detail —Display information about the master in detail. interface —Display information about the configured interface of the master. |
| Required Privilege Level | View |
| Output Fields | Table 148 on page 685 lists the output fields for the show ptp master command. Output fields are listed in the approximate order in which they appear. |

Table 148: show ptp master Output Fields

| Field Name | Field Description |
|--|--|
| Interface | Name of the interface configured for Precision Time Protocol (PTP) on the master. |
| Status | Status of the Precision Time Protocol master: <ul style="list-style-type: none"> • Master or Slave • Active or Inactive • Initializing or Down |
| Local IP | IP address of the configured master clock. |
| Status (Local IP Address Status) | Status of the local IP address of the interface: <ul style="list-style-type: none"> • Configured or Not configured • Master or Slave • Active or Inactive |
| Total Remote Slaves | Number of remote slaves. |
| Slave IP | IP address of the slave. |
| Status (Slave IP Address Status) | Status of the IP address of the slave: <ul style="list-style-type: none"> • Configured or Not configured • Master or Slave • Active or Inactive or Ready |

Sample Output

show ptp master

```
user@host> run show ptp master brief
PTP Master Interface Configured:

Master Interface      Status
ge-7/0/2.0           Master, Active
```

show ptp master detail

```
user@host> run show ptp master detail
PTP Master Interface Details:
Interface   : ge-7/0/2.0
Status      : Master, Active
Clock Info :
  Local IP: 10.0.0.1           Status: Configured, Master, Active
  Total Remote Slaves: 0
  Slave IP: 10.0.0.2          Status: Configured, Slave, Active
```

show ptp interface ge-7/0/2.0

```
user@host> run show ptp master interface ge-7/0/2.0
PTP Master Interface Configured:

Master Interface      Status
ge-7/0/2.0           Master, Active
```


show ptp port

| | |
|---------------------------------|--|
| Syntax | <code>show ptp port</code> <code><brief detail></code> |
| Release Information | Command introduced in Junos OS Release 12.2. |
| Description | (MX80, MX240, MX480, and MX960 routers) Display information about the number of ports created according to the configuration. For each unique local IP address, one Precision Time Protocol port is created. |
| Options | brief —Display information about the PTP port in brief. detail —Display information about the PTP port in detail. |
| Required Privilege Level | View |
| Output Fields | Table 149 on page 687 lists the output fields for the show ptp port command. Output fields are listed in the approximate order in which they appear. |

Table 149: show ptp port Output Fields

| Field Name | Field Description |
|--------------------|--|
| Local IP | IP address of the interface acting as the slave. |
| Remote IP | IP address of the remote node. |
| Clock Stream | Unique index for each session created. |
| Clock Identity | IP address of the slave. |
| Port State | Status of the port: PTP listening or PTP initializing . |
| Delay Req Interval | Interval in seconds between the delay request messages sent by the slave to the master. |
| Announce Interval | Logarithmic mean interval for the announce messages to be sent by the master. |
| Announce Timeout | Number of times the announce interval message has to pass between the slave and the master without receipt of an announce message. |
| Sync Interval | Logarithmic mean interval for sync interval messages to be sent by the master. |
| Delay Mechanism | Type of delay mechanism used. |
| Port Number | PTP port number. |
| Operating Mode | Clock mode of the node. |
| Master Clock ID | Unique clock-identity of the master. |

Table 149: show ptp port Output Fields (*continued*)

| Field Name | Field Description |
|----------------------------|-----------------------------|
| Previous Announce Messages | Previous announce messages. |
| Current Announce Message | Current announce messages. |

Sample Output

show ptp port brief

```

user@host> run show ptp port brief
PTP port-data:
Local IP      : 10.0.0.1      Remote IP      : 10.0.0.2
Clock Stream  : 1            Clock Identity  : 00:05:85:ff:fe:73:ef:d0
Port State    : Listening     Delay Req Interval: -4
Announce Interval : 1        Announce Timeout : 3
Sync Interval : -6           Delay Mechanism  : End-to-end
Port Number   : 2            Operating Mode   : Master only

Local IP      : 10.10.1.10    Remote IP      : 10.10.1.2
Clock Stream  : 0            Clock Identity  : 00:05:85:ff:fe:73:ef:d0
Port State    : Listening     Delay Req Interval: -4
Announce Interval : 1        Announce Timeout : 3
Sync Interval : -6           Delay Mechanism  : End-to-end
Port Number   : 1            Operating Mode   : BMC Mode

```

show ptp port detail

```

user@host> run show ptp port detail
PTP port-data:
Local IP      : 10.0.0.1      Remote IP      : 10.0.0.2
Clock Stream  : 1            Clock Identity  : 00:05:85:ff:fe:73:ef:d0
Port State    : Listening     Delay Req Interval: -4
Announce Interval : 1        Announce Timeout : 3
Sync Interval : -6           Delay Mechanism  : End-to-end
Port Number   : 2            Operating Mode   : Master only

Local IP      : 10.10.1.10    Remote IP      : 10.10.1.2
Clock Stream  : 0            Clock Identity  : 00:05:85:ff:fe:73:ef:d0
Port State    : Listening     Delay Req Interval: -4
Announce Interval : 1        Announce Timeout : 3
Sync Interval : -6           Delay Mechanism  : End-to-end
Port Number   : 1            Operating Mode   : BMC Mode

Foreign Master Clock Details:
Master Clock Id      : 00:18:0b:ff:ff:20:01:62
Previous Announce Messages : 8
Current Announce Messages   : 1

```

show ptp slave

| | |
|---------------------------------|---|
| Syntax | show ptp slave <brief detail interface> |
| Release Information | Command introduced in Junos OS Release 12.2. |
| Description | (MX80, MX240, MX480, and MX960 routers) Display information about the configured slave and the status of the slave. |
| Options | brief —Display information about the slave in detail. detail —Display information about the slave in detail. interface —Display information about the configured interface of the slave. |
| Required Privilege Level | View |
| Output Fields | Table 150 on page 689 lists the output fields for the show ptp slave command. Output fields are listed in the approximate order in which they appear. |

Table 150: show ptp slave Output Fields

| Field Name | Field Description |
|---|--|
| Interface | Name of the interface configured for Precision Time Protocol. |
| Status | Status of the Precision Time Protocol slave: <ul style="list-style-type: none"> • Master or Slave • Active or Inactive • Initializing or Down |
| Interface | Interface configured on the slave. |
| Local IP | IP address of the local interface. |
| Status (Local IP address Status) | Status of the IP address of the interface acting as the slave: <ul style="list-style-type: none"> • Configured or Unconfigured • Master or Slave • Active or Inactive or Ready |
| Total Remote Masters | Number of remote masters. |
| Remote Master | IP address of the remote node. |

Table 150: show ptp slave Output Fields (*continued*)

| Field Name | Field Description |
|---|---|
| Status (Slave IP Address Status) | Status of the IP address of the master: <ul style="list-style-type: none"> • Configured or Unconfigured • Master or Slave • Active or Inactive |

Sample Output

show ptp slave

```

user@host> run show ptp slave
PTP Slave Interfaces Configured:

Slave Interface      Status
ge-7/0/0.0           Slave, Active

```

show ptp slave detail

```

user@host> run show ptp slave detail
PTP Slave Interface Details:

Interface           : ge-7/0/0.0
Status              : Slave, Active
Clock Info
  Local IP : 10.10.1.10           Status: Configured, Slave, Active
  Total Remote Masters: 0
  Remote Master: 10.10.1.2       Status: Configured, Master, Active

```

RIP Operational Mode Commands

Table 151 on page 691 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot the Routing Information Protocol (RIP). Commands are listed in alphabetical order.

Table 151: RIP Operational Mode Commands

| Task | Command |
|---|---|
| Clear RIP general statistics. | <code>clear rip general-statistics</code> |
| Clear RIP statistics. | <code>clear rip statistics</code> |
| Display brief RIP statistics. | <code>show rip general-statistics</code> |
| Display information about RIP neighbors. | <code>show rip neighbor</code> |
| Display RIP statistics about messages sent and received on an interface, as well as information received through advertisements from other routers. | <code>show rip statistics</code> |



NOTE: For more RIP-related commands, such as `show route protocol`, `show route instance`, and `show route table`, see Protocol-Independent Routing Operational Mode Commands.

For information about how to configure RIP, see the *Junos Routing Protocols Configuration Guide*.

clear rip general-statistics

| | |
|---|---|
| Syntax | clear rip general-statistics <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switches and QFX Series) | clear rip general-statistics |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series. |
| Description | Clear RIP general statistics. |
| Options | none —Clear RIP general statistics. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• show rip general-statistics on page 694 |
| List of Sample Output | clear rip general-statistics on page 692 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

```
clear rip  
general-statistics      user@host> clear rip general-statistics
```

clear rip statistics

| | |
|---|--|
| Syntax | clear rip statistics <instance (all <i>instance-name</i>)> <logical-system (all <i>logical-system-name</i>)> <neighbor> <peer (all <i>address</i>)> |
| Syntax (EX Series Switches and QFX Series) | clear rip statistics <instance (all <i>instance-name</i>)> <neighbor> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series. |
| Description | Clear RIP statistics. |
| Options | <p>none—Reset RIP counters for all neighbors for all routing instances.</p> <p>instance (all <i>instance-name</i>)—(Optional) Clear RIP statistics for all instances or for the specified routing instance only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>neighbor—(Optional) Clear RIP statistics for the specified neighbor only.</p> <p>peer (all <i>address</i>)—(Optional) Clear RIP statistics for a single peer or all peers.</p> |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none"> • show rip statistics on page 698 |
| List of Sample Output | clear rip statistics on page 693 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

```
clear rip statistics      user@host> clear rip statistics
```

show rip general-statistics

| | |
|---|--|
| Syntax | show rip general-statistics <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switches and QFX Series) | show rip general-statistics |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series. |
| Description | Display brief RIP statistics. |
| Options | <p>none—Display brief RIP statistics.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • clear rip general-statistics on page 692 |
| List of Sample Output | show rip general-statistics on page 695 |
| Output Fields | Table 152 on page 694 lists the output fields for the show rip general-statistics command. Output fields are listed in the approximate order in which they appear. |

Table 152: show rip general-statistics Output Fields

| Field Name | Field Description |
|--------------------|--|
| bad msgs | Number of invalid messages received. |
| no rcv intf | Number of packets received with no matching interface. |
| curr memory | Amount of memory currently used by RIP. |
| max memory | Most memory used by RIP. |

Sample Output

```
show rip
general-statistics

user@host> show rip general-statistics
RIPv2 I/O info:
  bad msgs      :      0
  no recv intf  :      0
  curr memory   :      0
  max memory    :      0
```

show rip neighbor

| | |
|---|--|
| Syntax | show rip neighbor <instance (all <i>instance-name</i>)> <logical-system (all <i>logical-system-name</i>)> < <i>name</i> > |
| Syntax (EX Series Switches and QFX Series) | show rip neighbor <instance (all <i>instance-name</i>)> < <i>name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series. |
| Description | Display information about RIP neighbors. |
| Options | <p>none—Display information about all RIP neighbors for all instances.</p> <p>instance (all <i>instance-name</i>)—(Optional) Display RIP neighbor information for all instances or for only the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><i>name</i>—(Optional) Display detailed information about only the specified RIP neighbor.</p> |
| Required Privilege Level | view |
| List of Sample Output | show rip neighbor on page 697 show rip neighbor (With Demand Circuits Configured) on page 697 |
| Output Fields | Table 153 on page 696 lists the output fields for the show rip neighbor command. Output fields are listed in the approximate order in which they appear. |

Table 153: show rip neighbor Output Fields

| Field Name | Field Description |
|-----------------|--|
| Neighbor | <p>Name of the RIP neighbor.</p> <p>NOTE: Beginning with Junos OS Release 11.1, when you configure demand circuits, the output displays a demand circuit (DC) flag next to neighbor interfaces configured for demand circuits.</p> <p>If you configure demand circuits at the [edit protocols rip group <i>group-name</i> neighbor <i>neighbor-name</i>] hierarchy level, the output shows only the neighboring interface that you specifically configured as a demand circuit. If you configure demand circuits at the [edit protocols rip group <i>group-name</i>] hierarchy level, all of the interfaces in the group are configured as demand circuits. Therefore, the output shows all of the interfaces in that group as demand circuits.</p> |

Table 153: show rip neighbor Output Fields (*continued*)

| Field Name | Field Description |
|----------------------------|--|
| State | State of the connection: Up or Dn (Down). |
| Source Address | Address of the port on the local router. |
| Destination Address | Address of the port on the remote router. |
| Send Mode | Send options: broadcast , multicast , none , or version 1 . |
| Receive Mode | Type of packets to accept: both , none , version 1 , or version 2 . |
| In Met | Metric added to incoming routes when advertising into RIP routes that were learned from other protocols. |

Sample Output

show rip neighbor

```

user@host> show rip neighbor
Neighbor      Local State Source Address Destination Address Send Mode Receive Mode In Met
-----
ge-2/3/0.0    Up    192.168.9.105 192.168.9.107 bcast both      1
at-5/1/1.42   Dn    (null)        (null)        mcast v2 only    3
at-5/1/0.42   Dn    (null)        (null)        mcast both    3
at-5/1/0.0    Up    20.0.0.1      224.0.0.9     mcast both    3
so-0/0/0.0    Up    192.168.9.97 224.0.0.9     mcast both    3

```

show rip neighbor (With Demand Circuits Configured)

```

user@host> show rip neighbor
Neighbor      Local State Source Address Destination Address Send Mode Receive Mode In Met
-----
so-0/1/0.0(DC) Up    10.10.10.2    224.0.0.9     mcast both      1
so-0/2/0.0(DC) Up    13.13.13.2    224.0.0.9     mcast both      1

```

show rip statistics

| | |
|---|--|
| Syntax | <code>show rip statistics</code> <code><instance (all <i>instance-name</i>)></code> <code><logical-system (all <i>logical-system-name</i>)></code> <code><name></code> <code><peer (all <i>address</i>)></code> |
| Syntax (EX Series Switches and QFX Series) | <code>show rip statistics</code> <code><instance (all <i>instance-name</i>)></code> <code><name></code> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series. |
| Description | Display RIP statistics about messages sent and received on an interface, as well as information received from advertisements from other routing devices. |
| Options | none —Display RIP statistics for all routing instances. instance (all <i>instance-name</i>) —(Optional) Display RIP statistics for all instances or for only the specified routing instance. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. name —(Optional) Display detailed information about only the specified RIP neighbor. peer (all <i>address</i>) —(Optional) Display RIP statistics for a single peer or all peers. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">• clear rip statistics on page 693 |
| List of Sample Output | show rip statistics on page 700 |
| Output Fields | Table 154 on page 699 lists the output fields for the show rip statistics command. Output fields are listed in the approximate order in which they appear. |

Table 154: show rip statistics Output Fields

| Field Name | Field Description |
|---------------------------------|--|
| RIP info | <p>Information about RIP on the specified interface:</p> <ul style="list-style-type: none"> • port—UDP port number used for RIP. • update interval—Interval between routing table updates, in seconds. • holddown—Hold-down interval, in seconds. • timeout—Timeout interval, in seconds. • restart in progress—Graceful restart status. Displayed when RIP is or has been in the process of graceful restart. • restart time—Estimated time for the graceful restart to finish, in seconds. • restart will complete in—Remaining time for the graceful restart to finish, in seconds. • rts learned—Number of routes learned through RIP. • rts held down—Number of routes held down by RIP. • rqsts dropped—Number of received request packets that were dropped. • resps dropped—Number of received response packets that were dropped. |
| <i>logical-interface</i> | <p>Name of the logical interface and its statistics:</p> <ul style="list-style-type: none"> • routes learned—Number of routes learned on the logical interface. • routes advertised—Number of routes advertised by the logical interface. |
| Counter | <p>List of counter types:</p> <ul style="list-style-type: none"> • Updates Sent—Number of update messages sent. • Triggered Updates Sent—Number of triggered update messages sent. • Responses Sent—Number of response messages sent. • Bad Messages—Number of invalid messages received. • RIPv1 Updates Received—Number of RIPv1 update messages received. • RIPv1 Bad Route Entries—Number of RIPv1 invalid route entry messages received. • RIPv1 Updates Ignored—Number of RIPv1 update messages ignored. • RIPv2 Updates Received—Number of RIPv2 update messages received. • RIPv2 Bad Route Entries—Number of RIPv2 invalid route entry messages received. • RIPv2 Updates Ignored—Number of RIPv2 update messages ignored. • Authentication Failures—Number of received update messages that failed authentication. • RIP Requests Received—Number of RIP request messages received. • RIP Requests Ignored—Number of RIP request messages ignored. |
| Total | Total number of packets for the selected counter. |
| Last 5 min | Number of packets for the selected counter in the most recent 5-minute period. |
| Last minute | Number of packets for the selected counter in the most recent 1-minute period. |

Sample Output

show rip statistics

```

user@host> show rip statistics so-0/0/0.0
RIP info: port 520; update interval: 30s; holddown 180s; timeout 120s
restart in progress: restart time 60s; restart will complete in 55s
    rts learned  rts held down  rqsts dropped  resps dropped
              0              0              0              0
so-0/0/0.0: 0 routes learned; 501 routes advertised
Counter          Total    Last 5 min  Last minute
-----
Updates Sent      0          0          0
Triggered Updates Sent  0          0          0
Responses Sent    0          0          0
Bad Messages      0          0          0
RIPv1 Updates Received  0          0          0
RIPv1 Bad Route Entries  0          0          0
RIPv1 Updates Ignored  0          0          0
RIPv2 Updates Received  0          0          0
RIPv2 Bad Route Entries  0          0          0
RIPv2 Updates Ignored  0          0          0
Authentication Failures  0          0          0
RIP Requests Received  0          0          0
RIP Requests Ignored   0          0          0

```

RIPng Operational Mode Commands

Table 155 on page 701 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot the Routing Information Protocol next generation (RIPng). Commands are listed in alphabetical order.

Table 155: RIPng Operational Mode Commands

| Task | Command |
|-----------------------------|---|
| Clear general statistics. | <code>clear ripng general-statistics</code> |
| Clear statistics. | <code>clear ripng statistics</code> |
| Display general statistics. | <code>show ripng general-statistics</code> |
| Display RIPng neighbors. | <code>show ripng neighbor</code> |
| Display statistics. | <code>show ripng statistics</code> |



NOTE: For more RIPng-related commands, such as `show route protocol`, `show route instance`, and `show route table`, see Protocol-Independent Routing Operational Mode Commands.

For information about how to configure RIPng, see the *Junos Routing Protocols Configuration Guide*.

clear ripng general-statistics

| | |
|------------------------------------|---|
| Syntax | clear ripng general-statistics <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switches) | clear ripng general-statistics |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. |
| Description | Clear RIP next generation (RIPng) general statistics. |
| Options | none —Clear RIPng general statistics. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• show ripng general-statistics on page 704 |
| List of Sample Output | clear ripng general-statistics on page 702 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

```
clear ripng  
general-statistics      user@host> clear ripng general-statistics
```


clear ripng statistics

| | |
|----------------------------------|---|
| Syntax | clear ripng statistics <instance name> <logical-system (all logical-system-name)> |
| Syntax (EX Series Switch) | clear ripng statistics <instance name> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. |
| Description | Clear RIP next-generation (RIPng) statistics. |
| Options | <p>none—Reset RIPng counters for all neighbors for all routing instances.</p> <p>instance—(Optional) Reset RIPng counters for the specified instance.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>name—(Optional) Reset RIPng counters for the specified neighbor.</p> |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none"> • show ripng statistics on page 707 |
| List of Sample Output | clear ripng statistics on page 703 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear ripng statistics user@host> clear ripng statistics

show ripng general-statistics

| | |
|----------------------------------|---|
| Syntax | show ripng general-statistics <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switch) | show ripng general-statistics |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. |
| Description | Display general RIP next-generation (RIPng) statistics. |
| Options | none —Display general RIPng statistics. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> clear ripng general-statistics on page 702 |
| List of Sample Output | show ripng general-statistics on page 704 |
| Output Fields | Table 156 on page 704 lists the output fields for the show ripng general-statistics command. Output fields are listed in the approximate order in which they appear. |

Table 156: show ripng general-statistics Output Fields

| Field Name | Field Description |
|-------------|--|
| bad msgs | Number of invalid messages received. |
| no rcv intf | Number of packets received with no matching interface. |
| curr memory | Amount of memory currently used by RIPng. |
| max memory | Most memory used by RIPng. |

Sample Output

```

show ripng
general-statistics

user@host> show ripng general-statistics
RIPng I/O info:
  bad msgs      :      0
  no rcv intf   :      0
  curr memory   :      0
  max memory    :      0

```

show ripng neighbor

| | |
|----------------------------------|--|
| Syntax | show ripng neighbor <logical-system (all <i>logical-system-name</i>)> < <i>name</i> > |
| Syntax (EX Series Switch) | show ripng neighbor < <i>name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. |
| Description | Display information about RIP next-generation (RIPng) neighbors. |
| Options | <p>none—Display information about all RIPng neighbors.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><i>name</i>—(Optional) Display detailed information about a specific RIPng neighbor.</p> |
| Required Privilege Level | view |
| List of Sample Output | show ripng neighbor on page 706 |
| Output Fields | Table 157 on page 705 lists the output fields for the show ripng neighbor command. Output fields are listed in the approximate order in which they appear. |

Table 157: show ripng neighbor Output Fields

| Field Name | Field Description |
|----------------------------|--|
| Neighbor | Name of RIPng neighbor. |
| State | State of the connection: Up or Dn (Down). |
| Source Address | Source address. |
| Destination Address | Destination address. |
| Send | Send options: broadcast , multicast , none , version 1 , or yes . |
| Recv | Type of packets to accept: both , none , version 1 , or yes . |
| In Met | Metric added to incoming routes when advertising into RIPng routes that were learned from other protocols. |

Sample Output

`show ripng neighbor`

`user@host> show ripng neighbor`

| Neighbor | State | Source Address | Dest Address | Send | Recv | In Met |
|------------|-------|--------------------------|-----------------|-------|-------|-----------|
| ----- | ----- | ----- | ----- | ----- | ----- | ----- |
| fe-0/0/2.0 | Up | fe80::290:69ff:fe68:b002 | ff02::9 | yes | yes | 1 |

show ripng statistics

| | |
|----------------------------------|---|
| Syntax | show ripng statistics <logical-system (all <i>logical-system-name</i>)> < <i>name</i> > |
| Syntax (EX Series Switch) | show ripng statistics < <i>name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. |
| Description | Display RIP next generation (RIPng) statistics about messages sent and received on an interface, as well as information received from advertisements from other routing devices. |
| Options | <p>none—Display RIPng statistics for all neighbors.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><i>name</i>—(Optional) Display detailed information about a specific RIPng neighbor.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • clear ripng statistics on page 703 |
| List of Sample Output | show ripng statistics on page 708 |
| Output Fields | Table 158 on page 707 lists the output fields for the show ripng statistics command. Output fields are listed in the approximate order in which they appear. |

Table 158: show ripng statistics Output Fields

| Field Name | Field Description |
|------------|--|
| RIPng info | <p>Information about RIPng on the specified interface:</p> <ul style="list-style-type: none"> • port—UDP port number used for RIPng. • holddown—Hold-down interval, in seconds. • rts learned—Number of routes learned through RIPng. • rts held down—Number of routes held down by RIPng. • rqsts dropped—Number of received request packets that were dropped. • resps dropped—Number of received response packets that were dropped. • restart—Graceful restart status. Displayed when RIPng is or has been in the process of graceful restart. |

Table 158: show ripng statistics Output Fields (*continued*)

| Field Name | Field Description |
|--------------------------|---|
| <i>logical-interface</i> | Name of the logical interface and its statistics: <ul style="list-style-type: none"> • routes learned—Number of routes learned on the logical interface. • routes advertised—Number of routes advertised by the logical interface. • timeout—Timeout interval, in seconds. • update interval—Interval between routing table updates, in seconds. |
| Counter | List of counter types: <ul style="list-style-type: none"> • Updates Sent—Number of update messages sent. • Triggered Updates Sent—Number of triggered update messages sent. • Responses Sent—Number of response messages sent. • Bad Messages—Number of invalid messages received. • Updates Received—Number of RIPng update messages received. • Bad Route Entries—Number of RIPng invalid route entry messages received. • Updates Ignored—Number of RIPng update messages ignored. • RIPng Requests Received—Number of RIPng request messages received. • RIPng Requests Ignored—Number of RIPng request messages ignored. |
| Total | Total number of packets for the selected counter. |
| Last 5 min | Number of packets for the selected counter in the most recent 5-minute period. |
| Last minute | Number of packets for the selected counter in the most recent 1-minute period. |

Sample Output

```

user@host> show ripng statistics
RIPng info: port 521; holddown 120s;
      rts learned  rts held down  rqsts dropped  resps dropped
              0              0              0              0

so-0/1/3.0: 0 routes learned; 1 routes advertised; timeout 180s; update interval
20s
Counter                Total    Last 5 min  Last minute
-----
Updates Sent            934         16          4
Triggered Updates Sent    1          0          0
Responses Sent           0          0          0
Bad Messages             0          0          0
Updates Received          0          0          0
Bad Route Entries         0          0          0
Updates Ignored           0          0          0
RIPng Requests Received   0          0          0
RIPng Requests Ignored    0          0          0

```

PART 2

Policy Framework

- [Firewall Filter Operational Mode Commands on page 711](#)
- [Forwarding Operational Mode Commands on page 731](#)
- [Routing Policy Operational Mode Commands on page 765](#)

Firewall Filter Operational Mode Commands

Table 159 on page 711 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot firewall filters. Commands are listed in alphabetical order.

Table 159: Firewall Filter Operational Mode Commands


| Task | Command |
|---|--|
| Clear firewall filter counters. | <code>clear firewall</code> |
| Operational statistics for firewall filters. | <code>show firewall</code> |
| Version number of installed firewall filters. | <code>show firewall filter version</code> |
| Firewall filter log information. | <code>show firewall log</code> |
| Prefix-action statistics for firewall filters. | <code>show firewall prefix-action-stats</code> |
| Names of configured filter templates in use by dynamic subscribers and number of times each template is referenced. | <code>show firewall templates-in-use</code> |
| Counters for policers. | <code>show policer</code> |



NOTE: For information about how to configure firewall filters, see the *Junos Policy Framework Configuration Guide*.

For information about the related operational mode commands, `show interfaces filters` and `show interfaces policers`, see the *Junos Interfaces Command Reference*.

clear firewall

| | |
|------------------------------------|--|
| Syntax | clear firewall (all counter <i>counter-name</i> filter <i>filter-name</i> logical-system <i>logical-system-name</i>) |
| Syntax (EX Series Switches) | clear firewall (all counter <i>counter-name</i> filter <i>filter-name</i> policer counter (all counter-id <i>counter-index</i>)) |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. logical-system option introduced in Junos OS Release 9.3. |
| Description | <p>Clear statistics about configured firewall filters.</p> <p>When you clear the counters of a filter, this impacts not only the counters shown by the CLI, but also the ones tracked by SNMP2.</p> |
| | <div>  <p>NOTE: The clear firewall command cannot be used to clear the Routing Engine filter counters on a backup Routing Engine that is enabled for graceful Routing Engine switchover (GRES).</p> </div> |
| | <p>If you clear statistics for firewall filters that are applied to Trio-based DPCs and that also use the prefix-action action on matched packets, wait at least 5 seconds before you enter the show firewall prefix-action-stats command. A 5-second pause between issuing the clear firewall and show firewall prefix-action-stats commands avoids a possible timeout of the show firewall prefix-action-stats command.</p> |
| Options | <p>all—Clear the packet and byte counts for all filters. On EX Series switches, this option also clears the packet counts for all policer counters.</p> <p>counter <i>counter-name</i>—Clear the packet and byte counts for a filter counter that has been configured with the counter firewall filter action.</p> <p>filter <i>filter-name</i>—Clear the packet and byte counts for the specified firewall filter.</p> <p>logical-system <i>logical-system-name</i>—Clear the packet and byte counts for the specified logical system.</p> <p>policer counter (all counter-id <i>counter-index</i>)—(EX8200 switches only) Clear all policer counters using the policer counter all command, or clear a specific policer counter using the policer counter counter-id <i>counter-index</i> command. The value of <i>counter-index</i> can be 0, 1, or 2.</p> |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none"> • show firewall on page 714 |

List of Sample Output [clear firewall all on page 713](#)
 [clear firewall \(counter counter-name\) on page 713](#)
 [clear firewall \(filter filter-name\) on page 713](#)
 [clear firewall \(policer counter all\) \(EX8200 Switch\) on page 713](#)
 [clear firewall \(policer counter counter-id counter-index\) \(EX8200 Switch\) on page 713](#)

Sample Output

[clear firewall all](#) user@host> clear firewall all

[clear firewall \(counter counter-name\)](#) user@host> clear firewall counter port-filter-counter

[clear firewall \(filter filter-name\)](#) user@host> clear firewall filter ingress-port-filter

[clear firewall \(policer counter all\) \(EX8200 Switch\)](#) user@switch> clear firewall policer counter all

[clear firewall \(policer counter counter-id counter-index\) \(EX8200 Switch\)](#) user@switch> clear firewall policer counter counter-id 0

show firewall

| | |
|------------------------------------|--|
| Syntax | <pre>show firewall <counter <i>counter-name</i>> <filter <i>filter-name</i>> <log> <logical-system (all <i>logical-system-name</i>)> <terse></pre> |
| Syntax (EX Series Switches) | <pre>show firewall <counter <i>counter-name</i>> <detail> <filter <i>filter-name</i>> <log <(detail interface <i>interface-name</i>)>> <policer counters <(detail counter-id <i>counter-index</i> <detail>)>> <terse></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>logical-system option introduced in Junos OS Release 9.3.</p> <p>terse option introduced in Junos OS Release 9.4.</p> <p>policer counters option introduced in Junos OS Release 12.2 for EX Series switches.</p> <p>detail option introduced in Junos OS Release 12.3.</p> |
| Description | Display statistics about configured firewall filters. |
| Options | <p>none—(Optional) Display statistics about all configured firewall filters and counters. For EX Series switches, this command also displays statistics about all configured policers.</p> <p>counter <i>counter-name</i>—(Optional) Name of a filter counter.</p> <p>detail—(EX Series switches only) (Optional) Display firewall filter statistics with enhanced policer.</p> <p>filter <i>filter-name</i>—(Optional) Name of a configured filter.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>log—(Optional) Display log entries for firewall filters.</p> <p>log <(detail interface <i>interface-name</i>)>—(EX Series switches only) (Optional) Display detailed log entries of firewall activity or log information about a specific interface.</p> <p>policer counters <(detail counter-id <i>counter-index</i> <detail>)>—(EX8200 switches only) (Optional) Display policer counter statistics in brief or in detail.</p> <p>terse—(Optional) Display firewall filter names only.</p> |
| Required Privilege Level | view |

- Related Documentation**
- [clear firewall on page 712](#)
 - [show firewall log on page 722](#)
 - Verifying That Firewall Filters Are Operational
 - Verifying That Policers Are Operational

- List of Sample Output**
- [show firewall filter \(MX Series\) on page 717](#)
 - [show firewall filter \(non MX Series Router\) on page 717](#)
 - [show firewall filter \(Hierarchical Policier, MX Series with MPC\) on page 717](#)
 - [show firewall filter \(Dynamic Input Filter\) on page 717](#)
 - [show firewall \(Logical Systems\) on page 717](#)
 - [show firewall \(counter counter-name\) on page 718](#)
 - [show firewall log on page 718](#)
 - [show firewall policer counters \(EX8200 Switch\) on page 718](#)
 - [show firewall policer counters \(detail\) \(EX8200 Switch\) on page 719](#)
 - [show firewall policer counters \(counter-id counter-index\) \(EX8200 Switch\) on page 719](#)
 - [show firewall policer counters \(counter-id counter-index detail\) \(EX8200 Switch\) on page 719](#)
 - [show firewall detail on page 720](#)

- Output Fields** [Table 160 on page 715](#) lists the output fields for the **show firewall** command. Output fields are listed in the approximate order in which they appear.

Table 160: show firewall Output Fields

| Field Name | Field Description |
|-----------------|--|
| Filter | <p>Name of a filter that has been configured with the filter statement at the [edit firewall] hierarchy level.</p> <p>Except on EX Series switches:</p> <ul style="list-style-type: none"> • When an interface-specific filter is displayed, the name of the filter is followed by the full interface name and by either -i for an input filter or -o for an output filter. • When dynamic filters are displayed, the name of the filter is followed by the full interface name and by either -in for an input filter or -out for an output filter. When a logical system-specific filter is displayed, the name of the filter is prefixed with two underscore (__) characters and the name of the logical system (for example, __ls1/filter1). |
| Counters | <p>Display filter counter information:</p> <ul style="list-style-type: none"> • Name—Name of a filter counter that has been configured with the counter firewall filter action. • Bytes—Number of bytes that match the filter term under which the counter action is specified. • Packets—Number of packets that matched the filter term under which the counter action is specified. |

Table 160: show firewall Output Fields (*continued*)

| Field Name | Field Description |
|------------------------------|--|
| Policers | <p>Display policer information:</p> <ul style="list-style-type: none"> • Name—Name of policer. • Bytes—(For two-color policers on MX Series routers, and for hierarchical policers on interfaces hosted on MICs and MPCs in MX Series routers) Number of bytes that match the filter term under which the policer action is specified. This is only the number out-of-specification (out-of-spec) byte counts, not all the bytes in all packets policed by the policer. For other platforms, this field is blank. • Packets—Number of packets that matched the filter term under which the policer action is specified. This is only the number of out-of-specification (out-of-spec) packet counts, not all packets policed by the policer. |
| Policer Counter Index | (EX8200 switch only) Global management counter ID. The counter ID value (<i>counter-index</i>) can be 0, 1, or 2. |
| Green | (EX8200 switch only) Number of packets within the limits. The number of packets is smaller than the committed information rate (CIR). |
| Yellow | (EX8200 switch only) Number of packets partially within the limits. The number of packets is greater than the CIR, but the burst size is within the excess burst size (EBS) limit. |
| Discard | (EX8200 switch only) Number of discarded packets. |
| Bytes | (EX8200 switch only) Number of green, yellow, red, or discarded packets in bytes. |
| Packets | (EX8200 switch only) Number of green, yellow, red, or discarded packets. |
| Filter name | (EX8200 switch only) Name of the filter with a term associated to a policer. |
| Term name | (EX8200 switch only) Name of the term associated with a policer. |
| Policer name | (EX8200 switch only) Name of the policer that is associated with a global management counter. |

Sample Output

show firewall filter (MX Series)

```
user@host> show firewall filter test
Filter: test
Counters:
Name          Bytes      Packets
Counter-1      0           0
Counter-2      0           0
Policers:
Name          Bytes      Packets
Policer-1     2770       70
```

show firewall filter (non MX Series Router)

```
user@host> show firewall filter test
Filter: test
Counters:
Name          Bytes      Packets
Counter-1      0           0
Counter-2      0           0
Policers:
Name          Bytes      Packets
Policer-1     2770       70
```

show firewall filter (Hierarchical Policers, MX Series with MPC)

```
user@host> show firewall filter
FL_V4_PHY-HP-EF-AWARE-Gold=400k-MCAST=200k-Total=1M-ds-10/0/0:2:1-i

Filter: FL_V4_PHY-HP-EF-AWARE-Gold=400k-MCAST=200k-Total=1M-ds-10/0/0:2:1-i
Counters:
Name          Bytes      Packets
AF1x_counter-ds-10/0/0:2:1-i      0           0
AF2x_counter-ds-10/0/0:2:1-i    25529445976  24500428
AF3x_counter-ds-10/0/0:2:1-i    2182022      39482
AF4x_counter-ds-10/0/0:2:1-i      0           0
BE_counter-ds-10/0/0:2:1-i        0           0
EF_counter-ds-10/0/0:2:1-i    14817044120  12265765
STD_counter-ds-10/0/0:2:1-i      0           0
Policers:
Name          Bytes      Packets
POL_CE-PE_M=200k-filter-ds-10/0/0:2:1-i  5948099658  5708349
POL_CE-PE_G=400K_R=1M-filter-ds-10/0/0:2:1-i  ??????????  3572794
???????????? ?????????? ??????????
```

show firewall filter (Dynamic Input Filter)

```
user@host> show firewall filter dfwd-ge-5/0/0.1-in
Filter: dfwd-ge-5/0/0.1-in
Counters:
Name          Bytes      Packets
c1-ge-5/0/0.1-in      0           0
```

show firewall (Logical Systems)

```
user@host> show firewall

Filter: __lr1/test
Counters:
Name          Bytes      Packets
icmp          420         5
Filter: __default_bpdu_filter__
```

```

Filter: __lr1/inet_filter1
Counters:
Name                               Bytes      Packets
inet_tcp_count                     0          0
inet_udp_count                     0          0
Filter: __lr1/inet_filter2
Counters:
Name                               Bytes      Packets
inet_icmp_count                    0          0
inet_pim_count                     0          0
Filter: __lr2/inet_filter1
Counters:
Name                               Bytes      Packets
inet_tcp_count                     0          0
inet_udp_count                     0          0

```

show firewall (counter counter-name)

```

user@host> show firewall counter icmp-counter
Filter: ingress-port-voip-class-filter
Counters:
Name                               Bytes      Packets
icmp-counter                       0          0

```

show firewall log

```

user@host> show firewall log
Log :

Time      Filter  Action Interface  Protocol  Src Addr
Dest Addr
08:00:53  pfe      R      ge-1/0/1.0    ICMP      192.168.3.5
192.168.3.4
08:00:52  pfe      R      ge-1/0/1.0    ICMP      192.168.3.5
192.168.3.4
08:00:51  pfe      R      ge-1/0/1.0    ICMP      192.168.3.5
192.168.3.4
08:00:50  pfe      R      ge-1/0/1.0    ICMP      192.168.3.5
192.168.3.4
08:00:49  pfe      R      ge-1/0/1.0    ICMP      192.168.3.5
192.168.3.4
08:00:48  pfe      R      ge-1/0/1.0    ICMP      192.168.3.5
192.168.3.4
08:00:47  pfe      R      ge-1/0/1.0    ICMP      192.168.3.5
192.168.3.4

```

show firewall policer counters (EX8200 Switch)

```

user@switch> show firewall policer counters
Policer Counter Index 0:
Bytes      Packets
Green:      73      15914
Yellow:     9      1962
Discard:    119     25942

Policer Counter Index 1:
Bytes      Packets
Green:      0      0
Yellow:     0      0
Discard:    0      0

Policer Counter Index 2:
Bytes      Packets

```



```

Green:                0                0
Yellow:               0                0
Discard:              0                0

```

**show firewall policer
counters (detail)
(EX8200 Switch)**

```
user@switch> show firewall policer counters detail
```

```
Policer Counter Index 0:
```

```

          Bytes          Packets
Green:         73         15914
Yellow:         9         1962
Discard:       119        25942

```

```

Filter name      Term name      Policer name
myfilter        polcr-term-1    myfilter-polcr-1
inet-filter-ae   ae-snmp          policer-1
inet-filter-ae   ae-ssh           policer-2

```

```
Policer Counter Index 1:
```

```

          Bytes          Packets
Green:         0          0
Yellow:         0          0
Discard:         0          0

```

```

Filter name      Term name      Policer name

```

```
Policer Counter Index 2:
```

```

          Bytes          Packets
Green:         0          0
Yellow:         0          0
Discard:         0          0

```

```

Filter name      Term name      Policer name

```

**show firewall policer
counters (counter-id
counter-index)
(EX8200 Switch)**

```
user@switch> show firewall policer counters counter-id 0
```

```
Policer Counter Index 0:
```

```

          Bytes          Packets
Green:         73         15914
Yellow:         9         1962
Discard:       119        25942

```

**show firewall policer
counters (counter-id**

```
user@switch> show firewall policer counters counter-id 0 detail
```

```
Policer Counter Index 0:
```

```

          Bytes          Packets

```

counter-index detail
(EX8200 Switch)

| | | |
|----------|-----|-------|
| Green: | 73 | 15914 |
| Yellow: | 9 | 1962 |
| Discard: | 119 | 25942 |

| Filter name | Term name | Policer name |
|----------------|--------------|------------------|
| myfilter | polcr-term-1 | myfilter-polcr-1 |
| inet-filter-ae | ae-snmp | policer-1 |
| inet-filter-ae | ae-ssh | policer-2 |

show firewall detail

```
user@host> show firewall detail
Filter: __default_bpdu_filter__

Filter: foo
Counters:
Name                               Bytes          Packets
c1                                  17652140       160474
Policers:
Name                               Bytes          Packets
P1-t1
  OOS                               0              18286
  Offered                           0 18446744073709376546
  Transmitted                       0 18446744073709358260
```

show firewall filter version

| | |
|---------------------------------|--|
| Syntax | show firewall filter version <filter-name> |
| Release Information | Command introduced in Junos OS Release 10.2R2. |
| Description | Display the version number of the installed firewall filter in the Routing Engine. |
| Options | <p>none—(Optional) Display the version number of all installed firewall filters.</p> <p>filter-name—(Optional) Name of a configured filter. If you specify the name of a filter, only the version number of that filter is displayed.</p> |
| Additional Information | The initial version number is 1. This number increments by one when you modify the firewall filter settings or an associated prefix action. The maximum version number is 4,294,967,295. When the version number reaches 4,294,967,295, this number is reset to 1. |
| Required Privilege Level | view |
| List of Sample Output | show firewall filter version on page 721 |
| Output Fields | Table 161 on page 721 lists the output fields for the show firewall filter version command. Output fields are listed in the approximate order in which they appear. |

Table 161: show firewall filter version Output Fields

| Field Name | Field Description |
|------------|---|
| Filter | Name of a filter that has been configured with the filter statement at the [edit firewall] hierarchy level. |
| Version | Display the version number of the firewall filter. |

Sample Output

```

show firewall filter version
user@host> show firewall filter version
Filter version information :
Filter                                     Version
test                                     10

```

show firewall log

| | |
|------------------------------------|---|
| Syntax | show firewall log <detail> <interface <i>interface-name</i> > <logical-system (<i>logical-system-name</i> all)> |
| Syntax (EX Series Switches) | show firewall log <detail> <interface <i>interface-name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. logical-system option introduced in Junos OS Release 9.3. |
| Description | Display log information about firewall filters. |
| Options | <p>none—Display log information about firewall filters.</p> <p>detail—(Optional) Display detailed information.</p> <p>interface <i>interface-name</i>—(Optional) Display log information about a specific interface.</p> <p>logical-system (<i>logical-system-name</i> all)—(Optional) Perform this operation on all logical systems or on a particular system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show firewall log on page 724 show firewall log detail on page 724 |
| Output Fields | Table 162 on page 722 lists the output fields for the show firewall log command. Output fields are listed in the approximate order in which they appear. |

Table 162: show firewall log Output Fields

| Field Name | Field Description |
|--------------------|---|
| Time of Log | Time that the event occurred. |
| Filter | <p>Name of a filter that has been configured with the filter statement at the [edit firewall] hierarchy level.</p> <ul style="list-style-type: none"> A hyphen (-) indicates that the packet was handled by the Packet Forwarding Engine. A space (no hyphen) indicates the packet was handled by the Routing Engine. The notation pfe indicates packets logged by the Packet Forwarding Engine hardware filters. |

Table 162: show firewall log Output Fields (*continued*)

| Field Name | Field Description |
|---------------------|---|
| Filter Action | Filter action: <ul style="list-style-type: none">• A—Accept• D—Discard• R—Reject |
| Name of Interface | Ingress interface for the packet. |
| Name of protocol | Packet's protocol name: egp , gre , icmp , ipip , ospf , pim , rsvp , tcp , or udp . |
| Packet length | Length of the packet. |
| Source address | Packet's source address. |
| Destination address | Packet's destination address and port. |

Sample Output

show firewall log

```
user@host>show firewall log
```

| Time | Filter | Action | Interface | Protocol | Src Addr | Dest Addr |
|----------|--------|--------|-----------|----------|-------------|-------------|
| 13:10:12 | pfe | D | rlsq0.902 | ICMP | 180.1.177.2 | 180.1.177.1 |
| 13:10:11 | pfe | D | rlsq0.902 | ICMP | 180.1.177.2 | 180.1.177.1 |

show firewall log detail

```
user@host> show firewall log detail
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0Name of protocol: TCP, Packet Length: 50824, Source address:
172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 1020, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
....
```

show firewall prefix-action-stats

| | |
|---------------------------------|--|
| Syntax | show firewall prefix-action-stats filter <i>filter-name</i> prefix-action <i>prefix-action-name</i> <from <i>number</i> to <i>number</i> > <logical-system (<i>logical-system-name</i> all)> |
| Release Information | Command introduced before Junos OS Release 7.4. logical-system option introduced in Junos OS Release 9.3. |
| Description | Display prefix action statistics about configured firewall filters. If you clear statistics for firewall filters that are applied to Trio-based DPCs and that also use the prefix-action action on matched packets, wait at least 5 seconds before you enter the show firewall prefix-action-stats command. A 5-second pause between issuing the clear firewall and show firewall prefix-action-stats commands avoids a possible timeout of the show firewall prefix-action-stats command. |
| Options | filter <i>filter-name</i> —Name of a filter. prefix-action <i>prefix-action-name</i> —Name of a prefix action. from <i>number</i> to <i>number</i> —(Optional) Starting and ending counter or policer. logical-system (<i>logical-system-name</i> all) —(Optional) Perform this operation on all logical systems or on a particular system. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • clear firewall on page 712 |
| List of Sample Output | show firewall prefix-action-stats on page 725 |
| Output Fields | Table 163 on page 725 lists the output fields for the show firewall prefix-action-stats command. Output fields are listed in the approximate order in which they appear. |

Table 163: show firewall prefix-action-stats Output Fields

| Field Name | Field Description |
|---------------|---|
| Filter | Filter name. Filters configured for logical systems include the name of the filter prefixed with the two underscore characters (__) and the name of the logical system (for example, __ls1/filter1). |

Sample Output

show firewall
prefix-action-stats

```
user@host> show firewall prefix-action-stats filter test prefix-action act1
Filter: __ls2/test
```

show firewall templates-in-use

| | |
|---------------------------------|--|
| Syntax | show firewall templates-in-use |
| Release Information | Command introduced in Junos OS Release 12.3. |
| Description | Display the names of configured filter templates that are currently in use by dynamic subscribers and the number of times each template is referenced. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">• clear firewall on page 712• show firewall log on page 722• Verifying That Firewall Filters Are Operational |
| List of Sample Output | show firewall templates-in-use on page 727 |
| Output Fields | Table 164 on page 726 lists the output fields for the show firewall templates-in-use command. Output fields are listed in the approximate order in which they appear. |

Table 164: show firewall templates-in-use Output Fields

| Field Name | Field Description |
|------------------------|---|
| Filter Template | Name of a filter that has been configured using the filter statement at either the [edit firewall] or [edit dynamic-profiles <i>profile-name</i> firewall] hierarchy and is being used as a template for dynamic subscriber filtering. |
| Reference Count | Number of times the filter has been referenced by subscribers accessing the network. |

Sample Output

show firewall
templates-in-use

```
user@host> show firewall templates-in-use
```

| Filter Template | Dynamic | Subscribers | Reference Counts |
|-----------------|---------|-------------|------------------|
| ----- | | | ----- |
| egressFilter | | | 10 |
| ingressFilter | | | 10 |
| dfilter | | | 5 |
| dfilter-pol | | | 5 |

show policer

| | |
|---------------------------------|--|
| Syntax | show policer <detail> <policer-name> |
| Release Information | Command introduced before Junos OS Release 7.4. Option detail introduced in Junos OS Release 12.3. |
| Description | Display the number of policed packets for a given policer or an aggregate policer. An aggregate policer is an aggregate of different policers on the same logical interface. |
| Options | none —Display the number of policed packets for all configured policers. detail —(Optional) Display enhanced statistics for policers. policer-name —(Optional) Display the number of policed packets for the specified policer. |
| Required Privilege Level | view |
| List of Sample Output | show policer (MX Series) on page 729 show policer (non MX Series Router) on page 729 show policer (Aggregate Policar, non MX Series Router) on page 729 show policer detail on page 729 |
| Output Fields | Table 165 on page 728 lists the output fields for the show policer command. Output fields are listed in the approximate order in which they appear. |

Table 165: show policer Output Fields

| Field Name | Field Description |
|----------------|--|
| Name | Name of the policer. |
| Bytes | (For two-color policers on MX Series routers, and for hierarchical policers on interfaces hosted on MICs and MPCs in MX Series routers) Total number of bytes policed by the specified policer. For other platforms, this field is blank. |
| Packets | Total number of packets policed by the specified policer. |

Sample Output

show policer (MX Series)

```
user@host> show policer
Policers:
Name                                     Bytes      Packets
__default_arp_policer__                 314520      5242
pol-2M-ge-1/2/0.1-inet-i                10372300    103723
pol-2M-ge-1/2/0.1-inet6-i               7727800     77278
pol-2M-ge-1/2/0.1-mp1s-i                 7070336     67984
pol-2M-ge-1/2/0.1001-vpls-i             65153700    651537
pol-2M-ge-1/2/0.2001-vpls-i             65180900    651809
pol-2M-ge-1/2/0.3001-ccc-i              62202144    647939
```

show policer (non MX Series Router)

```
user@host> show policer
Policers:
Name                                     Bytes      Packets
__default_arp_policer__                 5242
pol-2M-ge-1/2/0.1-inet-i                103723
pol-2M-ge-1/2/0.1-inet6-i               77278
pol-2M-ge-1/2/0.1-mp1s-i                 67984
pol-2M-ge-1/2/0.1001-vpls-i             651537
pol-2M-ge-1/2/0.2001-vpls-i             651809
pol-2M-ge-1/2/0.3001-ccc-i              647939
```

show policer (Aggregate Policar, non MX Series Router)

```
user@host> show policer
Policers:
Name                                     Bytes      Packets
__default_arp_policer__                  0
P1-ae0.0-log_int-o                      0
P2-ge-7/0/2.0-inet-o                     0
P2-ge-7/0/2.0-inet6-o                    0
__policer_tmpl__-term                    0
__policer_tmpl__-fc0                     0
__policer_tmpl__-fc0                     0
__policer_tmpl__-fc1                     0
__policer_tmpl__-fc0                     0
__policer_tmpl__-fc1                     0
__policer_tmpl__-fc2                     0
__policer_tmpl__-fc0                     0
__policer_tmpl__-fc1                     0
__policer_tmpl__-fc2                     0
__policer_tmpl__-fc3                     0
```

show policer detail

```
user@host> show policer detail
Policers:
Name                                     Bytes      Packets
__default_arp_policer__
  OOS                                     0          0
  Offered                                0         496
  Transmitted                             0         496
P1-xe-1/0/0.0-inet-i
  OOS                                     0        11329
  Offered                                0       111188
  Transmitted                             0       99859
```


Forwarding Operational Mode Commands

Table 166 on page 731 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot forwarding options. Commands are listed in alphabetical order.

Table 166: Forwarding Operational Mode Commands

| Task | Command |
|---|--|
| Clear the binding state of a Dynamic Host Configuration Protocol (DHCP) client from the client table. | <code>clear dhcp relay binding</code> |
| Clear all DHCP relay statistics. | <code>clear dhcp relay statistics</code> |
| Clear the binding state of a DHCPv6 client from the client table. | <code>clear dhcpv6 relay binding</code> |
| Clear all DHCPv6 relay statistics. | <code>clear dhcpv6 relay statistics</code> |
| Clear statistic counters in the User Datagram Protocol (UDP) forwarding process. | <code>clear helper statistics</code> |
| Display the address bindings in the DHCP client table. | <code>show dhcp relay binding</code> |
| Display DHCP relay statistics. | <code>show dhcp relay statistics</code> |
| Display the address bindings in the DHCPv6 client table. | <code>show dhcpv6 relay binding</code> |
| Display DHCPv6 relay statistics. | <code>show dhcpv6 relay statistics</code> |
| Display statistics collected by the UDP forwarding process. | <code>show helper statistics</code> |

clear dhcp relay binding

| | |
|---------------------------------|--|
| Syntax | clear dhcp relay binding <address> <all> <interface <i>interface-name</i>> <interfaces-vlan> <interfaces-wildcard> <logical-system <i>logical-system-name</i>> <routing-instance <i>routing-instance-name</i>> |
| Release Information | Command introduced in Junos OS Release 8.3. Options all and interface added in Junos OS Release 8.4. Options interfaces-vlan and interfaces-wildcard added in Junos OS Release 12.1. Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Switches. |
| Description | Clear the binding state of a Dynamic Host Configuration Protocol (DHCP) client from the client table. |
| Options | <p>address—(Optional) Clear the binding state for the DHCP client, using one of the following entries:</p> <ul style="list-style-type: none">• ip-address—The specified IP address.• mac-address—The specified MAC address.• session-id—The specified session ID. <p>all—(Optional) Clear the binding state for all DHCP clients.</p> <p>interface <i>interface-name</i>—(Optional) Clear the binding state for DHCP clients on the specified interface.</p> <p>interfaces-vlan—(Optional) Clear the binding state on the interface VLAN ID and S-VLAN ID.</p> <p>interfaces-wildcard—(Optional) The set of interfaces on which to clear bindings. This option supports the use of the wildcard character (*).</p> <p>logical-system <i>logical-system-name</i>—(Optional) Clear the binding state for DHCP clients on the specified logical system.</p> <p>routing-instance <i>routing-instance-name</i>—(Optional) Clear the binding state for DHCP clients on the specified routing instance.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">• Clearing DHCP Bindings for Subscriber Access• show dhcp relay binding on page 745 |

List of Sample Output [clear dhcp relay binding on page 733](#)
[clear dhcp relay binding all on page 733](#)
[clear dhcp relay binding interface on page 733](#)
[clear dhcp relay binding <interfaces-vlan> on page 733](#)
[clear dhcp relay binding <interfaces-wildcard> on page 733](#)

Output Fields See [show dhcp relay binding](#) for an explanation of output fields.

Sample Output

clear dhcp relay binding

The following sample output displays the address bindings in the DHCP client table before and after the **clear dhcp relay binding** command is issued.

```
user@host> show dhcp relay binding
IP address      Hardware address  Type    Lease expires at
100.20.32.1     90:00:00:01:00:01 active    2007-02-08 16:41:17 EST
192.168.14.8    90:00:01:01:02:01 active    2007-02-10 10:01:06 EST
```

```
user@host> clear dhcp relay binding 100.20.32.1
```

```
user@host> show dhcp relay binding
IP address      Hardware address  Type    Lease expires at
192.168.14.8    90:00:01:01:02:01 active    2007-02-10 10:01:06 EST
```

clear dhcp relay binding all

The following command clears all DHCP relay agent bindings:

```
user@host> clear dhcp relay binding all
```

clear dhcp relay binding interface

The following command clears DHCP relay agent bindings on a specific interface:

```
user@host> clear dhcp relay binding interface fe-0/0/3
```

clear dhcp relay binding <interfaces-vlan>

The following command uses the *interfaces-vlan* option to clear all DHCP relay agent bindings on top of the underlying interface **ae0**, which clears DHCP bindings on all demux VLANs on top of **ae0**:

```
user@host> clear dhcp relay binding interface ae0
```

clear dhcp relay binding <interfaces-wildcard>

The following command uses the *interfaces-wildcard* option to clear all DHCP relay agent bindings over a specific interface:

```
user@host> clear dhcp relay binding ge-1/0/0.*
```

clear dhcp relay statistics

| | |
|--------------------------|--|
| Syntax | <code>clear dhcp relay statistics</code> <code><logical-system <i>logical-system-name</i>></code> <code><routing-instance <i>routing-instance-name</i>></code> |
| Syntax | Syntax for EX Series switches: <code>show dhcp relay statistics</code> <code><routing-instance <i>routing-instance-name</i>></code> |
| Release Information | Command introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 12.1 for EX Series switches. Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Switches. |
| Description | Clear all Dynamic Host Configuration Protocol (DHCP) relay statistics. |
| Options | logical-system <i>logical-system-name</i> —(On routers only) (Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are cleared for the default logical system. routing-instance <i>routing-instance-name</i> —(Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are cleared for the default routing instance. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">• show dhcp relay statistics on page 750 |
| List of Sample Output | clear dhcp relay statistics on page 736 |
| Output Fields | Table 167 on page 735 lists the output fields for the <code>clear dhcp relay statistics</code> command. |

Table 167: clear dhcp relay statistics Output Fields

| Field Name | Field Description |
|--------------------------|---|
| Packets dropped | <p>Number of packets discarded by the extended DHCP relay agent application due to errors. Only nonzero statistics appear in the Packets dropped output. When all of the Packets dropped statistics are 0 (zero), only the Total field appears.</p> <ul style="list-style-type: none"> • Total—Total number of packets discarded by the extended DHCP relay agent application. • Bad hardware address—Number of packets discarded because an invalid hardware address was specified. • Bad opcode—Number of packets discarded because an invalid operation code was specified. • Bad options—Number of packets discarded because invalid options were specified. • Invalid server address—Number of packets discarded because an invalid server address was specified. • No available addresses—Number of packets discarded because there were no addresses available for assignment. • No interface match—Number of packets discarded because they did not belong to a configured interface. • No routing instance match—Number of packets discarded because they did not belong to a configured routing instance. • No valid local address—Number of packets discarded because there was no valid local address. • Packet too short—Number of packets discarded because they were too short. • Read error—Number of packets discarded because of a system read error. • Send error—Number of packets that the extended DHCP relay application could not send. • Option 60—Number of packets discarded containing DHCP option 60 vendor-specific information. • Option 82—Number of packets discarded because DHCP option 82 information could not be added. |
| Messages received | <p>Number of DHCP messages received.</p> <ul style="list-style-type: none"> • BOOTREQUEST—Number of BOOTP protocol data units (PDUs) received • DHCPDECLINE—Number of DHCP PDUs of type DECLINE received • DHCPDISCOVER—Number of DHCP PDUs of type DISCOVER received • DHCPINFORM—Number of DHCP PDUs of type INFORM received • DHCPRELEASE—Number of DHCP PDUs of type RELEASE received • DHCPREQUEST—Number of DHCP PDUs of type REQUEST received |
| Messages sent | <p>Number of DHCP messages sent.</p> <ul style="list-style-type: none"> • BOOTREPLY—Number of BOOTP PDUs transmitted • DHCPOFFER—Number of DHCP OFFER PDUs transmitted • DHCPACK—Number of DHCP ACK PDUs transmitted • DHCPNACK—Number of DHCP NACK PDUs transmitted |

Sample Output

clear dhcp relay statistics

The following sample output displays the DHCP relay statistics before and after the **clear dhcp relay statistics** command is issued.

```
user@host> show dhcp relay statistics
Packets dropped:
  Total                      0

Messages received:
  BOOTREQUEST                116
  DHCPDECLINE                 0
  DHCPDISCOVER                11
  DHCPINFORM                  0
  DHCPRELEASE                 0
  DHCPREQUEST                105

Messages sent:
  BOOTREPLY                   44
  DHCPOFFER                   11
  DHCPACK                     11
  DHCPNAK                     11
```

```
user@host> clear dhcp relay statistics
```

```
user@host> show dhcp relay statistics
Packets dropped:
  Total                      0

Messages received:
  BOOTREQUEST                0
  DHCPDECLINE                 0
  DHCPDISCOVER                0
  DHCPINFORM                  0
  DHCPRELEASE                 0
  DHCPREQUEST                0

Messages sent:
  BOOTREPLY                   0
  DHCPOFFER                   0
  DHCPACK                     0
  DHCPNAK                     0
```

clear dhcpv6 relay binding

| | |
|---------------------------------|---|
| Syntax | <pre>clear dhcpv6 relay binding <address> <all> <interface interface-name> <interfaces-vlan> <interfaces-wildcard> <logical-system logical-system-name> <routing-instance routing-instance-name></pre> |
| Release Information | <p>Command introduced in Junos OS Release 11.4.</p> <p>Options <i>interfaces-vlan</i> and <i>interfaces-wildcard</i> added in Junos OS Release 12.1.</p> <p>Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Switches.</p> |
| Description | Clear the binding state of Dynamic Host Configuration Protocol for IPv6 (DHCPv6) clients from the client table. |
| Options | <p>address—(Optional) Clear the binding state for the DHCPv6 client, using one of the following entries:</p> <ul style="list-style-type: none"> • <i>CID</i>—The specified Client ID (CID). • <i>ipv6-prefix</i>—The specified IPv6 prefix. • <i>session-id</i>—The specified session ID. <p>all—(Optional) Clear the binding state for all DHCPv6 clients.</p> <p>interfaces-vlan—(Optional) Clear the binding state on the interface VLAN ID and S-VLAN ID.</p> <p>interfaces-wildcard—(Optional) The set of interfaces on which to clear bindings. This option supports the use of the wildcard character (*).</p> <p>interface interface-name—(Optional) Clear the binding state for DHCPv6 clients on the specified interface.</p> <p>logical-system logical-system-name—(Optional) Clear the binding state for DHCPv6 clients on the specified logical system.</p> <p>routing-instance routing-instance-name—(Optional) Clear the binding state for DHCPv6 clients on the specified routing instance.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • Clearing DHCP Bindings for Subscriber Access • show dhcpv6 relay binding on page 753 |

- List of Sample Output** [clear dhcpv6 relay binding on page 739](#)
 [clear dhcpv6 relay binding <prefix> on page 739](#)
 [clear dhcpv6 relay binding all on page 739](#)
 [clear dhcpv6 relay binding interface on page 739](#)
 [clear dhcpv6 relay binding <interfaces-vlan> on page 739](#)
 [clear dhcpv6 relay binding <interfaces-wildcard> on page 739](#)
- Output Fields** See [show dhcpv6 relay binding](#) for an explanation of output fields.

Sample Output

clear dhcpv6 relay binding

The following sample output displays the DHCPv6 bindings before and after the **clear dhcpv6 relay binding** command is issued.

```
user@host> show dhcpv6 relay binding
```

| Prefix | Session Id | Expires | State | Interface | Client DUID |
|---|------------|---------|-------|------------|-------------|
| 2001:bd8:3c4d:15::/64 | 1 | 83720 | BOUND | ge-1/0/0.0 | |
| LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01 | | | | | |
| 2001:bd8:3c4d:16::/64 | 2 | 83720 | BOUND | ge-1/0/0.0 | |
| LL_TIME0x1-0x4bfa26af-00:10:94:00:00:02 | | | | | |
| 2001:bd8:3c4d:17::/64 | 3 | 83720 | BOUND | ge-1/0/0.0 | |
| LL_TIME0x1-0x4bfa26af-00:10:94:00:00:03 | | | | | |
| 2001:bd8:3c4d:18::/64 | 4 | 83720 | BOUND | ge-1/0/0.0 | |
| LL_TIME0x1-0x4bfa26af-00:10:94:00:00:04 | | | | | |
| 2001:bd8:3c4d:19::/64 | 5 | 83720 | BOUND | ge-1/0/0.0 | |
| LL_TIME0x1-0x4bfa26af-00:10:94:00:00:05 | | | | | |
| 2001:bd8:3c4d:20::/64 | 6 | 83720 | BOUND | ge-1/0/0.0 | |
| LL_TIME0x1-0x4bfa26af-00:10:94:00:00:06 | | | | | |

clear dhcpv6 relay binding <prefix>

```
user@host> clear dhcpv6 relay binding 2001:bd8:3c4d:15::/64
```

```
user@host> show dhcpv6 relay binding
```

| Prefix | Session Id | Expires | State | Interface | Client DUID |
|---|------------|---------|-------|------------|-------------|
| 2001:bd8:3c4d:16::/64 | 2 | 83720 | BOUND | ge-1/0/0.0 | |
| LL_TIME0x1-0x4bfa26af-00:10:94:00:00:02 | | | | | |
| 2001:bd8:3c4d:17::/64 | 3 | 83720 | BOUND | ge-1/0/0.0 | |
| LL_TIME0x1-0x4bfa26af-00:10:94:00:00:03 | | | | | |
| 2001:bd8:3c4d:18::/64 | 4 | 83720 | BOUND | ge-1/0/0.0 | |
| LL_TIME0x1-0x4bfa26af-00:10:94:00:00:04 | | | | | |
| 2001:bd8:3c4d:19::/64 | 5 | 83720 | BOUND | ge-1/0/0.0 | |
| LL_TIME0x1-0x4bfa26af-00:10:94:00:00:05 | | | | | |
| 2001:bd8:3c4d:20::/64 | 6 | 83720 | BOUND | ge-1/0/0.0 | |
| LL_TIME0x1-0x4bfa26af-00:10:94:00:00:06 | | | | | |

clear dhcpv6 relay binding all

The following command clears all DHCP relay agent bindings:

```
user@host> clear dhcpv6 relay binding all
```

clear dhcpv6 relay binding interface

The following command clears DHCPv6 relay agent bindings on a specific interface:

```
user@host> clear dhcpv6 relay binding interface fe-0/0/2
```

clear dhcpv6 relay binding <interfaces-vlan>

The following command uses the *interfaces-vlan* option to clear all DHCPv6 relay agent bindings on top of the underlying interface **ae0**, which clears DHCPv6 bindings on all demux VLANs on top of **ae0**:

```
user@host> clear dhcpv6 relay binding interface ae0
```

The following command uses the *interfaces-wildcard* option to clear all DHCPv6 relay

clear dhcpv6 relay agent bindings over a specific interface:
binding
<interfaces-wildcard> user@host> clear dhcpv6 relay binding ge-1/0/0.*

clear dhcpv6 relay statistics

| | |
|---------------------------------|--|
| Syntax | <code>clear dhcpv6 relay statistics</code> <code><logical-system <i>logical-system-name</i>></code> <code><routing-instance <i>routing-instance-name</i>></code> |
| Release Information | Command introduced in Junos OS Release 11.4. Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Switches. |
| Description | Clear all Dynamic Host Configuration Protocol for IPv6 (DHCPv6) relay statistics. |
| Options | logical-system <i>logical-system-name</i> —(Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are cleared for the default logical system. routing-instance <i>routing-instance-name</i> —(Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are cleared for the default routing instance. |
| Required Privilege Level | view |
| List of Sample Output | clear dhcpv6 relay statistics on page 742 |
| Output Fields | See show dhcpv6 relay statistics for an explanation of output fields. |

Sample Output

`clear dhcpv6 relay statistics`

The following sample output displays the DHCPv6 relay statistics before and after the `clear dhcpv6 relay statistics` command is issued.

```
user@host> show dhcpv6 relay statistics
DHCPv6 Packets dropped:
  Total 0
```

```
Messages received:
  DHCPV6_DECLINE 0
  DHCPV6_SOLICIT 10
  DHCPV6_INFORMATION_REQUEST 0
  DHCPV6_RELEASE 0
  DHCPV6_REQUEST 10
  DHCPV6_CONFIRM 0
  DHCPV6_RENEW 0
  DHCPV6_REBIND 0
  DHCPV6_RELAY_REPL 0
```

```
Messages sent:
  DHCPV6_ADVERTISE 0
  DHCPV6_REPLY 0
  DHCPV6_RECONFIGURE 0
  DHCPV6_RELAY_FORW 0
```

```
user@host> clear dhcpv6 relay statistics
user@host> show dhcpv6 relay statistics
DHCPv6 Packets dropped:
  Total 0
```

```
Messages received:
  DHCPV6_DECLINE 0
  DHCPV6_SOLICIT 0
  DHCPV6_INFORMATION_REQUEST 0
  DHCPV6_RELEASE 0
  DHCPV6_REQUEST 0
  DHCPV6_CONFIRM 0
  DHCPV6_RENEW 0
  DHCPV6_REBIND 0
  DHCPV6_RELAY_REPL 0
```

```
Messages sent:
  DHCPV6_ADVERTISE 0
  DHCPV6_REPLY 0
  DHCPV6_RECONFIGURE 0
  DHCPV6_RELAY_FORW 0
```


clear helper statistics

| | |
|---------------------------------|--|
| Syntax | clear helper statistics |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Clear statistic counters in the User Datagram Protocol (UDP) forwarding process. |
| Options | This command has no options. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">• show helper statistics on page 762 |
| List of Sample Output | clear helper statistics on page 744 |
| Output Fields | See show helper statistics for an explanation of output fields. |

Sample Output

clear helper statistics The following sample output displays statistics counters before and after the **clear helper statistics** command is issued:

```
user@host> show helper statistics
domain:
  Received packets: 63
  Forwarded packets: 61
  Dropped packets: 2
    Due to no interface in fud database: 0
    Due to an error during packet read: 1
    Due to an error during packet send: 1
tftp:
  Received packets: 5
  Forwarded packets: 5
  Dropped packets: 0
    Due to no interface in fud database: 0
    Due to an error during packet read: 0
    Due to an error during packet send: 0
```

```
user@host> clear helper statistics
```

```
user@host> show helper statistics
domain:
  Received packets: 0
  Forwarded packets: 0
  Dropped packets: 0
    Due to no interface in fud database: 0
    Due to an error during packet read: 0
    Due to an error during packet send: 0
tftp:
  Received packets: 0
  Forwarded packets: 0
  Dropped packets: 0
    Due to no interface in fud database: 0
    Due to an error during packet read: 0
    Due to an error during packet send: 0
```

show dhcp relay binding

| | |
|----------------------------|--|
| Syntax | <pre> show dhcp relay binding <address> <brief> <detail> <interface interface-name> <interfaces-vlan> <interfaces-wildcard> <ip-address mac-address> <logical-system logical-system-name> <routing-instance routing-instance-name> <summary> </pre> |
| Release Information | <p>Command introduced in Junos OS Release 8.3.</p> <p>Options interface and mac-address added in Junos OS Release 8.4.</p> <p>Options interfaces-vlan and interfaces-wildcard added in Junos OS Release 12.1.</p> <p>Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Switches.</p> |
| Description | Display the address bindings in the Dynamic Host Configuration Protocol (DHCP) client table. |
| Options | <p>address—(Optional) Display DHCP binding information for a specific client identified by one of the following entries:</p> <ul style="list-style-type: none"> • <i>ip-address</i>—The specified IP address. • <i>mac-address</i>—The specified MAC address. • <i>session-id</i>—The specified session ID. <p>brief—(Optional) Display brief information about the active client bindings. This is the default, and produces the same output as show dhcp relay binding.</p> <p>detail—(Optional) Display detailed client binding information.</p> <p>interface interface-name—(Optional) Perform this operation on the specified interface. You can optionally filter on VLAN ID and SVLAN ID.</p> <p>interfaces-vlan—(Optional) Show the binding state information on the interface VLAN ID and S-VLAN ID.</p> <p>interfaces-wildcard—(Optional) The set of interfaces on which to show binding state information. This option supports the use of the wildcard character (*).</p> <p>logical-system logical-system-name—(Optional) Perform this operation on the specified logical system.</p> <p>routing-instance routing-instance-name—(Optional) Perform this operation on the specified routing instance.</p> <p>summary—(Optional) Display a summary of DHCP client information.</p> |

Required Privilege Level view

Related Documentation

- Clearing DHCP Bindings for Subscriber Access
- [clear dhcp relay binding on page 732](#)

List of Sample Output

- [show dhcp relay binding on page 748](#)
- [show dhcp relay binding detail on page 748](#)
- [show dhcp relay binding interface on page 748](#)
- [show dhcp relay binding interface vlan-id on page 748](#)
- [show dhcp relay binding interface svlan-id on page 748](#)
- [show dhcp relay binding ip-address on page 749](#)
- [show dhcp relay binding mac-address on page 749](#)
- [show dhcp relay binding session-id on page 749](#)
- [show dhcp relay binding <interfaces-vlan> on page 749](#)
- [show dhcp relay binding <interfaces-wildcard> on page 749](#)
- [show dhcp relay binding summary on page 749](#)

Output Fields Table 168 on page 746 lists the output fields for the **show dhcp relay binding** command. Output fields are listed in the approximate order in which they appear.

Table 168: show dhcp relay binding Output Fields

| Field Name | Field Description | Level of Output |
|--|--|--------------------|
| <i>number</i> clients, (<i>number</i> init, <i>number</i> bound, <i>number</i> selecting, <i>number</i> requesting, <i>number</i> renewing, <i>number</i> rebinding, <i>number</i> releasing) | Summary counts of the total number of DHCP clients and the number of DHCP clients in each state. | summary |
| IP address | IP address of the DHCP client. | briefdetail |
| Session Id | Session ID of the subscriber session. | briefdetail |
| Hardware address | Hardware address of the DHCP client. | briefdetail |
| Expires | Number of seconds in which the lease expires. | briefdetail |
| State | State of the DHCP relay address binding table on the DHCP client: <ul style="list-style-type: none"> • BOUND—Client has an active IP address lease. • INIT—Initial state. • REBINDING—Client is broadcasting a request to renew the IP address lease. • RELEASE—Client is releasing the IP address lease. • RENEWING—Client is sending a request to renew the IP address lease. • REQUESTING—Client is requesting a DHCP server. • SELECTING—Client is receiving offers from DHCP servers. | briefdetail |

Table 168: show dhcp relay binding Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|----------------------------------|---|-----------------|
| Interface | Incoming client interface. | brief |
| Lease Expires | Date and time at which the client's IP address lease expires. | detail |
| Lease Expires in | Number of seconds in which the lease expires. | detail |
| Lease Start | Date and time at which the client's IP address lease started. | detail |
| Incoming Client Interface | Client's incoming interface. | detail |
| Server IP Address | IP address of the DHCP server. | detail |
| Server Interface | Interface of the DHCP server. | detail |
| Bootp Relay Address | IP address of BOOTP relay. | detail |
| Type | Type of DHCP packet processing performed on the router: <ul style="list-style-type: none"> • active—Router actively processes and relays DHCP packets. • passive—Router passively snoops DHCP packets passing through the router. | All levels |
| Lease expires at | Date and time at which the client's IP address lease expires. | All levels |

Sample Output

show dhcp relay binding

```
user@host> show dhcp relay binding
```

| IP address | Session Id | Hardware address | Expires | State | Interface |
|--------------|------------|-------------------|---------|-------|------------|
| 100.20.32.11 | 41 | 00:10:94:00:00:01 | 86371 | BOUND | ge-1/0/0.0 |
| 100.20.32.12 | 42 | 00:10:94:00:00:02 | 86371 | BOUND | ge-1/0/0.0 |
| 100.20.32.13 | 43 | 00:10:94:00:00:03 | 86371 | BOUND | ge-1/0/0.0 |
| 100.20.32.14 | 44 | 00:10:94:00:00:04 | 86371 | BOUND | ge-1/0/0.0 |
| 100.20.32.15 | 45 | 00:10:94:00:00:05 | 86371 | BOUND | ge-1/0/0.0 |

show dhcp relay binding detail

```
user@host> show dhcp relay binding detail
```

Client IP Address: 100.20.32.11

```

Hardware Address: 00:10:94:00:00:01
State: BOUND(DHCP_RELAY_STATE_BOUND_ON_INTF_DELETE)
Lease Expires: 2009-07-21 11:00:06 PDT
Lease Expires in: 86361 seconds
Lease Start: 2009-07-20 11:00:06 PDT
Last Packet Received: 2009-07-20 11:00:06 PDT
Incoming Client Interface: ge-1/0/0.0
Server Ip Address: 100.20.22.2
Server Interface: none
Bootp Relay Address: 100.20.32.2
Session Id: 41

```

Client IP Address: 100.20.32.12

```

Hardware Address: 00:10:94:00:00:02
State: BOUND(DHCP_RELAY_STATE_BOUND_ON_INTF_DELETE)
Lease Expires: 2009-07-21 11:00:06 PDT
Lease Expires in: 86361 seconds
Lease Start: 2009-07-20 11:00:06 PDT
Last Packet Received: 2009-07-20 11:00:06 PDT
Incoming Client Interface: ge-1/0/0.0
Server Ip Address: 100.20.22.2
Server Interface: none
Bootp Relay Address: 100.20.32.2
Session Id: 42

```

show dhcp relay binding interface

```
user@host> show dhcp relay binding interface fe-0/0/2
```

| IP address | Hardware address | Type | Lease expires at |
|-------------|-------------------|--------|-------------------------|
| 100.20.32.1 | 90:00:00:01:00:01 | active | 2007-03-27 15:06:20 EDT |

show dhcp relay binding interface vlan-id

```
user@host> show dhcp relay binding interface ge-1/1/0:100
```

| IP address | Session Id | Hardware address | Expires | State | Interface |
|--------------|------------|-------------------|---------|-------|--------------|
| 200.20.20.15 | 6 | 00:10:94:00:00:01 | 86124 | BOUND | ge-1/1/0:100 |

```
user@host> show dhcp relay binding interface ge-1/1/0:10-100
```

```
show dhcp relay
binding interface
svlan-id
```

| IP address | Session Id | Hardware address | Expires | State | Interface |
|-----------------|------------|-------------------|---------|-------|-----------|
| 200.20.20.16 | 7 | 00:10:94:00:00:02 | 86124 | BOUND | |
| ge-1/1/0:10-100 | | | | | |

```
show dhcp relay
binding ip-address
```

```
user@host> show dhcp relay binding 100.20.32.13
```

| IP address | Session Id | Hardware address | Expires | State | Interface |
|--------------|------------|-------------------|---------|-------|------------|
| 100.20.32.13 | 43 | 00:10:94:00:00:03 | 86293 | BOUND | ge-1/0/0.0 |

```
show dhcp relay
binding mac-address
```

```
user@host> show dhcp relay binding 00:10:94:00:00:05
```

| IP address | Session Id | Hardware address | Expires | State | Interface |
|--------------|------------|-------------------|---------|-------|------------|
| 100.20.32.15 | 45 | 00:10:94:00:00:05 | 86279 | BOUND | ge-1/0/0.0 |

```
show dhcp relay
binding session-id
```

```
user@host> show dhcp relay binding 41
```

| IP address | Session Id | Hardware address | Expires | State | Interface |
|--------------|------------|-------------------|---------|-------|------------|
| 100.20.32.11 | 41 | 00:10:94:00:00:01 | 86305 | BOUND | ge-1/0/0.0 |

```
show dhcp relay
binding
<interfaces-vlan>
```

```
user@host> show dhcp relay binding ge-1/0/0:100-200
```

| IP address | Session Id | Hardware address | Expires | State | Interface |
|---------------------|------------|-------------------|---------|-------|-----------|
| 192.168.0.17 | 42 | 00:10:94:00:00:02 | 86346 | BOUND | |
| ge-1/0/0.1073741827 | | | | | |
| 192.168.0.16 | 41 | 00:10:94:00:00:01 | 86346 | BOUND | |
| ge-1/0/0.1073741827 | | | | | |

```
show dhcp relay
binding
<interfaces-wildcard>
```

```
user@host> show dhcp relay binding ge-1/3/*
```

| IP address | Session Id | Hardware address | Expires | State | Interface |
|--------------|------------|-------------------|---------|-------|-----------|
| 192.168.0.9 | 24 | 00:10:94:00:00:04 | 86361 | BOUND | |
| ge-1/3/0.110 | | | | | |
| 192.168.0.8 | 23 | 00:10:94:00:00:03 | 86361 | BOUND | |
| ge-1/3/0.110 | | | | | |
| 192.168.0.7 | 22 | 00:10:94:00:00:02 | 86361 | BOUND | |
| ge-1/3/0.110 | | | | | |

```
show dhcp relay
binding summary
```

```
user@host> show dhcp relay binding summary
```

3 clients, (2 init, 1 bound, 0 selecting, 0 requesting, 0 renewing, 0 rebinding, 0 releasing)

show dhcp relay statistics

| | |
|--------------------------|--|
| Syntax | <code>show dhcp relay statistics</code> <code><logical-system <i>logical-system-name</i>></code> <code><routing-instance <i>routing-instance-name</i>></code> |
| Syntax | Syntax for EX Series switches: <code>show dhcp relay statistics</code> <code><routing-instance <i>routing-instance-name</i>></code> |
| Release Information | Command introduced in Junos OS Release 8.3. Command introduced in Junos OS Release 12.1 for EX Series switches. Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Switches. |
| Description | Display Dynamic Host Configuration Protocol (DHCP) relay statistics. |
| Options | <code>logical-system <i>logical-system-name</i></code> —(On routers only) (Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are displayed for the default logical system. <code>routing-instance <i>routing-instance-name</i></code> —(Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are displayed for the default routing instance. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">• clear dhcp relay statistics on page 734 |
| List of Sample Output | show dhcp relay statistics on page 752 |
| Output Fields | Table 169 on page 751 lists the output fields for the <code>show dhcp relay statistics</code> command. Output fields are listed in the approximate order in which they appear. |

Table 169: show dhcp relay statistics Output Fields

| Field Name | Field Description |
|--------------------------|---|
| Packets dropped | <p>Number of packets discarded by the extended DHCP relay agent application due to errors. Only nonzero statistics appear in the Packets dropped output. When all of the Packets dropped statistics are 0 (zero), only the Total field appears.</p> <ul style="list-style-type: none"> • Total—Total number of packets discarded by the extended DHCP relay agent application. • Bad hardware address—Number of packets discarded because an invalid hardware address was specified. • Bad opcode—Number of packets discarded because an invalid operation code was specified. • Bad options—Number of packets discarded because invalid options were specified. • Invalid server address—Number of packets discarded because an invalid server address was specified. • No available addresses—Number of packets discarded because there were no addresses available for assignment. • No interface match—Number of packets discarded because they did not belong to a configured interface. • No routing instance match—Number of packets discarded because they did not belong to a configured routing instance. • No valid local address—Number of packets discarded because there was no valid local address. • Packet too short—Number of packets discarded because they were too short. • Read error—Number of packets discarded because of a system read error. • Send error—Number of packets that the extended DHCP relay application could not send. • Option 60—Number of packets discarded containing DHCP option 60 vendor-specific information. • Option 82—Number of packets discarded because DHCP option 82 information could not be added. |
| Messages received | <p>Number of DHCP messages received.</p> <ul style="list-style-type: none"> • BOOTREQUEST—Number of BOOTP protocol data units (PDUs) received • DHCPDECLINE—Number of DHCP PDUs of type DECLINE received • DHCPDISCOVER—Number of DHCP PDUs of type DISCOVER received • DHCPINFORM—Number of DHCP PDUs of type INFORM received • DHCPRELEASE—Number of DHCP PDUs of type RELEASE received • DHCPREQUEST—Number of DHCP PDUs of type REQUEST received |
| Messages sent | <p>Number of DHCP messages sent.</p> <ul style="list-style-type: none"> • BOOTREPLY—Number of BOOTP PDUs transmitted • DHCPOFFER—Number of DHCP OFFER PDUs transmitted • DHCPACK—Number of DHCP ACK PDUs transmitted • DHCPNACK—Number of DHCP NACK PDUs transmitted • DHCPFORCERENEW—Number of DHCP FORCERENEW PDUs transmitted |
| Packets forwarded | <p>Number of packets forwarded.</p> <ul style="list-style-type: none"> • BOOTREQUEST—Number of BOOTREQUEST protocol data units (PDUs) forwarded • BOOTREPLY—Number of BOOTREPLY protocol data units (PDUs) forwarded |

Sample Output

**show dhcp relay
statistics**

```
user@host> show dhcp relay statistics
Packets dropped:
  Total                  30
  Bad hardware address   1
  Bad opcode             1
  Bad options            3
  Invalid server address  5
  No available addresses  1
  No interface match     2
  No routing instance match 9
  No valid local address  4
  Packet too short       2
  Read error             1
  Send error             1
  Option 60              1
  Option 82              2

Messages received:
  BOOTREQUEST           116
  DHCPDECLINE           0
  DHCPDISCOVER          11
  DHCPINFORM            0
  DHCPRELEASE           0
  DHCPREQUEST           105

Messages sent:
  BOOTREPLY             0
  DHCPOFFER             2
  DHCPACK               1
  DHCPNAK               0
  DHCPFORCERENEW        0

Packets forwarded:
  Total                  4
  BOOTREQUEST           2
  BOOTREPLY             2
```

show dhcpv6 relay binding

| | |
|----------------------------|---|
| Syntax | <pre>show dhcpv6 relay binding <address client-id session-id> <brief> <detail> <interface interface-name> <interfaces-vlan> <interfaces-wildcard> <logical-system logical-system-name> <routing-instance routing-instance-name> <summary></pre> |
| Release Information | <p>Command introduced in Junos OS Release 11.4.</p> <p>Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Switches.</p> <p><i>interfaces-vlan</i> and <i>interfaces-wildcard</i> options introduced in Junos OS Release 12.1.</p> |
| Description | Display the DHCPv6 address bindings in the Dynamic Host Configuration Protocol (DHCP) client table. |
| Options | <p>address—(Optional) Clear the binding state for the DHCPv6 client, using one of the following entries:</p> <ul style="list-style-type: none"> • <i>CID</i>—The specified Client ID (CID). • <i>ipv6-prefix</i>—The specified IPv6 prefix. • <i>session-id</i>—The specified session ID. <p>brief—(Optional) Display brief information about the active client bindings. This is the default, and produces the same output as show dhcpv6 relay binding.</p> <p>detail—(Optional) Display detailed client binding information.</p> <p>interface interface-name—(Optional) Perform this operation on the specified interface. You can optionally filter on VLAN ID and S-VLAN ID.</p> <p>interfaces-vlan—(Optional) Show the binding state information on the interface VLAN ID and S-VLAN ID.</p> <p>interfaces-wildcard—(Optional) The set of interfaces on which to show binding state information. This option supports the use of the wildcard character (*).</p> <p>logical-system logical-system-name—(Optional) Perform this operation on the specified logical system.</p> <p>routing-instance routing-instance-name—(Optional) Perform this operation on the specified routing instance.</p> <p>summary—(Optional) Display a summary of DHCPv6 client information.</p> |

Required Privilege Level view

Related Documentation

- Clearing DHCP Bindings for Subscriber Access
- [clear dhcpv6 relay binding on page 737](#)

List of Sample Output

- [show dhcpv6 relay binding on page 756](#)
- [show dhcpv6 relay binding \(address\) on page 756](#)
- [show dhcpv6 relay binding \(client-id\) on page 756](#)
- [show dhcpv6 relay binding detail on page 756](#)
- [show dhcpv6 relay binding detail \(Multi-Relay Topology\) on page 757](#)
- [show dhcpv6 relay binding \(session-id\) on page 758](#)
- [show dhcpv6 relay binding \(interfaces-vlan\) on page 758](#)
- [show dhcpv6 relay binding \(interfaces-wildcard\) on page 758](#)
- [show dhcpv6 relay binding \(interfaces-wildcard\) on page 758](#)
- [show dhcpv6 relay binding summary on page 758](#)

Output Fields Table 170 on page 754 lists the output fields for the **show dhcpv6 relay binding** command. Output fields are listed in the approximate order in which they appear.

Table 170: show dhcpv6 relay binding Output Fields

| Field Name | Field Description | Level of Output |
|--|--|--------------------|
| <i>number</i> clients, (<i>number</i> init, <i>number</i> bound, <i>number</i> selecting, <i>number</i> requesting, <i>number</i> renewing, <i>number</i> rebinding, <i>number</i> releasing) | Summary counts of the total number of DHCPv6 clients and the number of DHCPv6 clients in each state. | summary |
| Client IPv6 Prefix | Prefix of the DHCPv6 client. | briefdetail |
| Client DUID | DHCP for IPv6 Unique Identifier (DUID) of the client. | briefdetail |
| Session Id | Session ID of the subscriber session. | briefdetail |
| Expires | Number of seconds in which the lease expires. | briefdetail |
| State | State of the DHCPv6 relay address binding table on the DHCPv6 client: <ul style="list-style-type: none"> • BOUND—Client has an active IP address lease. • INIT—Initial state. • REBINDING—Client is broadcasting a request to renew the IP address lease. • RELEASE—Client is releasing the IP address lease. • RENEWING—Client is sending a request to renew the IP address lease. • REQUESTING—Client is requesting a DHCPv6 server. • SELECTING—Client is receiving offers from DHCPv6 servers. | briefdetail |
| Interface | Incoming client interface. | brief |

Table 170: show dhcpv6 relay binding Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-------------------------------------|--|-----------------|
| Lease Expires | Date and time at which the client's IP address lease expires. | detail |
| Lease Expires in | Number of seconds in which the lease expires. | detail |
| Lease Start | Date and time at which the client's IP address lease started. | detail |
| Incoming Client Interface | Client's incoming interface. | detail |
| Server Address | IP address of the DHCPv6 server. Displays unknown for a DHCPv6 relay agent in a multi-relay topology that is not directly adjacent to the DHCPv6 server and does not detect the IP address of the server. In that case, the output instead displays the Next Hop Server Facing Relay field. | detail |
| Next Hop Server Facing Relay | Next-hop address in the direction of the DHCPv6 server. | detail |
| Server Interface | Interface of the DHCPv6 server. | detail |
| Relay Address | IP address of the relay. | detail |
| Client Pool Name | Address pool that granted the client lease. | |
| Client ID Length | Length of client ID. | All levels |
| Client Id | Client ID. | All levels |

Sample Output

show dhcpv6 relay binding

```
user@host> show dhcpv6 relay binding
Prefix                Session Id Expires State Interface Client DUID
2001:bd8:3c4d:15::/64 1          83720 BOUND ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01
2001:bd8:3c4d:16::/64 2          83720 BOUND ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:02
2001:bd8:3c4d:17::/64 3          83720 BOUND ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:03
2001:bd8:3c4d:18::/64 4          83720 BOUND ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:04
2001:bd8:3c4d:19::/64 5          83720 BOUND ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:05
2001:bd8:3c4d:20::/64 6          83720 BOUND ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:06
```

show dhcpv6 relay binding (address)

```
user@host> show dhcpv6 relay binding 2001:bd8:1111:2222::/64 detail
Session Id: 1
  Client IPv6 Prefix:          2001:bd8:3c4d:15::/64
  Client DUID:                  LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01

  State:                        BOUND(RELAY_STATE_BOUND)
  Lease Expires:                 2011-05-25 07:12:09 PDT
  Lease Expires in:              77115 seconds
  Lease Start:                   2011-05-24 07:12:09 PDT
  Incoming Client Interface:     ge-1/0/0.0
  Server Address:                2008:aaaa:bbbb::1
  Server Interface:              none
  Relay Address:                 2001:bd8:1111:2222::
  Client Pool Name:              pool-25
  Client Id Length:              14
  Client Id:                     /0x00010001/0x4bfa26af/0x00109400/0x0001
```

show dhcpv6 relay binding (client-id)

```
user@host> show dhcpv6 relay binding 14/0x00010001/0x4bfa26af/0x00109400/0x0001 detail
Session Id: 1
  Client IPv6 Prefix:          2001:bd8:3c4d:15::/64
  Client DUID:                  LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01

  State:                        BOUND(RELAY_STATE_BOUND)
  Lease Expires:                 2011-05-25 07:12:09 PDT
  Lease Expires in:              77115 seconds
  Lease Start:                   2011-05-24 07:12:09 PDT
  Incoming Client Interface:     ge-1/0/0.0
  Server Address:                2008:aaaa:bbbb::1
  Server Interface:              none
  Relay Address:                 2001:bd8:1111:2222::
  Client Pool Name:              pool-25
  Client Id Length:              14
  Client Id:                     /0x00010001/0x4bfa26af/0x00109400/0x0001
```

show dhcpv6 relay

```
user@host> show dhcpv6 relay binding detail
Session Id: 1
```

binding detail

```

Client IPv6 Prefix:      2001:bd8:3c4d:15::/64
Client DUID:             LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01

State:                   BOUND(RELAY_STATE_BOUND)
Lease Expires:           2011-05-25 07:12:09 PDT
Lease Expires in:        77115 seconds
Lease Start:             2011-05-24 07:12:09 PDT
Incoming Client Interface: ge-1/0/0.0
Server Address:           2008:aaaa:bbbb::1
Server Interface:         none
Relay Address:            2001:bd8:1111:2222::
Client Pool Name:         pool-25
Client Id Length:         14
Client Id:                /0x00010001/0x4bfa26af/0x00109400/0x0001

```

**show dhcpv6 relay
binding detail**

```

user@host > show dhcpv6 relay binding detail
Session Id: 13
Client IPv6 Prefix:

```

```

3000:0:0:8001::5/128

```

**(Multi-Relay
Topology)**

```

Client DUID:                LL0x1-00:00:65:03:01:02
State:                      BOUND(DHCPV6_RELAY_STATE_BOUND)
Lease Expires:              2011-11-21 06:14:50 PST
Lease Expires in:          293 seconds
Lease Start:                2011-11-21 06:09:50 PST
Incoming Client Interface:  ge-1/0/0.0
Server Address:             unknown
Next Hop Server Facing Relay: 4000::2
Server Interface:           none
Client Id Length:          10
Client Id:                  /0x00030001/0x00006503/0x0102

```

**show dhcpv6 relay
binding (session-id)**

```

user@host> show dhcpv6 relay binding 41
Prefix                Session Id  Expires  State  Interface  Client DUID
2001:bd8:3c4d:15::/64  41        78837    BOUND  ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01

```

**show dhcpv6 relay
binding
(interfaces-vlan)**

```

user@host> show dhcpv6 relay binding ge-1/0/0:100-200
Prefix                Session Id  Expires  State  Interface  Client DUID
2001:DB8::/32        11        87583    BOUND  ge-1/0/0.1073741827
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB9::/32        12        87583    BOUND  ge-1/0/0.1073741827
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01

```

**show dhcpv6 relay
binding
(interfaces-wildcard)**

```

user@host> show dhcpv6 relay binding demux0
Prefix                Session Id  Expires  State  Interface  Client DUID
2001:DB8::/32        30        79681    BOUND  demux0.1073741824
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB9::/32        31        79681    BOUND  demux0.1073741825
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:CB9::/32        32        79681    BOUND  demux0.1073741826
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01

```

**show dhcpv6 relay
binding
(interfaces-wildcard)**

```

user@host> show dhcpv6 relay binding ge-1/3/*
Prefix                Session Id  Expires  State  Interface  Client DUID
2001:DB8::/32        22        79681    BOUND  ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB9::/32        33        79681    BOUND  ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:CB9::/32        24        79681    BOUND  ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01

```

**show dhcpv6 relay
binding summary**

```

user@host> show dhcpv6 relay binding summary
5 clients, (0 init, 5 bound, 0 selecting, 0 requesting, 0 renewing, 0 releasing)

```


show dhcpv6 relay statistics

| | |
|---------------------------------|---|
| Syntax | <code>show dhcpv6 relay statistics</code> <code><logical-system <i>logical-system-name</i>></code> <code><routing-instance <i>routing-instance-name</i>></code> |
| Release Information | Command introduced in Junos OS Release 11.4. Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Switches. |
| Description | Display Dynamic Host Configuration Protocol for IPv6 (DHCPv6) relay statistics. |
| Options | <p>logical-system <i>logical-system-name</i>—(Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are displayed for the default logical system.</p> <p>routing-instance <i>routing-instance-name</i>—(Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are displayed for the default routing instance.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • clear dhcpv6 relay statistics on page 741 |
| List of Sample Output | show dhcpv6 relay statistics on page 761 |
| Output Fields | Table 171 on page 759 lists the output fields for the show dhcpv6 relay statistics command. Output fields are listed in the approximate order in which they appear. |

Table 171: show dhcpv6 relay statistics Output Fields

| Field Name | Field Description |
|------------------------|---|
| DHCPv6 Packets dropped | <p>Number of packets discarded by the extended DHCPv6 relay agent application due to errors. Only nonzero statistics appear in the Packets dropped output. When all of the Packets dropped statistics are 0 (zero), only the Total field appears.</p> <ul style="list-style-type: none"> • Total—Total number of packets discarded by the DHCPV6 relay agent application. • Bad options—Number of packets discarded because invalid options were specified. • Bad send—Number of packets that the extended DHCP relay application could not send. • Bad src address—Number of packets discarded because the family type was not AF_INET6. • No client id—Number of packets discarded because they could not be matched to a client. • No safd—Number of packets discarded because they arrived on an unconfigured interface. • Short packet—Number of packets discarded because they were too short. • Relay hop count—Number of packets discarded because the hop count in the packet exceeded 32. |

Table 171: show dhcpv6 relay statistics Output Fields (*continued*)

| Field Name | Field Description |
|--------------------------|---|
| Messages received | <p>Number of DHCPv6 messages received.</p> <ul style="list-style-type: none"> • DHCPV6_DECLINE—Number of DHCPv6 PDUs of type DECLINE received • DHCPV6_SOLICIT—Number of DHCPv6 PDUs of type SOLICIT received • DHCPV6_INFORMATION_REQUEST—Number of DHCPv6 PDUs of type INFORMATION-REQUEST received • DHCPV6_RELEASE—Number of DHCPv6 PDUs of type RELEASE received • DHCPV6_REQUEST—Number of DHCPv6 PDUs of type REQUEST received • DHCPV6_CONFIRM—Number of DHCPv6 PDUs of type CONFIRM received • DHCPV6_RENEW—Number of DHCPv6 PDUs of type RENEW received • DHCPV6_REBIND—Number of DHCPv6 PDUs of type REBIND received • DHCPV6_RELAY_REPL—Number of DHCPv6 PDUs of type RELAY-REPL received |
| Messages sent | <p>Number of DHCPv6 messages sent.</p> <ul style="list-style-type: none"> • DHCPV6_ADVERTISE—Number of DHCPv6 ADVERTISE PDUs transmitted • DHCP_REPLY—Number of DHCPv6 REPLY PDUs transmitted • DHCP_RECONFIGURE—Number of DHCPv6 RECONFIGURE PDUs transmitted • DHCP_RELAY_FORW—Number of DHCPv6 RELAY-FORW PDUs transmitted |
| Packets forwarded | <p>Number of packets forwarded by the extended DHCPv6 relay agent application.</p> <ul style="list-style-type: none"> • FWD REQUEST—Number of DHCPv6 REQUEST packets forwarded • FWD REPLY—Number of DHCPv6 REPLY packets forwarded |

Sample Output

`show dhcpv6 relay statistics`

```
user@host> show dhcpv6 relay statistics
DHCPv6 Packets dropped:
  Total 0

Messages received:
  DHCPV6_DECLINE 0
  DHCPV6_SOLICIT 10
  DHCPV6_INFORMATION_REQUEST 0
  DHCPV6_RELEASE 0
  DHCPV6_REQUEST 10
  DHCPV6_CONFIRM 0
  DHCPV6_RENEW 0
  DHCPV6_REBIND 0
  DHCPV6_RELAY_REPL 0

Messages sent:
  DHCPV6_ADVERTISE 0
  DHCPV6_REPLY 0
  DHCPV6_RECONFIGURE 0
  DHCPV6_RELAY_FORW 0

Packets forwarded:
  Total 4
  FWD REQUEST 2
  FWD REPLY 2
```

show helper statistics

| | |
|---------------------------------|--|
| Syntax | show helper statistics |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Show statistics collected by the UDP forwarding process. |
| Options | This command has no options. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • clear helper statistics on page 743 |
| List of Sample Output | show helper statistics on page 763 |
| Output Fields | Table 172 on page 762 lists the output fields for the show helper statistics command. Output fields are listed in the approximate order in which they appear. |

Table 172: show helper statistics Output Fields

| Field Name | Field Description |
|------------|--|
| domain | <p>Statistics for Domain Name System (DNS) forwarding:</p> <ul style="list-style-type: none"> • Received packets—Packets received for this service. • Forwarded packets—Packets forwarded for this service. • Dropped packets—Total number of packets dropped for this service. • Due to no interface in fud database—Number of packets dropped because the packet came in on an interface that the UDP forwarding process did not identify as active. • Due to an error during packet read—Number of packets dropped because an error occurred when the packet was read from the wire. • Due to an error during packet send—Number of packets dropped because an error occurred when the packet was sent to the wire. |

Table 172: show helper statistics Output Fields (*continued*)

| Field Name | Field Description |
|------------|--|
| tftp | <p>Statistics for Trivial File Transfer Protocol (TFTP) forwarding:</p> <ul style="list-style-type: none"> • Received packets—Packets received for this service. • Forwarded packets—Packets forwarded for this service. • Dropped packets—Total number of packets dropped for this service. <p>Reasons for dropped packets include:</p> <ul style="list-style-type: none"> • Due to no interface in fud database—Number of packets dropped because the packet came in on an interface that the UDP forwarding process did not identify as active. • Due to no matching routing instance—Number of packets dropped because the packet had no matching routing instance. • Due to an error during packet read—Number of packets dropped because an error occurred when the packet was read from the wire. • Due to an error during packet send—Number of packets dropped because an error occurred when the packet was sent to the wire. • Due to invalid server address—Number of packets dropped because the packet contained an invalid server address. • Due to no valid local address—Number of packets dropped because the packet contained no local address. • Due to no route to server/client—Number of packets dropped because the packet contained no route to the server or the client. |

Sample Output

```

show helper statistics  user@host> show helper statistics
                        domain: Received packets: 0
                        Forwarded packets: 0
                        Dropped packets: 0
                          Due to no interface in fud database: 0
                          Due to an error during packet read: 0
                          Due to an error during packet send: 0
tftp: Received packets: 0
tftp: Forwarded packets: 0
tftp: Dropped packets: 0
  Due to no interface in fud database: 0
  Due to no matching routing instance: 0
  Due to an error during packet read: 0
  Due to an error during packet send: 0
  Due to invalid server address: 0
  Due to no valid local address: 0
  Due to no route to server/client: 0

```


Routing Policy Operational Mode Commands

Table 173 on page 765 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot routing policy filters. Commands are listed in alphabetical order.

Table 173: Routing Policy Operational Mode Commands

| Task | Command |
|---|--|
| Display configured routing policies. | show policy |
| Display configured policy conditions and associated routes. | show policy conditions |
| Test import and export policies. | test policy |



NOTE: For information about how to configure routing policy filters, see the *Junos Policy Framework Configuration Guide*.

show policy

| | |
|------------------------------------|--|
| Syntax | show policy <logical-system (all <i>logical-system-name</i>)> < <i>policy-name</i> > |
| Syntax (EX Series Switches) | show policy < <i>policy-name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. |
| Description | Display information about configured routing policies. |
| Options | <p>none—List the names of all configured routing policies.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><i>policy-name</i>—(Optional) Show the contents of the specified policy.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • show policy damping on page 74 |
| List of Sample Output | show policy on page 767 show policy policy-name on page 767 show policy (Multicast Scoping) on page 767 |
| Output Fields | Table 174 on page 766 lists the output fields for the show policy command. Output fields are listed in the approximate order in which they appear. |

Table 174: show policy Output Fields

| Field Name | Field Description |
|--------------------|---------------------------------|
| <i>policy-name</i> | Name of the policy listed. |
| <i>term</i> | Policy term listed. |
| <i>from</i> | Match condition for the policy. |
| <i>then</i> | Action for the policy. |

Sample Output

show policy

```
user@host> show policy
Configured policies:
__vrf-export-red-internal__
__vrf-import-red-internal__
red-export
all_routes
```

**show policy
policy-name**

```
user@host> show policy test-statics
Policy test-statics:
  from
    3.0.0.0/8  accept
    3.1.0.0/16  accept
  then reject
```

**show policy (Multicast
Scoping)**

```
user@host> show policy test-statics
Policy test-statics:
  from
    multicast-scoping == 8
```

show policy conditions

| | |
|------------------------------------|--|
| Syntax | <pre>show policy conditions <condition-name> <detail> <dynamic> <logical-system (all logical-system-name)></pre> |
| Syntax (EX Series Switches) | <pre>show policy conditions <condition-name> <detail> <dynamic></pre> |
| Release Information | <p>Command introduced in Junos OS Release 9.0.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> |
| Description | <p>Display all the configured conditions as well as the routing tables with which the configuration manager is interacting. If the detail keyword is included, the output also displays dependent routes for each condition.</p> |
| Options | <p>none—Display all configured conditions and associated routing tables.</p> <p>condition-name—(Optional) Display information about the specified condition only.</p> <p>detail—(Optional) Display the specified level of output.</p> <p>dynamic—(Optional) Display information about the conditions in the dynamic database.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show policy conditions detail on page 769 |
| Output Fields | <p>Table 175 on page 768 lists the output fields for the show policy conditions command. Output fields are listed in the approximate order in which they appear.</p> |

Table 175: show policy conditions Output Fields

| Field Name | Field Description | Level of Output |
|-------------------------|--|-----------------|
| Condition | Name of configured condition. | All levels |
| event | Condition type. If the if-route-exists option is configured, the event type is: Existence of a route in a specific routing table. | All levels |
| Dependent routes | List of routes dependent on the condition, along with the latest generation number. | detail |
| Condition tables | List of routing tables associated with the condition, along with the latest generation number and number of dependencies. | All levels |

Table 175: show policy conditions Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|----------------------------|---|-----------------|
| If-route-exists conditions | List of conditions configured to look for a route in the specified table. | All levels |

Sample Output

show policy conditions detail

```

user@host> show policy conditions detail
Configured conditions:
Condition cond1, event: Existence of a route in a specific routing table
Dependent routes:
  4.4.4.4/32, generation 3
  6.6.6.6/32, generation 3
  10.10.10.10/32, generation 3

Condition cond2, event: Existence of a route in a specific routing table
Dependent routes:
None

Condition tables:
Table inet.0, generation 4, dependencies 3, If-route-exists conditions: cond1
(static) cond2 (static)

```

test policy

| | |
|---------------------------------|--|
| Syntax | <code>test policy <i>policy-name</i> <i>prefix</i></code> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. |
| Description | Test a policy configuration to determine which prefixes match routes in the routing table. |
| Options | <i>policy-name</i> —Name of a policy. <i>prefix</i> —Destination prefix to match. |
| Additional Information | All prefixes in the default unicast routing table (inet.0) that match prefixes that are the same as or longer than the specific prefix are processed by the from clause in the specified policy. All prefixes accepted by the policy are displayed. The test policy command evaluates a policy differently from the BGP import process. When testing a policy that contains an interface match condition in the from clause, the test policy command uses the match condition. In contrast, BGP does not use the interface match condition when evaluating the policy against routes learned from internal BGP (IBGP) or external BGP (EGBP) multihop peers. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">• show policy damping on page 74 |
| List of Sample Output | test policy on page 771 |
| Output Fields | For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command. |

Sample Output

test policy

```
user@host> test policy test-statics 3.0.0.1/8
inet.0: 44 destinations, 44 routes (44 active, 0 holddown, 0 hidden)
Prefixes passing policy:

3.0.0.0/8      *[BGP/170] 16:22:46, localpref 100, from 10.255.255.41
               AS Path: 50888 I
               > to 10.11.4.32 via en0.2, label-switched-path l2
3.3.3.1/32    *[IS-IS/18] 2d 00:21:46, metric 0, tag 2
               > to 10.0.4.7 via fxp0.0
3.3.3.2/32    *[IS-IS/18] 2d 00:21:46, metric 0, tag 2
               > to 10.0.4.7 via fxp0.0
3.3.3.3/32    *[IS-IS/18] 2d 00:21:46, metric 0, tag 2
               > to 10.0.4.7 via fxp0.0
3.3.3.4/32    *[IS-IS/18] 2d 00:21:46, metric 0, tag 2
               > to 10.0.4.7 via fxp0.0
Policy test-statics: 5 prefixes accepted, 0 prefixes rejected
```


PART 3

MPLS

- [LDP Operational Mode Commands on page 775](#)
- [MPLS Operational Mode Commands on page 807](#)
- [RSVP Operational Mode Commands on page 865](#)

CHAPTER 18

LDP Operational Mode Commands

Table 176 on page 775 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot the Label Distribution Protocol (LDP). Commands are listed in alphabetical order.

Table 176: LDP Operational Mode Commands

| Task | Command |
|---|--|
| Clear LDP neighbors. | <code>clear ldp neighbor</code> |
| Clear LDP sessions. | <code>clear ldp session</code> |
| Clear LDP statistics. | <code>clear ldp statistics</code> |
| Display entries in the LDP database. | <code>show ldp database</code> |
| Display forwarding equivalence class filters. | <code>show ldp fec-filters</code> |
| Display the status of interfaces on which LDP is running. | <code>show ldp interface</code> |
| Display LDP neighbors. | <code>show ldp neighbor</code> |
| Display the configured named paths that are used by LDP. | <code>show ldp path</code> |
| Display LDP routing table entries. | <code>show ldp route</code> |
| Display currently active LDP sessions. | <code>show ldp session</code> |
| Display LDP statistics. | <code>show ldp statistics</code> |
| Display LDP traffic statistics. | <code>show ldp traffic-statistics</code> |



.....

NOTE: For more LDP-related commands, such as `show route protocol`, `show route instance`, and `show route table`, see Protocol-Independent Routing Operational Mode Commands.

For information about how to configure LDP, see the *Junos MPLS Applications Configuration Guide*.

.....

clear ldp neighbor

| | |
|---------------------------------|--|
| Syntax | clear ldp neighbor <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> < <i>neighbor</i> > |
| Description | Tear down Label Distribution Protocol (LDP) neighbor connections. |
| Options | <p>none—Tear down connections with all LDP neighbors for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Clear the LDP session for the specified routing instance only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><i>neighbor</i>—(Optional) Clear an LDP session for the specified neighbor (IP address) only.</p> |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none"> • show ldp neighbor on page 787 |
| List of Sample Output | clear ldp neighbor on page 777 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear ldp neighbor user@host> clear ldp neighbor

clear ldp session

| | |
|---------------------------------|--|
| Syntax | <code>clear ldp session</code> <code><destination></code> <code><instance <i>instance-name</i>></code> <code><logical-system (all <i>logical-system-name</i>)></code> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Clear Label Distribution Protocol (LDP) sessions. |
| Options | <p>none—Clear LDP sessions for all destinations for all routing instances.</p> <p><i>destination</i>—(Optional) Clear an LDP session for the specified destination (IP address).</p> <p><i>instance instance-name</i>—(Optional) Clear the LDP session for the specified routing instance only.</p> <p><i>logical-system (all logical-system-name)</i>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• show ldp session on page 795 |
| List of Sample Output | clear ldp session on page 778 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

`clear ldp session` `user@host> clear ldp session`

clear ldp statistics

| | |
|---------------------------------|--|
| Syntax | clear ldp statistics <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Set all Label Distribution Protocol (LDP) statistics to zero. |
| Options | <p>none—Set all LDP statistics to zero for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Clear the LDP session for the specified routing instance only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none"> • show ldp statistics on page 801 • show ldp traffic-statistics on page 805 |
| List of Sample Output | clear ldp statistics on page 779 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear ldp statistics user@host> clear ldp statistics

show ldp database

| | |
|---------------------------------|--|
| Syntax | <pre>show ldp database <brief detail extensive> <inet l2circuit> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <session <i>session</i>></pre> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display entries in the Label Distribution Protocol (LDP) database. |
| Options | <p>none—Display standard information about all entries in the LDP database for all routing instances.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>inet l2circuit—(Optional) Display only IPv4 or Layer 2 circuit bindings.</p> <p>instance <i>instance-name</i>—(Optional) Display routing instance information for the specified instance only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>session <i>session</i>—(Optional) Display database for the specified session only. <i>session</i> is the destination address of the LDP session.</p> |
| Required Privilege Level | view |
| List of Sample Output | show ldp database on page 783 show ldp database l2circuit detail on page 783 show ldp database session on page 783 |
| Output Fields | Table 177 on page 780 describes the output fields for the show ldp database command. Output fields are listed in the approximate order in which they appear. |

Table 177: show ldp database Output Fields

| Field Name | Field Description | Level of Output |
|---------------------------|--|-----------------|
| Input label database | Label received from the other router. | All levels |
| Output label database | Label advertised to the other router. | All levels |
| <i>session-identifier</i> | Session identifier, which includes the local and remote label space identifiers. | All levels |
| Label | Label binding to a route prefix. | All levels |

Table 177: show ldp database Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------------------|--|-----------------|
| Prefix | <p>Route prefix. It can be either the IP prefix or the Layer 2 encapsulation type in the format L2CKT control word status encapsulation-type vc-number, for example, L2CKT CtlfWord FRAME RELAY VC 2</p> <ul style="list-style-type: none"> • control-word-status—Displays whether the use of the control word has been negotiated for this virtual circuit: <ul style="list-style-type: none"> • NoCtrlWord • CtrlWord • encapsulation-type—Encapsulation type: <ul style="list-style-type: none"> • FRAME RELAY • ATM AAL5 • ATM CELL • VLAN • ETHERNET • CISCO_HDLC • PPP • VC number—Virtual circuit number. It can have any numeric value. • (Stale)—When you display the LDP database for the neighbor of a restarting router, the bindings learned from the restarting neighbor are displayed as (Stale). Stale bindings are deleted if they are not refreshed within the recovery time. | All levels |
| MTU | MTU of the Layer 2 circuit. MTU is displayed for all encapsulation types except ATM cell encapsulations. | detail |
| VCCV Control Channel types | <p>Virtual Circuit Connection Verification (VCCV) control channel types</p> <ul style="list-style-type: none"> • MPLS router alert label • MPLS PW label with TTL=1 | extensive |
| VCCV Control Verification types | The only valid VCCV control verification type is LSP ping . | extensive |
| TDM payload size | Size of the Time Division Multiplex (TDM) payload. | All levels |
| TDM bitrate | Bit rate for the TDM traffic. | All levels |
| Requested VLAN ID | (VLANs) VLAN identifier of the Layer 2 circuit. | detail |
| Cell bundle size | (ATM cell encapsulations) Maximum number of cells that the Layer 2 circuit can receive in a packet. | detail |

Table 177: show ldp database Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------|---|-----------------|
| State | State of the label binding: <ul style="list-style-type: none">• Active—Label binding has been installed and distributed appropriately. A label binding is almost always in this state.• New—New label that has not yet been distributed.<ul style="list-style-type: none">• MapRcv—Waiting to receive a label mapping message.• MapSend—Waiting to send a label mapping message.• RelRcv—Waiting to receive a label release message.• RelRsnd—Waiting to receive a label release message before resending label mapping message.• RelSend—Waiting to send a label release message.• ReqSend—Waiting to send a label request message.• W/dSend—Waiting to send a label withdrawal message. | detail |
| Age | Time elapsed since the binding was created. | detail |

Sample Output

show ldp database

```
user@host> show ldp database
Input label database, 10.255.245.222:0--10.255.245.221:0
  Label    Prefix
  3        10.255.245.221/32 (Stale)
100018     10.255.245.222/32
100011     L2CKT FRAME RELAY VC 11
Output label database, 10.255.245.222:0--10.255.245.221:0
  Label    Prefix
  3        10.255.245.221/32
100018     10.255.245.222/32
100011     L2CKT FRAME RELAY VC 1
```

show ldp database l2circuit detail

```
user@host> show ldp database l2circuit detail
Input label database, 10.255.245.44:0--10.255.245.45:0
  Label    Prefix
  100176    L2CKT CtrlWord ATM CELL (VC Mode) VC 100
            Cell bundle size: 80
            State: Active
            Age: 9:48
  100256    L2CKT CtrlWord FRAME RELAY VC 101
            MTU: 4470
            State: Active
            Age: 9:48

Output label database, 10.255.245.44:0--10.255.245.45:0
  Label    Prefix
  100048    L2CKT CtrlWord ATM CELL (VC Mode) VC 100
            Cell bundle size: 80
            State: Active
            Age: 9:48
  100112    L2CKT CtrlWord FRAME RELAY VC 101
            MTU: 4470
            State: Active
            Age: 9:48
```

show ldp database session

```
user@host> show ldp database session 10.1.1.195
Input label database, 10.0.0.194:0--10.1.1.195:0
  Label    Prefix
  100002    10.255.245.197/32
  100003    10.255.245.196/32
  100004    10.0.0.194/32
  3         10.1.1.195/32
  100000    L2CKT NoCtrlWord FRAME RELAY VC 1
  100001    L2CKT CtrlWord FRAME RELAY VC 2
Output label database, 10.0.0.194:0--10.1.1.195:0
  Label    Prefix
  100003    10.255.245.197/32
  100004    10.1.1.195/32
  100002    10.255.245.196/32
  3         10.0.0.194/32
  100000    L2CKT CtrlWord FRAME RELAY VC 2
  100001    L2CKT NoCtrlWord FRAME RELAY VC 1
```

show ldp fec-filters

| | |
|---------------------------------|--|
| Syntax | show ldp fec-filters <fec> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display information about configured Label Distribution Protocol (LDP) forwarding equivalence class (FEC) filters. |
| Options | <p>fec—(Optional) Display FEC filter information for the specified FEC.</p> <p>instance <i>instance-name</i>—(Optional) Display FEC filter information for the specified instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show ldp fec-filters on page 784 |
| Output Fields | Table 178 on page 784 lists the output fields for the show ldp fec-filters command. Output fields are listed in the approximate order in which they appear. |

Table 178: show ldp fec-filters Output Fields

| Field Name | Field Description |
|------------|--|
| Ingress | Names of the FEC filters on the ingress routers. |
| Transit | Names of the FEC filters on the transit routers. |

Sample Output

```

show ldp fec-filters
user@host> show ldp fec-filters 10/8
10.22.1.2/32
  Ingress: f1-10.22.1.2/32 (index: 3)
  Transit: (null) (index: 0)

```

show ldp interface

| | |
|---------------------------------|--|
| Syntax | show ldp interface <brief detail extensive> <interface-name> <instance instance-name> <logical-system (all logical-system-name)> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display the status of Label Distribution Protocol (LDP)-enabled interfaces. |
| Options | <p>none—Display standard status information about all LDP-enabled interface for all routing instances.</p> <p>interface-name—(Optional) Display information for the specified interface.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>instance instance-name—(Optional) Display information for the specified routing instance.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show ldp interface extensive on page 786 |
| Output Fields | Table 179 on page 785 describes the output fields for the show ldp interface command. Output fields are listed in the approximate order in which they appear. |

Table 179: show ldp interface Output Fields

| Field Name | Field Description | Level of Output |
|----------------|--|---------------------|
| Interface | Interface name. | All levels |
| Label space ID | Label space identifier that the router is advertising on the interface. | All levels |
| Nbr count | Number of neighbors on the interface. | All levels |
| Next hello | How long until the next hello packet is sent on this interface, in seconds. | All levels |
| Hello interval | One-third of the negotiated hold time (in seconds). If the user-configured value for the hello interval is smaller than the computed value, the user-configured value is used. | detail extensive |
| Hold time | Configured hold time, in seconds. | detail extensive |

Table 179: show ldp interface Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|----------------------|---|-----------------|
| Transport address | Address to which the neighbor wants the local route to establish the LDP session. | extensive |
| Local hello interval | Locally configured hello interval. | extensive |

Sample Output

show ldp interface
extensive

```
user@host> show ldp interface extensive
Interface          Label space ID      Nbr count  Next hello
fe-0/0/3.0         10.255.245.6:0      2          0
  Hello interval: 1, Hold time: 15, Transport address: 10.255.245.6
  Local hello interval: 2, Index: 69
```

show ldp neighbor

| | |
|---------------------------------|--|
| Syntax | show ldp neighbor <brief detail extensive> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> <neighbor-address> |
| Release Information | Command introduced before Junos OS Release 7.4. neighbor-address option added in Junos OS Release 8.5. |
| Description | Display Label Distribution Protocol (LDP) neighbor information. |
| Options | <p>none—Display standard information about LDP neighbors for all routing instances.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>instance <i>instance-name</i>—(Optional) Display information for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>neighbor-address—(Optional) Display information about the specified LDP neighbor.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • clear ldp neighbor on page 777 |
| List of Sample Output | show ldp neighbor extensive on page 788 |
| Output Fields | Table 180 on page 787 describes the output fields for the show ldp neighbor command. Output fields are listed in the approximate order in which they appear. |

Table 180: show ldp neighbor Output Fields

| Field Name | Field Description | Level of Output |
|------------------------|---|-----------------|
| Address | IP address of the neighbor. | All levels |
| Interface | Interface over which the neighbor was discovered. | All levels |
| Label space ID | Label space identifier advertised by the neighbor. | All levels |
| Hold time | Remaining hold time before the neighbor expires, in seconds. | All levels |
| Transport address | Address to which the neighbor wants the local route to establish the LDP session. | detail |
| Configuration sequence | Counter that increments whenever the neighbor changes its configuration. | detail |

Table 180: show ldp neighbor Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------|---|-------------------------|
| Up for | Length of time the LDP neighbor has been in operation. | detail extensive |
| Reference count | Reference count for the LDP neighbor. | extensive |
| Hold time | Displays the neighbor's hold time. The hold time is the proposed hold times for the local and peer routers. | extensive |
| Proposed local/peer | Hold time value proposed by the local router and the peer router. | extensive |

Sample Output

**show ldp neighbor
extensive**

```
user@host> show ldp neighbor extensive
Address          Interface      Label space ID      Hold Time
192.168.37.23    so-1/0/0.0    10.255.245.5:0      44
  Transport address: 10.255.245.5, Configuration sequence: 6
  Up for 00:03:37
  Reference count: 1
  Hold time: 45, Proposed local/peer: 15/45
```

show ldp path

| | |
|---------------------------------|---|
| Syntax | <pre>show ldp path <brief detail extensive> <destination> <instance instance-name> <logical-system (all logical-system-name)></pre> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display Label Distribution Protocol (LDP) label-switched paths (LSPs). |
| Options | <p>none—Display standard information about all LDP LSPs for all routing instances.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>destination—(Optional) Restrict the output to entries that match the specified destination prefix.</p> <p>instance instance-name—(Optional) Display information for the specified routing instance only.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show ldp path extensive on page 790 |
| Output Fields | Table 181 on page 789 describes the output fields for the show ldp path command. Output fields are listed in the approximate order in which they appear. |

Table 181: show ldp path Output Fields

| Field Name | Field Description |
|-------------------------------|--|
| Output Session (label) | Session ID and labels that this system has sent using LDP. These correspond to MPLS packets received. |
| Input Session (label) | Session ID and labels that this system has received using LDP. These correspond to MPLS packets transmitted. |
| route | MPLS route. |
| Attached route | Route corresponding to the LSP. |
| Ingress route | The router acts as the ingress for the LSP. |
| Reference count | Reference count for the LDP neighbor. |
| Transit route | Names of the forwarding equivalence class (FEC) filters on the transit routers. |

Table 181: show ldp path Output Fields (*continued*)

| Field Name | Field Description |
|--------------|-----------------------------------|
| Global label | MPLS label that is used globally. |

Sample Output

**show ldp path
extensive**

```
user@host> show ldp path extensive
Output Session (label)      Input Session (label)
10.255.14.220:0(3)          ( )
  Attached route: 10.255.14.221/32
  Reference count: 3, Global label: 3
10.255.14.220:0(100000)      10.255.14.220:0(3)
  Attached route: 10.255.14.220/32, Ingress route
  Reference count: 2, Transit route, Global label: 100000
10.255.14.220:0(100001)      10.255.14.220:0(100001)
  Attached route: 10.255.14.214/32, Ingress route
  Reference count: 2, Transit route, Global label: 100001
```


show ldp route

| | |
|---------------------------------|---|
| Syntax | <pre>show ldp route <brief detail extensive> <destination> <instance instance-name> <logical-system (all logical-system-name)></pre> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display the entries in the Label Distribution Protocol (LDP) internal topology table. The internal topology table contains routes from inet.0 and inet.3 and is used when binding a label to a forwarding equivalence class (FEC). |
| Options | <p>none—Display standard information about all entries in the LDP internal topology table for all routing instances.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>destination—(Optional) Restrict the output to entries that are longer than the specified destination prefix and prefix length.</p> <p>instance instance-name—(Optional) Display entries for the specified routing instance only.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show ldp route detail on page 793 show ldp route extensive on page 793 |
| Output Fields | Table 182 on page 791 describes the output fields for the show ldp route command. Output fields are listed in the approximate order in which they appear. |

Table 182: show ldp route Output Fields

| Field Name | Field Description |
|--------------------------------|---|
| Destination | Destination prefix. |
| Next-hop intf/lsp/table | Interface that is the next hop to the destination prefix. |
| Next-hop address | IP address of the next hop. |
| Session ID | LDP session ID. |
| Route flags | Information about the route. For example, the Ingress TTL propagate flag indicates that the time-to-live (TTL) value is being propagated with the route. |

Table 182: show ldp route Output Fields (*continued*)

| Field Name | Field Description |
|--------------------------------|--|
| Bound to outgoing label | The route has been bound to LSPs with the label being distributed for that LSP. |
| Topology entry | The topology that the route is bound to. |
| Ingress route status | Status of the ingress route. For example, it could be Active or Inactive . |
| Last modified | The length of time since the ingress route status last changed. |

Sample Output

show ldp route detail

```

user@host> show ldp route 10.255.8.5 detail
Destination      Next-hop intf/lsp      Next-hop address
10.255.8.5/32    f1
  Session ID 10.255.170.84:0--10.255.170.92:0
    fe-0/0/0.0          192.168.100.2
  Session ID 10.255.170.84:0--10.255.8.5:0
    so-0/2/1.0
  Session ID 10.255.170.84:0--10.255.8.5:0
    so-0/2/2.0
  Session ID 10.255.170.84:0--10.255.8.3:0
  Bound to outgoing label 299776, Topology entry: 0x8c38a80
  BFD dest addr  BFD state LSP-ping Next-hop addr  Next-hop intf/lsp
127.0.0.64      up        up        192.168.100.2    fe-0/0/0.0
127.0.1.64      up        up        so-0/2/1.0
127.0.2.64      up        up        so-0/2/2.0
127.0.3.64      up        up        f1
.....

```

show ldp route extensive

```

user@host> show ldp route extensive

user@host> show ldp route extensive
Destination      Next-hop intf/lsp/table  Next-hop address
10.0.0.0/30      ge-1/2/0.18             10.0.0.17
  Session ID 192.168.0.6:0--192.168.0.5:0
  Route flags: None
Destination      Next-hop intf/lsp/table  Next-hop address
10.0.0.4/30      ge-1/2/0.18             10.0.0.17
  Session ID 192.168.0.6:0--192.168.0.5:0
  Route flags: None
Destination      Next-hop intf/lsp/table  Next-hop address
10.0.0.8/30      ge-1/2/1.21             10.0.0.22
  Session ID 192.168.0.6:0--192.168.0.4:0
  Route flags: None
Destination      Next-hop intf/lsp/table  Next-hop address
10.0.0.12/30     ge-1/2/1.21             10.0.0.22
  Session ID 192.168.0.6:0--192.168.0.4:0
  Route flags: None
Destination      Next-hop intf/lsp/table  Next-hop address
10.0.0.16/30     ge-1/2/0.18             10.0.0.17
  Route flags: None
Destination      Next-hop intf/lsp/table  Next-hop address
10.0.0.18/32     ge-1/2/0.18             10.0.0.17
  Route flags: None
Destination      Next-hop intf/lsp/table  Next-hop address
10.0.0.20/30     ge-1/2/1.21             10.0.0.22
  Route flags: None
Destination      Next-hop intf/lsp/table  Next-hop address
10.0.0.21/32     ge-1/2/1.21             10.0.0.22
  Route flags: None
Destination      Next-hop intf/lsp/table  Next-hop address
192.168.0.1/32   ge-1/2/0.18             10.0.0.17
  Session ID 192.168.0.6:0--192.168.0.5:0
  Route flags: None
Destination      Next-hop intf/lsp/table  Next-hop address
192.168.0.2/32   ge-1/2/1.21             10.0.0.22
  Session ID 192.168.0.6:0--192.168.0.4:0
    ge-1/2/0.18          10.0.0.17

```

```
Session ID 192.168.0.6:0--192.168.0.5:0
Route flags: None
Destination      Next-hop intf/lsp/table      Next-hop address
192.168.0.3/32   ge-1/2/1.21                 10.0.0.22
Session ID 192.168.0.6:0--192.168.0.4:0
Route flags: None
Destination      Next-hop intf/lsp/table      Next-hop address
192.168.0.4/32   ge-1/2/1.21                 10.0.0.22
Session ID 192.168.0.6:0--192.168.0.4:0
Bound to outgoing label 299808, Topology entry: 0x92a483c
Ingress route status: Active, Last modified: 00:01:19 ago
Route flags: Ingress TTL propagate, Transit TTL propagate
Destination      Next-hop intf/lsp/table      Next-hop address
192.168.0.5/32   ge-1/2/0.18                 10.0.0.17
Session ID 192.168.0.6:0--192.168.0.5:0
Bound to outgoing label 299792, Topology entry: 0x92a47f8
Ingress route status: Active, Last modified: 00:01:19 ago
Route flags: Ingress TTL propagate, Transit TTL propagate
Destination      Next-hop intf/lsp/table      Next-hop address
192.168.0.6/32   lo0.6                       10.0.0.6
Bound to outgoing label 3, Topology entry: 0x92a4a5c
Ingress route status: Inactive
Route type: Egress route
Route flags: None
```

show ldp session

| | |
|---------------------------------|---|
| Syntax | show ldp session <brief detail extensive> <destination> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display information about Label Distribution Protocol (LDP) sessions. |
| Options | <p>none—Display standard information about all LDP sessions for all routing instances.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>destination—(Optional) Restrict LDP session display to the specified address.</p> <p>instance <i>instance-name</i>—(Optional) Display routing instance information for the specified instance. If <i>instance-name</i> is omitted, information is displayed for the master instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • clear ldp session on page 778 |
| List of Sample Output | show ldp session brief on page 799 show ldp session detail on page 799 show ldp session extensive on page 799 |
| Output Fields | Table 183 on page 795 describes the output fields for the show ldp session command. Output fields are listed in the approximate order in which they appear. |

Table 183: show ldp session Output Fields

| Field Name | Field Description | Level of Output |
|------------|---|------------------|
| Address | Transport address of the session. | any |
| State | State of the session: Nonexistent , Connecting , Initialized , OpenRec , OpenSent , Operational , or Closing . The states correspond to the state diagram specified in Internet Draft LDP Specification draft-ietf-mpls-rfc3036bis-01.txt. | any |
| Connection | TCP connection state: Closed , Opening , or Open . | any |
| Hold time | Time remaining until the session will be closed, in seconds. | any |
| Session ID | LDP identifiers of the peers of this session. | detail extensive |

Table 183: show ldp session Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-------------------------------|---|------------------|
| Next keepalive | Time until next keepalive is sent, in seconds. | detail extensive |
| Active | Whether the local router is playing the active role in the session and during session establishment. | detail extensive |
| Passive | Whether the local router is playing the passive role in the session and during session establishment. | detail extensive |
| Maximum PDU | Maximum protocol data unit (PDU) size (packet size) for the session. | detail extensive |
| Hold time | Time remaining until the session will be closed, in seconds. This value corresponds to the one configured using the keepalive-timeout statement configured at the [edit protocols ldp] hierarchy level. | detail extensive |
| Neighbor count | Number of neighbors that are contributing to the session. | detail extensive |
| Keepalive interval | Keepalive interval, in seconds. | detail extensive |
| Connect retry interval | TCP connection retry interval, in seconds. | detail extensive |
| Local address | Local transport address. | detail extensive |
| Remote address | Remote transport address. | detail extensive |
| Up for | Time that this session has been up. | detail extensive |
| Last down | Time since the session last went down. | detail extensive |
| Reason | Reason the session went down: <ul style="list-style-type: none"> • Aborted graceful restart • Authentication key was changed • Bad type length value (TLV) • Bad protocol data unit (PDU) packets • Command-line interface (CLI) command • Connect time expired • Connection error • Connection reset • Error during initialization • Hold time expired • No adjacency or all adjacencies down • Notification received • Received notification from peer • Unexpected End of File (EOF) • Unknown reason | detail extensive |

Table 183: show ldp session Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-------------------------------------|--|------------------|
| Number of session flaps | Number of times the session changes from up to down. | detail extensive |
| Restarting | LDP is in the process of gracefully restarting. | detail extensive |
| Capabilities advertised | LDP capabilities advertised to a peer. | detail extensive |
| Capabilities received | LDP capabilities received from a peer. | detail extensive |
| Protection | Information about the status of MPLS LDP session protection. | detail extensive |
| restart complete in <i>nnn msec</i> | Amount of time (in milliseconds) remaining until graceful restart is declared complete. | detail extensive |
| Local | <p>Information about graceful restart for the local end of an LDP session. Graceful restart and helper mode are independent.</p> <ul style="list-style-type: none"> • Restart—Status of the graceful restart feature at the local end of the LDP session: enabled or disabled. • Helper mode—Status of the helper mode feature at the local end of the LDP session: enabled or disabled. When this feature is enabled, the local end of the LDP session can help the restarting router with its LDP restart procedures. • Reconnect time—Amount of time to wait from when a restart is initiated until the router can exchange LDP messages with its neighbors. The default is 60000 msec and is not configurable. (Reconnect timeout refers to "FT Reconnect timeout" in draft-ietf-mpls-ldp-restart-06, <i>Internet Draft Graceful Restart Mechanism for LDP</i>.) | detail extensive |
| Remote | <p>Information about graceful restart at the remote end of an LDP session. Graceful restart and helper mode are independent.</p> <ul style="list-style-type: none"> • Restart—Status of the graceful restart feature at the remote end of the LDP session: enabled or disabled. • Helper mode—Status of the helper mode feature at the remote end of the LDP session: enabled or disabled. When this feature is enabled, the remote end of the LDP session can help the restarting router with its LDP restart procedures. • Reconnect time—Amount of time in milliseconds from when a restart is initiated until the remote router can exchange LDP messages with its neighbors. | detail extensive |
| Local maximum recovery time | Amount of time during which the restarting node attempts to recover its lost states with help from its neighbors (in milliseconds). | detail extensive |
| Next-hop addresses received | Next-hop addresses received on the session. | detail extensive |
| Queue depth | Number of messages that are queued for sending to the peers in the group. | extensive |

Table 183: show ldp session Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------|--|------------------|
| Message type | <p>Type of message being sent:</p> <ul style="list-style-type: none"> • Initialization—Session initialization negotiation messages sent by an LSR to an LDP peer when the transport connection is established. • Keepalive—Keepalive timer messages sent by an LSR to an LDP peer to keep the session active when there is no information or PDU exchanged between them. • Notification—Notification messages (such as state of the LDP session) or error information (such as bad PDU length) sent by an LSR to an LDP peer. • Address—Message sent by an LSR to an LDP peer to advertise interface addresses. • Address withdraw—Message sent by an LSR to an LDP peer to withdraw a previously advertised interface address. • Label mapping—Message sent by an LSR to an LDP peer to advertise label mapping for a forwarding equivalence class (FEC). • Label request—Message sent by an LSR to an LDP peer to request a label mapping for an FEC. • Label withdraw—Message sent by an LSR to an LDP peer to withdraw a previously advertised FEC-label mapping. • Label release—Message sent by an LSR to an LDP peer to notify the peer that a specific FEC-label mapping has been released. • Label abort—Message sent by an LSR to an LDP peer to abort a label request message. • Total—Messages sent and received during the lifetime of the session. • Last 5 seconds—Messages sent and received during the current session. | extensive |

Sample Output

```

show ldp session brief  user@host> show ldp session brief
                        Address      State      Connection      Hold time
10.255.72.160          Operational Open          21
10.255.72.164          Operational Open          20
10.255.72.172          Operational Open          21

show ldp session detail user@host> show ldp session detail
Address: 192.168.0.3, State: Operational, Connection: Open, Hold time: 27
Session ID: 192.168.0.2:0--192.168.0.3:0
Next keepalive in 7 seconds
Passive, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
Neighbor types: discovered
Keepalive interval: 10, Connect retry interval: 1
Local address: 192.168.0.2, Remote address: 192.168.0.3
Up for 00:00:02
Capabilities advertised: none
Capabilities received: none
Protection: disabled
Local - Restart: enabled, Helper mode: enabled, Reconnect time: 60000
Remote - Restart: enabled, Helper mode: enabled, Reconnect time: 60000
Local maximum neighbor reconnect time: 120000 msec
Local maximum neighbor recovery time: 240000 msec
Local Label Advertisement mode: Downstream unsolicited
Remote Label Advertisement mode: Downstream unsolicited
Negotiated Label Advertisement mode: Downstream unsolicited
Nonstop routing state: Not in sync
Next-hop addresses received:
  10.0.0.5
  10.0.0.33

show ldp session extensive user@host> show ldp session extensive
Address: 192.168.0.3, State: Operational, Connection: Open, Hold time: 22
Session ID: 192.168.0.2:0--192.168.0.3:0
Next keepalive in 2 seconds
Passive, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
Neighbor types: discovered
Keepalive interval: 10, Connect retry interval: 1
Local address: 192.168.0.2, Remote address: 192.168.0.3
Up for 00:05:37
Capabilities advertised: none
Capabilities received: none
Protection: disabled
Local - Restart: enabled, Helper mode: enabled, Reconnect time: 60000
Remote - Restart: enabled, Helper mode: enabled, Reconnect time: 60000
Local maximum neighbor reconnect time: 120000 msec
Local maximum neighbor recovery time: 240000 msec
Local Label Advertisement mode: Downstream unsolicited
Remote Label Advertisement mode: Downstream unsolicited
Negotiated Label Advertisement mode: Downstream unsolicited
Nonstop routing state: Not in sync
Next-hop addresses received:
  10.0.0.5
  10.0.0.33
Queue depth: 0
Message type              Total
                        Sent      Received
                        Last 5 seconds
                        Sent      Received

```

| | | | | |
|------------------|----|----|---|---|
| Initialization | 1 | 1 | 0 | 0 |
| Keepalive | 33 | 33 | 1 | 1 |
| Notification | 0 | 0 | 0 | 0 |
| Address | 1 | 1 | 0 | 0 |
| Address withdraw | 0 | 0 | 0 | 0 |
| Label mapping | 7 | 5 | 0 | 0 |
| Label request | 0 | 0 | 0 | 0 |
| Label withdraw | 3 | 1 | 0 | 0 |
| Label release | 1 | 3 | 0 | 0 |
| Label abort | 0 | 0 | 0 | 0 |

show ldp statistics

| | |
|---------------------------------|--|
| Syntax | show ldp statistics <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display Label Distribution Protocol (LDP) statistics. |
| Options | <p>none—Display LDP statistics for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Display information for the specified routing instance only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • clear ldp statistics on page 779 |
| List of Sample Output | show ldp statistics on page 804 |
| Output Fields | Table 184 on page 801 lists the output fields for the show ldp statistics command. Output fields are listed in the approximate order in which they appear. |

Table 184: show ldp statistics Output Fields

| Field Name | Field Description |
|-------------------------------|--|
| Total Sent, Received | Total number of each message type sent and received. |
| Last 5 seconds Sent, Received | Number of each message type sent and received in the last 5 seconds. |

Table 184: show ldp statistics Output Fields (*continued*)

| Field Name | Field Description |
|---------------------|--|
| Message type | <p>LDP message types:</p> <ul style="list-style-type: none"> • Hello—Messages that enable LDP nodes to discover one another and to detect the failure of a neighbor or of the link to the neighbor. • Initialization—Messages that indicate an LDP session has started. • Keepalive—Messages that ensure that the keepalive timeout is not exceeded. • Notification—Advisory information and signal error information. • Address—Messages with address information. • Address withdrawal—Messages regarding address withdrawal. • Label mapping—Messages with label mapping information. • Label request—Request for a label mapping from a neighboring router. • Label withdrawal—Withdrawal message sent by the downstream LSR to recall a label that it previously mapped. If an LSR that has received a label mapping subsequently determines that it no longer needs that label, it can send a label release message that frees the label for use. • Label release—Message sent by the downstream LSR to recall a label that it previously mapped. If an LSR that has received a label mapping subsequently determines that it no longer needs that label, it can send a label release message that frees the label for use. • Label abort—Messages about label interruptions. • All UDP—All hello messages sent by LSRs to the well-known UDP port, 646. • All TCP—All LDP session messages. |

Table 184: show ldp statistics Output Fields (*continued*)

| Field Name | Field Description |
|-----------------------|---|
| Event type | <p>LDP events and errors:</p> <ul style="list-style-type: none"> • Sessions opened—Number of LDP sessions that have been opened. • Sessions closed—Number of LDP sessions that have been closed. • Topology changes—Number of changes to the known LDP topology. • No interface—Number of missing interface address messages. When a new LDP session is initialized and before sending label lapping or label request messages, the LSR advertises its interface addresses with one or more address messages. • No session—Number of missing session messages. Session messages are used to establish, maintain, and terminate sessions between LDP peers. • No adjacency—The exchange of hello adjacency messages results in the creation of an adjacency. The LDP identifier, together with the sender's LDP identifier in the PDU header, enables the receiver to match the initialization message with one of its hello adjacencies. If there is no matching hello adjacency, the LSR sends a session the initialization message is rejected. • Unknown version—The LDP protocol version is not supported by the receiver, or it is supported but is not the version negotiated for the session during session establishment. • Malformed PDU—An LDP PDU received on a TCP connection for an LDP session is malformed if the LDP identifier in the PDU header is unknown to the receiver, or if it is known but is not the LDP identifier associated by the receiver with the LDP peer for this LDP session. An LDP PDU is considered to be malformed if the LDP protocol version is not supported by the receiver, or it is supported but is not the version negotiated for the session during session establishment. An LDP PDU is considered malformed if the PDU length field is too small (less than 14) or too large (greater than maximum PDU length). • Malformed message—Malformed LDP messages that are part of the LDP discovery mechanism are handled by silently discarding them. An LDP message is malformed if the message type is unknown. If the message type is less than 0x8000 (high order bit = 0), it is an error signaled by the unknown message type status code. An LDP message is considered to be malformed if the message length is too large, meaning that the message extends beyond the end of the containing LDP PDU. The LDP message is considered to be malformed if the message length is too small, meaning that it is smaller than the smallest possible value component. The LDP message is considered to be malformed if the message is missing one or more mandatory parameters. • Unknown message type—If the message type is less than 0x8000 (high order bit = 0) or greater than or equal to 0x8000 (high order bit = 1) it is considered to be an unknown message. • Inappropriate message—The message is not of the type that the receiver expects to receive. • Malformed TLV—The TLV Length is too large or the receiver cannot decode the TLV value. This can indicate an issue in either the sending or receiving LSR. • Bad TLV value—The TLV Length is too large. • Missing TLV—The TLV is missing one or more mandatory parameters. • PDU too large—The PDF is greater than the maximum PDU length. Section "Initialization Message" in RFC 5036 describes how the maximum PDU length for a session is determined. |
| Total | Total number of each event or error. |
| Last 5 seconds | Number of each event or error in the last 5 seconds. |

Sample Output

show ldp statistics

```

user@host> show ldp statistics
Message type          Total
                      Sent      Received
Hello                 265        263
Initialization        2          2
Keepalive             112       111
Notification          0          0
Address                2          2
Address withdraw       0          0
Label mapping          7          6
Label request          0          0
Label withdraw         2          0
Label release          0          2
Label abort            0          0
All UDP                265       263
All TCP                123       121

                      Last 5 seconds
                      Sent      Received
Hello                 2          2
Initialization        0          0
Keepalive             1          0
Notification          0          0
Address                0          0
Address withdraw       0          0
Label mapping          0          0
Label request          0          0
Label withdraw         0          0
Label release          0          0
Label abort            0          0
All UDP                2          2
All TCP                1          0

Event type              Total      Last 5 seconds
Sessions opened          2          0
Sessions closed          0          0
Topology changes        11          0

No interface             0          0
No session                0          0
No adjacency              0          0
Unknown version           0          0
Malformed PDU             0          0
Malformed message         0          0
Unknown message type       0          0
Inappropriate message      0          0
Malformed TLV              0          0
Bad TLV value              0          0
Missing TLV                0          0
PDU too large              0          0

```

show ldp traffic-statistics

| | |
|---------------------------------|---|
| Syntax | show ldp traffic-statistics <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> <p2mp> |
| Release Information | Command introduced before Junos OS Release 7.4. p2mp option added in Junos OS Release 11.2. |
| Description | Display Label Distribution Protocol (LDP) traffic statistics. |
| Options | <p>none—Display LDP traffic statistics for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Display LDP traffic statistics for the specified routing instance only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>p2mp—(Optional) Display only the data traffic statistics for a point-to-multipoint LSP.</p> |
| Additional Information | To obtain output from this command, you must configure the traffic-statistics statement for the LDP protocol. For more information, see the <i>Junos MPLS Applications Configuration Guide</i> . |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> clear ldp statistics on page 779 |
| List of Sample Output | show ldp traffic-statistics on page 806 show ldp traffic-statistics p2mp on page 806 |
| Output Fields | Table 185 on page 805 lists the output fields for the show ldp traffic-statistics command. Output fields are listed in the approximate order in which they appear. |

Table 185: show ldp traffic-statistics Output Fields

| Field Name | Field Description |
|--------------|---|
| Message type | LDP message types. |
| FEC | Forwarding equivalence class (FEC) for which LDP traffic statistics are collected. For P2MP LSPs, FEC appears as a combination of root address and the LSP ID (root_addr:lsp_id). |
| Type | Type of traffic originating from a router, either Ingress (originating from this router) or Transit (forwarded through this router). |
| Packets | Number of packets passed by the FEC since its LSP came up. |

Table 185: show ldp traffic-statistics Output Fields (*continued*)

| Field Name | Field Description |
|----------------|--|
| Bytes | Number of bytes of data passed by the FEC since its LSP came up. |
| Shared | Whether a label is shared by prefixes: Yes or No . A Yes value indicates that several prefixes are bound to the same label (for example, when several prefixes are advertised with an egress policy). The LDP traffic statistics for this case apply to all the prefixes and should be treated as such. |
| Nexthop | The next hop address for P2MP LSPs. |

Sample Output

show ldp
traffic-statistics

user@host> show ldp traffic-statistics

| FEC | Type | Packets | Bytes | Shared |
|------------------------|--------------|---------|----------|--------|
| 10.35.3.0/30 | Transit | 0 | 0 | Yes |
| | Ingress | 0 | 0 | No |
| 10.35.10.1/32 | Transit | 0 | 0 | Yes |
| | Ingress | 0 | 0 | No |
| 10.255.245.214/32 | Transit | 0 | 0 | No |
| | Ingress | 11 | 752 | No |
| 192.168.37.36/30 | Transit | 0 | 0 | Yes |
| | Ingress | 0 | 0 | No |
| FEC(root_addr:lsp_id) | Nexthop | Packets | Bytes | Shared |
| 10.255.72.160:16777217 | 192.168.8.81 | 152056 | 14597376 | No |
| | 192.168.8.1 | 152056 | 14597376 | No |
| | 192.168.8.65 | 152056 | 14597376 | No |

show ldp
traffic-statistics p2mp

user@host> show ldp traffic-statistics p2mp

| FEC(root_addr:lsp_id) | Nexthop | Packets | Bytes | Shared |
|------------------------|--------------|---------|----------|--------|
| 10.255.72.160:16777217 | 192.168.8.81 | 152056 | 14597376 | No |
| | 192.168.8.1 | 152056 | 14597376 | No |
| | 192.168.8.65 | 152056 | 14597376 | No |

CHAPTER 19

MPLS Operational Mode Commands

Table 186 on page 807 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot Multiprotocol Label Switching (MPLS). Commands are listed in alphabetical order.

Table 186: MPLS Operational Mode Commands

| Task | Command |
|--|--|
| Disconnect and restart dynamic LSPs that originate from this router. | <code>clear mpls lsp</code> |
| Manually trigger a bandwidth allocation adjustment for active LSP paths. | <code>request mpls lsp adjust-autobandwidth</code> |
| Display information about configured cross-connects. | <code>show connections</code> |
| Display peer and traffic engineering link information. | <code>show link-management</code> |
| Display peer link information. | <code>show link-management peer</code> |
| Display peer and traffic engineering link information (routing process). | <code>show link-management routing</code> |
| Display link management statistics. | <code>show link-management statistics</code> |
| Display traffic engineering link information. | <code>show link-management te-link</code> |
| Display MPLS administrative groups. | <code>show mpls admin-groups</code> |
| Display MPLS LSP call admission control (CAC) related information. | <code>show mpls call-admission-control</code> |
| Display CSPF statistics. | <code>show mpls cspf</code> |
| Display DiffServ traffic engineering classes. | <code>show mpls diffserv-te</code> |
| Display the status of interfaces on which MPLS is running. | <code>show mpls interface</code> |

Table 186: MPLS Operational Mode Commands (*continued*)

| Task | Command |
|---|--------------------------------------|
| Display configured LSPs on this router, as well as all ingress, transit, and egress LSPs. | show mpls lsp |
| Display configured named paths that are used in dynamic MPLS. | show mpls path |
| Display Shared Risk Link Group (SRLG) cost and value configuration information. | show mpls srlg |
| Display configured static LSPs on this router, as well as all ingress, transit, and egress static LSPs. | show mpls static-lsp |
| Display entries in the traffic engineering database. | show ted database |
| Display current traffic engineering database links. | show ted link |
| Display protocols contributing to the traffic engineering database. | show ted protocol |



NOTE: For more MPLS-related commands, such as `show route ccc`, `show route protocol`, `show route instance`, and `show route table`, see Protocol-Independent Routing Operational Mode Commands.

For information about how to configure MPLS, see the *Junos MPLS Applications Configuration Guide*.

clear mpls lsp

| | |
|------------------------------------|---|
| Syntax | clear mpls lsp <autobandwidth> <logical-system (all <i>logical-system-name</i>)> <name <i>name</i> > <optimize optimize-aggressive> <path <i>regular-expression</i> > <statistics> |
| Syntax (EX Series Switches) | clear mpls lsp <autobandwidth> <name <i>name</i> > <optimize optimize-aggressive> <path <i>regular-expression</i> > <statistics> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches. |
| Description | Release the routes and states associated with MPLS label-switched paths (LSPs), and start new LSPs. |



CAUTION: This command disconnects existing Resource Reservation Protocol (RSVP) sessions on the ingress routing device. If there is a time lag between the old path being torn down and the new path being set up, this command might impact traffic traveling along the LSPs.

| | |
|----------------|--|
| Options | <p>none—Reset and restart all LSPs that originated from this routing device; that is, all LSPs for which this routing device is the ingress routing device. Depending on the number of LSPs involved, it might take a while to restart all the LSPs.</p> <p>autobandwidth—(Optional) Clear LSP autobandwidth counters.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>name <i>name</i>—(Optional) Reset and restart the specified LSP or group of LSPs. You can include wildcard characters in the interface name, as described in the <i>Junos Network Interfaces Configuration Guide</i>.</p> <p>optimize optimize-aggressive—(Optional) Run nonpreemptive optimization or aggressive optimization computation now.</p> <p>path <i>regular-expression</i>—(Optional) Clear the specific LSP path matching the specified regular expression.</p> <p>statistics—(Optional) Clear LSP statistics. You cannot clear the MPLS LSP statistics using a regular expression (name and path options) on transit routers.</p> |
|----------------|--|

Required Privilege Level clear

Related Documentation

- [show mpls lsp on page 838](#)
- [show rsvp session on page 880](#)

List of Sample Output [clear mpls lsp on page 810](#)

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

Sample Output

`clear mpls lsp` user@host> clear mpls lsp

request mpls lsp adjust-autobandwidth

| | |
|------------------------------------|--|
| Syntax | request mpls lsp adjust-autobandwidth <logical-system (all <i>logical-system-name</i>)> <name <i>lsp-name</i> > |
| Syntax (EX Series Switches) | request mpls lsp adjust-autobandwidth <name <i>lsp-name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches. |
| Description | <p>Manually trigger a bandwidth allocation adjustment for active label-switched paths (LSPs).</p> <p>Without running this command, the bandwidth adjustment is recomputed at a configurable interval. The default interval is 5 minutes. If you do not want to wait for the periodic adjustment (for example, during a software demonstration), this command is useful.</p> <p>During bandwidth allocation adjustment, the LSP stays up to enable the bandwidth to be changed without dropping any traffic. This functionality is often referred to as <i>make-before-break</i>.</p> |
| Options | <p>none—Manually trigger a bandwidth allocation adjustment for all active LSP paths.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>name <i>lsp-name</i>—(Optional) Manually trigger a bandwidth allocation adjustment on the specified LSP only.</p> |
| Additional Information | <p>For this command to work properly, the following conditions must exist:</p> <ul style="list-style-type: none"> Automatic bandwidth allocation must be enabled on the LSP. The parameters for adjustment interval and maximum average bandwidth are not reset after you issue the request mpls lsp adjust-autobandwidth command. The difference between the adjusted bandwidth and the current LSP path bandwidth must be greater than the threshold limit. |
| Required Privilege Level | maintenance |
| Related Documentation | <ul style="list-style-type: none"> auto-bandwidth Configuring Automatic Bandwidth Allocation for LSPs |
| List of Sample Output | request mpls lsp adjust-auto-bandwidth on page 812 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

```
request mpls lsp          user@host> request mpls lsp adjust-auto-bandwidth
adjust-auto-bandwidth
```

show connections

| | |
|------------------------------------|--|
| Syntax | <pre>show connections <brief extensive> <all interface-switch lsp-switch p2mp-receive-switch p2mp-transmit-switch remote-interface-switch> <down up up-down> <history> <labels> <logical-system (all <i>logical-system-name</i>)> <name> <status></pre> |
| Syntax (EX Series Switches) | <pre>show connections <brief extensive> <all interface-switch lsp-switch p2mp-receive-switch p2mp-transmit-switch remote-interface-switch> <down up up-down> <history> <labels> <name> <status></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> |
| Description | Display information about the configured circuit cross-connect (CCC) connections. |
| Options | <p>none—Display the standard level of output for all configured CCC connections.</p> <p>all—(Optional) Display all connections.</p> <p>brief extensive—(Optional) Display the specified level of output. Use history to display information about connection history. Use labels to display labels used for transmit and receive LSPs. Use status to display information about the connection and interface status.</p> <p>interface-switch—(Optional) Display interface switch connections only.</p> <p>lsp-switch—(Optional) Display LSP switch connections only.</p> <p>p2mp-receive-switch—(Optional) Display point-to-multipoint LSP to local interfaces switch connections only.</p> <p>p2mp-transmit-switch—(Optional) Display local interface to point-to-multipoint LSP switch connections only.</p> <p>remote-interface-switch—(Optional) Display remote interface switch connections only.</p> <p>down up up-down—(Optional) Display nonoperational, operational, or both kinds of connections.</p> <p>history—(Optional) Display information about connection history.</p> |

labels—(Optional) Display labels used for transmit and receive.

logical-system (**all** | **logical-system-name**)—(Optional) Perform this operation on all logical systems or on a particular logical system.

name—(Optional) Display information about the specified connection only.

status—(Optional) Display information about the connection and interface status.

Required Privilege Level view

Output Fields [Table 187 on page 814](#) describes the output fields for the **show connections** command. Output fields are listed in the approximate order in which they appear.

Table 187: show connections Output Fields

| Field Name | Field Description |
|--|---|
| CCC and TCC connections [Link Monitoring On Off] | Whether link monitoring is enabled: On or Off . |
| Legend for Status (St) | Connection or circuit status. See the output's legend for an explanation of the status field values. |
| Legend for connection types | Type of connection: <ul style="list-style-type: none"> if-sw—Layer 2 switching cross-connect. rmt-if—Remote interface switch. While graceful restart is in progress, rmt-if will display a state (St) of Restart. lsp-sw—LSP stitching cross-connect. While graceful restart is in progress, lsp-sw will display a state (St) of Restart. |
| Legend for circuit types | Type of circuits: <ul style="list-style-type: none"> intf—Interface circuit. tlsp—Transmit LSP circuit. rlsp—Receive LSP circuit. |
| Connection/Circuit | Name of the configured CCC connection. |
| Type | Type of connection. |
| St | State of the connection. |
| Time last up | Time that the connection or circuit last transitioned to the Up (operational) state. |
| # Up trans | Number of times that the connection or circuit has transitioned to the Up (operational) state. |

Sample Output

show connections

```

user@switch> show connections
CCC and TCC connections [Link Monitoring On]
Legend for status (St)
UN -- uninitialized
NP -- not present
WE -- wrong encapsulation
DS -- disabled
Dn -- down
-> -- only outbound conn is up
<- -- only inbound conn is up
Up -- operational
RmtDn -- remote CCC down
Restart -- restarting

Legend for connection types
if-sw: interface switching
rmt-if: remote interface switching
lsp-sw: LSP switching

Legend for circuit types
intf -- interface
tlsp -- transmit LSP
rlsp -- receive LSP

CCC Graceful restart : Restarting

Connection/Circuit      Type  St    Time last up    # Up trans
IFSW-ed
  so-1/0/2.0             intf  Up      Aug  5 15:39:15      1
  t1-0/1/2.0             intf  Up
SW-db
  so-1/0/3.0             intf  Up
  pro4-ca                tlsp  Dn
  pro4-ac                rlsp  NP

```

show link-management

| | |
|---------------------------------|--|
| Syntax | show link-management |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches. |
| Description | Display Multiprotocol Label Switching (MPLS) peer and traffic engineering link information. |
| Options | This command has no options. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • show link-management peer on page 820 • show link-management routing on page 822 • show link-management statistics on page 825 • show link-management te-link on page 827 |
| List of Sample Output | show link-management on page 819 |
| Output Fields | Table 188 on page 816 describes the output fields for the show link-management command. Output fields are listed in the approximate order in which they appear. |

Table 188: show link-management Output Fields

| Field Name | Field Description |
|-------------------|---|
| Peer Name | Name of the peer. |
| System identifier | Internal identifier for the peer. The range of values is 0 through 64,000. |
| State | State of the peer: Up or Down . |
| Control address | Address to which a control channel is established. |
| CC local ID | Identifier assigned to the control channel by the local peer. The range of values is 1 through 4,294,967,296. |
| CC remote ID | Identifier assigned to the control channel by the remote peer. The range of values is 1 through 4,294,967,296. |
| State | State of the control channel: Up or Down . |
| TxSeqNum | Sequence number of the hello message being sent to the peer. The range of values is 1 through 4,294,967,295. |
| RcvSeqNum | Sequence number of the last hello message received from the peer. The range of values is 0 through 4,294,967,295. |

Table 188: show link-management Output Fields (*continued*)

| Field Name | Field Description |
|----------------------------|---|
| Flags | Code that provides information about the control channel. Currently supports only code value R , which indicates that the control channel is restarting after a failure in the control plane, as when the Link Management Protocol (LMP) process starts or restarts. |
| TE links | Traffic-engineered links that are managed by their peer. |
| TE link name | Name of the traffic-engineered link. |
| State | State of the traffic-engineered link: Up , Down , or Init . |
| Local identifier | Identifier of the local side of the link. |
| Remote identifier | Identifier of the remote side of the link. |
| Local address | Address of the local side of the link. |
| Remote address | Address of the remote side of the link. |
| Encoding | Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include SDH/SONET , Ethernet , Packet , and PDH . |
| Switching | Type of switching that can be performed on the traffic-engineered link. Supported values are PSC-1 and Packet . |
| Minimum bandwidth | Smallest single allocation of bandwidth possible on the traffic-engineered link. This number is equal to the smallest bandwidth interface that is a member of the traffic-engineered link (in bps). |
| Maximum bandwidth | Largest single allocation of bandwidth possible on the traffic-engineered link. This number is equal to the largest bandwidth interface that is a member of the link (in bps). |
| Total bandwidth | Sum of the bandwidth, in bits per second (bps) and megabits per second (Mbps), of all interfaces that are members of the link. |
| Available bandwidth | Sum of the bandwidths of all interfaces that are members of the link and that are not yet allocated (in bps). |
| Name | Name of the interface. |
| State | State of the interface: Up or Down . |
| Local ID | Identifier of the local side of the interface. |
| Remote ID | Identifier of the remote side of the interface. |
| Bandwidth | Bandwidth, in bps or Mbps, of the member interface. |
| Used | Whether the resource is allocated to an LSP: Yes or No . |

Table 188: show link-management Output Fields (*continued*)

| Field Name | Field Description |
|-----------------|-------------------|
| LSP-name | LSP name. |

Sample Output

show
link-management

```
user@host> show link-management
Peer name: PEER-A, System identifier: 11973
State: Up, Control address: 10.255.245.4
  CC local ID CC remote ID State      TxSeqNum  RcvSeqNum  Flags
    24547      24547 Up          1027      1026
TE links:
  pro4-ba

TE link name: pro4-ba, State: Init
Local identifier: 2662, Remote identifier: 0, Encoding: SDH/SONET, Switching:
PSC-1,
Minimum bandwidth: 155.52Mbps, Maximum bandwidth: 155.52Mbps, Total bandwidth:
155.52Mbps,
Available bandwidth: 155.52Mbps
  Name          State Local ID  Remote ID    Bandwidth Used  LSP-name
  so-1/0/2      Up          21271       0            155.52Mbps     No
```

show link-management peer

| | |
|---------------------------------|---|
| Syntax | show link-management peer <name <i>peer-name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches. |
| Description | Display Multiprotocol Label Switching (MPLS) peer link information. |
| Options | none —Display all peer link information. name <i>peer-name</i> —(Optional) Display information for the specified peer only. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • show link-management on page 816 • show link-management routing on page 822 • show link-management statistics on page 825 • show link-management te-link on page 827 |
| List of Sample Output | show link-management peer on page 821 |
| Output Fields | Table 189 on page 820 describes the output fields for the show link-management peer command. Output fields are listed in the approximate order in which they appear. |

Table 189: show link-management peer Output Fields

| Field Name | Field Description |
|---------------------|--|
| Peer Name | Name of the peer. |
| System identifier | Internal identifier for the peer. The range of values is 0 through 64,000. |
| State | State of the peer: Up or Down . |
| Control address | Address to which a control channel is established. |
| Hello interval | How often the routing device sends Link Management Protocol (LMP) hello packets. |
| Hello dead interval | How long LMP waits before declaring the control channel to be dead. This is an interval during which the routing device receives no LMP hello packets from the neighbor on a control that is active or up. |
| CC local ID | Identifier assigned to the control channel by the local peer. The range of values is 1 through 4,294,967,296. |
| CC remote ID | Identifier assigned to the control channel by the remote peer. The range of values is 1 through 4,294,967,296. |

Table 189: show link-management peer Output Fields (*continued*)

| Field Name | Field Description |
|------------------|---|
| State | State of the control channel: Up or Down . |
| TxSeqNum | Sequence number of the hello message being sent to the peer. The range of values is 1 through 4,294,967,295 . |
| RcvSeqNum | Sequence number of the last hello message received from the peer. The range of values is 0 through 4,294,967,295 . |
| Flags | Code that provides information about the control channel. Currently supports only code value R , which indicates that the control channel is restarting after a failure in the control plane, as when the Link Management Protocol (LMP) process starts or restarts. |
| TE links | Traffic-engineered links that are managed by their peer. |

Sample Output

```

show link-management peer
user@host> show link-management peer
Peer name: sonet, System identifier: 41448
State: Up, Control address: 70.70.70.70
Hello interval: 10000, Hello dead interval: 30000
  CC local ID CC remote ID State      TxSeqNum  RcvSeqNum  Flags
    3265          0 ConfSnd          1         0 R
TE links:
to-sonet

```

show link-management routing

| | |
|---------------------------------|--|
| Syntax | show link-management routing <peer <name <i>name</i> > te-link <name <i>name</i> >> <resource <name <i>name</i> >> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches. |
| Description | Display Multiprotocol Label Switching (MPLS) peer or traffic engineering link information from the routing process. |
| Options | <p>none—Display all peer and traffic-engineered link information.</p> <p>peer <name <i>name</i>>—(Optional) Display information for all peers or for the specified peer only.</p> <p>resource <name <i>name</i>>—(Optional) Display information for all resources or for the specified resource only.</p> <p>te-link <name <i>name</i>>—(Optional) Display information for all traffic-engineered forwarding paths or for the specified path only.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • show link-management on page 816 • show link-management peer on page 820 • show link-management statistics on page 825 • show link-management te-link on page 827 |
| List of Sample Output | show link-management routing on page 824 |
| Output Fields | Table 190 on page 822 describes the output fields for the show link-management routing command. Output fields are listed in the approximate order in which they appear. |

Table 190: show link-management routing Output Fields

| Field Name | Field Description |
|-------------------|--|
| Peer Name | Name of the peer. |
| System identifier | Internal identifier for the peer. The range of values is 0 through 64,000. |
| State | State of the peer: Up or Down. |
| Control address | Address to which a control channel is established. |
| Control channel | Interface over which control packets are sent. |

Table 190: show link-management routing Output Fields (*continued*)

| Field Name | Field Description |
|----------------------------|--|
| State | State of the control channel. |
| TE link name | Traffic-engineered link name. |
| State | State of the traffic-engineered link: Up or Down . |
| Local identifier | Identifier of the local side of the link. |
| Remote identifier | Identifier of the remote side of the link. |
| Local address | Address of the local side of the link. |
| Remote address | Address of the remote side of the link. |
| Encoding | Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include SDH/SONET , Ethernet , and Packet . |
| Minimum bandwidth | Smallest single allocation of bandwidth, in bits per second (bps) or megabits per second (Mbps), possible on the traffic-engineered link. This number is equal to the smallest bandwidth interface that is a member of the traffic-engineered link. |
| Maximum bandwidth | Largest single allocation of bandwidth, in bps or Mbps, possible on the traffic-engineered link. This number is equal to the largest bandwidth interface that is a member of the link (in bps). |
| Total bandwidth | Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link. |
| Available bandwidth | Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link and that are not yet allocated. |
| Resource | Forwarding adjacency LSP information. |
| Type | Type of resource. The type is always a forwarding adjacency LSP. |
| State | State of the LSP: Up or Down . |
| System Identifier | Internal identifier for the peer. The range of values is 0 through 64,000 . |
| Total bandwidth | Bandwidth resource, in bps or Mbps, on the TE-link learned from the routing process. |
| Traffic parameters | <ul style="list-style-type: none"> • Encoding—Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include SDH/SONET, Ethernet, and Packet. • Switching—Type of switching that can be performed on the traffic-engineered link: PSC-1 and Packet. • Granularity—Layer 2 data for switching Layer 2 LSPs for this resource. Not supported. This value is always unknown. |

Sample Output

show
link-management
routing

```
user@host> show link-management routing
Peer name: __rpd:fe-0/1/0.0, System identifier: 2147483649
State: Up, Control address: (null)
Control-channel          State
fe-0/1/0.0               Active

Peer name: __rpd:fe-0/1/2.0, System identifier: 2147483650
State: Up, Control address: (null)
Control-channel          State
fe-0/1/2.0               Active

Peer name: __rpd:so-0/2/0.0, System identifier: 2147483651
State: Down, Control address: (null)
Control-channel          State
so-0/2/0.0               State

Peer name: __rpd:so-0/2/1.0, System identifier: 2147483652
State: Down, Control address: (null)
Control-channel          State
so-0/2/1.0               State

...

TE link name: __rpd:fe-0/1/0.0, State: Up
Local identifier: 2147483649, Remote identifier: 0,
Local address: 192.168.37.66, Remote address: 192.168.37.66,
Encoding: Ethernet, Minimum bandwidth: Obps, Maximum bandwidth: 100Mbps,
Total bandwidth: 100Mbps, Available bandwidth: 100Mbps

TE link name: __rpd:fe-0/1/2.0, State: Up
Local identifier: 2147483650, Remote identifier: 0,
Local address: 192.168.37.73, Remote address: 192.168.37.73,
Encoding: Ethernet, Minimum bandwidth: Obps, Maximum bandwidth: 100Mbps,
Total bandwidth: 100Mbps, Available bandwidth: 100Mbps

TE link name: __rpd:so-0/2/0.0, State: Down
Local identifier: 2147483651, Remote identifier: 0,
Local address: 192.168.37.82, Remote address: 192.168.37.95,
Encoding: Ethernet, Minimum bandwidth: Obps, Maximum bandwidth: 155.52Mbps,
Total bandwidth: 155.52Mbps, Available bandwidth: 155.52Mbps

...

Resource: falsp-bd, Type: LSP, State: Dn System identifier: 2147483652,
Total bandwidth: Obps, Traffic parameters: Encoding: Packet, Switching: Packet,
Granularity: Unknown

Resource: falsp-be, Type: LSP, State: Up System identifier: 2147483654,
Total bandwidth: bw[1]=10Mbps, Traffic parameters: Encoding: Packet,
Switching: Packet, Granularity: Unknown
```

show link-management statistics

| | |
|---------------------------------|---|
| Syntax | show link-management statistics <peer <name <i>name</i> >> |
| Release Information | Command introduced in Junos OS Release 8.0. Command introduced in Junos OS Release 9.5 for EX Series switches. |
| Description | Display statistical information for Link Management Protocol (LMP) packets. |
| Options | none —Display information for all peers. peer <name <i>name</i>> —(Optional) Display information for all peers or for the specified peer only. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • show link-management on page 816 • show link-management peer on page 820 • show link-management routing on page 822 • show link-management te-link on page 827 |
| List of Sample Output | show link-management statistics on page 826 |
| Output Fields | Table 191 on page 825 describes the output fields for the show link-management statistics command. Output fields are listed in the approximate order in which they appear. |

Table 191: show link-management statistics Output Fields

| Field Name | Field Description |
|---------------------------|---|
| Received packets | Number of received packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed. |
| Received bad packets | Number of received bad packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed. |
| Small packets | Number of packets that are too small. |
| Wrong protocol version | Number of packets specifying the wrong LMP version. |
| Messages for unknown peer | Number of packets destined for an unknown peer. |
| Messages for bad state | Number of packets indicating a state that does not match the recipient. |
| Stale acknowledgments | Number of configAck and LinkSummaryAck packets received that have a stale message ID. |

Table 191: show link-management statistics Output Fields (*continued*)

| Field Name | Field Description |
|---------------------------------------|--|
| Stale negative acknowledgments | Number of configNack and LinkSummaryNack packets received that have a stale message ID. |
| Sent packets | Number of sent packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed. |
| Retransmitted packets | Number of retransmitted packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed. |
| Dropped packets | Number of packets sent, by message type, that have been dropped by the receiver after the LMP retransmission interval has been exceeded. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed. |

Sample Output

**show
link-management
statistics**

```
user@host> show link-management statistics peer pro4-a
Statistics for peer pro4-a
  Received packets
    Config: 1
    Hello: 2572
    Small packets: 0
    Wrong protocol version: 0
    Messages for unknown peer: 0
    Messages for bad state: 0
    Stale acknowledgments: 0
    Stale negative acknowledgments: 0
  Sent packets
    Config: 2
    ConfigAck: 1
    Hello: 2572
  Retransmitted packets
    Config: 1
```

show link-management te-link

| | |
|---------------------------------|--|
| Syntax | show link-management te-link <brief detail> <name <i>name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches. |
| Description | Display the resources used to set up Multiprotocol Label Switching (MPLS) traffic-engineered forwarding paths. |
| Options | none —Display information for all traffic-engineered links. brief detail —(Optional) Display the specified level of output. name <i>name</i> —(Optional) Display information for the specified traffic-engineered link only. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • show link-management on page 816 • show link-management peer on page 820 • show link-management routing on page 822 • show link-management statistics on page 825 |
| List of Sample Output | show link-management te-link on page 828 |
| Output Fields | Table 192 on page 827 describes the output fields for the show link-management te-link command. Output fields are listed in the approximate order in which they appear. |

Table 192: show link-management te-link Output Fields

| Field Name | Field Description |
|-------------------|---|
| TE link name | Traffic-engineered link name. |
| State | State of the traffic-engineered link: Up or Down . |
| Local identifier | Identifier of the local side of the link. |
| Remote identifier | Identifier of the remote side of the link. |
| Local address | Address of the local side of the link. |
| Remote address | Address of the remote side of the link. |
| Encoding | Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include SDH/SONET , Ethernet , Packet , and PDH . |

Table 192: show link-management te-link Output Fields (*continued*)

| Field Name | Field Description |
|----------------------------|---|
| Switching | Type of switching that can be performed on the traffic-engineered link. Supported values are PSC-1 and Packet . |
| Minimum bandwidth | Smallest single allocation of bandwidth, in bits per second (bps) or megabits per second (Mbps), possible on the traffic-engineered link. This number is equal to the smallest bandwidth interface that is a member of the traffic-engineered link. |
| Maximum bandwidth | Largest single allocation of bandwidth, in bps or Mbps, possible on the traffic-engineered link. This number is equal to the largest bandwidth interface that is a member of the link. |
| Total bandwidth | Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link (in bps). |
| Available Bandwidth | Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link and that are not yet allocated. |
| Name | Name of the interface. |
| State | State of the interface: Up or Down . |
| Local ID | Identifier of the local side of the interface. |
| Remote ID | Identifier of the remote side of the interface. |
| Bandwidth | Bandwidth, in bps or Mbps, of the member interface. |
| Used | Whether the resource is allocated to an LSP: Yes or No . |
| LSP-name | LSP name. |

Sample Output

**show
link-management
te-link**

```

user@host> show link-management te-link
TE link name: FA-bd, State: Up
  Local identifier: 4144, Remote identifier: 0, Local address: 2.2.2.1,
  Remote address: 2.2.2.2, Encoding: Ethernet, Switching: Packet,
  Minimum bandwidth: 0bps, Maximum bandwidth: 0bps, Total bandwidth: 0bps,
  Available bandwidth: 0bps
    Name      State Local ID  Remote ID    Bandwidth Used  LSP-name
    falsp-bd  Dn      43077      0            0bps No
TE link name: FA-be, State: Up
  Local identifier: 4145, Remote identifier: 0, Local address: 1.1.1.1,
  Remote address: 1.1.1.2, Encoding: Ethernet, Switching: Packet,
  Minimum bandwidth: 0bps, Maximum bandwidth: 10Mbps, Total bandwidth: 10Mbps,
  Available bandwidth: 8Mbps
    Name      State Local ID  Remote ID    Bandwidth Used  LSP-name
    falsp-be  Up      43076      0           10Mbps Yes   e2elisp-bf

```

show mpls admin-groups

| | |
|------------------------------------|--|
| Syntax | show mpls admin-groups <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switches) | show mpls admin-groups |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches. |
| Description | Display information about configured Multiprotocol Label Switching (MPLS) administrative groups. |
| Options | none —Display information about the configured MPLS administrative groups. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | view |
| List of Sample Output | show mpls admin-groups on page 829 |
| Output Fields | Table 193 on page 829 describes the output fields for the show mpls admin-groups command. Output fields are listed in the approximate order in which they appear. |

Table 193: show mpls admin-groups Output Fields

| Field Name | Field Description |
|------------|---|
| Group | Name of the administrative group. |
| Bit index | Value assigned to the administrative group. |

Sample Output

```

show mpls
admin-groups

user@host> show mpls admin-groups
Group      Bit index
black      3
blue       2
gold       1
green      0

```

show mpls call-admission-control

| | |
|------------------------------------|---|
| Syntax | show mpls call-admission-control <logical-system (all <i>logical-system-name</i>)> < <i>lsp-name</i> > |
| Syntax (EX Series Switches) | show mpls call-admission-control < <i>lsp-name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches. |
| Description | Display Multiprotocol Label Switching (MPLS) label-switched path (LSP) call admission control (CAC) information. |
| Options | <p>none—Display CAC information for all LSPs.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><i>lsp-name</i>—(Optional) Display CAC information for the specified LSP only.</p> |
| Additional Information | The available bandwidth on an LSP path at a particular class type is the total path bandwidth at that class type minus the total bandwidth reserved by any Layer 2 connection at that class type. |
| Required Privilege Level | view |
| List of Sample Output | show mpls call-admission-control on page 831 |
| Output Fields | Table 194 on page 830 describes the output fields for the show mpls call-admission-control command. Output fields are listed in the approximate order in which they appear. |

Table 194: show mpls call-admission-control Output Fields

| Field Name | Field Description |
|----------------------------|---|
| Available bandwidth | Current available bandwidth on each LSP path. Depending on whether the LSP is an E-LSP or a regular LSP, either per-class bandwidth or a single bandwidth value (corresponding to best-effort bandwidth at ct0) is displayed. The available bandwidth on an LSP path at a particular class type is the total path bandwidth at that class type minus the total bandwidth reserved by some Layer 2 connections at that class type. |
| Layer2 connections | Different Layer 2 connections that had some bandwidth requirement and were admitted into an LSP path. |
| LSP name | LSP pathname. |
| Neighbor address | Neighbor address from which CAC and bandwidth booking are configured for Layer 2 circuits. |
| Circuit | Interface name and circuit information. |

Table 194: show mpls call-admission-control Output Fields (*continued*)

| Field Name | Field Description |
|--------------|--|
| Primary | LSP's primary standby path. |
| Standby | LSP's secondary standby path. |
| VC bandwidth | Bandwidth constraints associated with a Layer 2 circuit route. |

Sample Output

**show mpls
call-admission-control**

```
user@host# show mpls call-admission-control
```

```
LSP name: pro1-be
```

```
*Primary
```

```
Available bandwidth: 0bps
```

```
LSP name: pro1-be-1
```

```
*Primary
```

```
Available bandwidth: 60kbps
```

```
LSP name: pro1-be-gold
```

```
*Primary
```

```
Available bandwidth: <ct0 50kbps> <ct1 20kbps> <ct2 30kbps> <ct3 0bps>
```

```
Layer2 connections:
```

```
Neighbor address: 10.255.245.215, Circuit: so-0/3/0.0(vc 5)
```

```
VC bandwidth: <ct0 50kbps> <ct1 40kbps> <ct2 40kbps>
```

```
LSP name: pro1-be-gold-2
```

```
*Primary
```

```
Available bandwidth: <ct0 0bps> <ct1 40kbps> <ct2 40kbps> <ct3 0bps>
```

```
LSP name: pro1-be-silver
```

```
*Primary prim1
```

```
Available bandwidth: <ct0 10kbps> <ct1 20kbps> <ct2 0bps> <ct3 40kbps>
```

```
Layer2 connections:
```

```
Neighbor address: 10.255.245.215, Circuit: so-0/3/0.1(vc 3)
```

```
VC bandwidth: <ct0 20kbps> <ct1 20kbps>
```

```
Standby sec1
```

```
Available bandwidth: <ct0 10kbps> <ct1 10kbps> <ct2 20kbps> <ct3 0bps>
```

```
Layer2 connections:
```

```
Neighbor address: 10.255.245.215, Circuit: so-0/3/0.1(vc 3)
```

```
VC bandwidth: <ct0 20kbps> <ct1 20kbps>
```

show mpls cspf

| | |
|------------------------------------|---|
| Syntax | show mpls cspf <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switches) | show mpls cspf |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches. |
| Description | Display Multiprotocol Label Switching (MPLS) Constrained Shortest Path First (CSPF) statistics. |
| Options | none —Display MPLS CSFP statistics. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | view |
| List of Sample Output | show mpls cspf on page 833 |
| Output Fields | Table 195 on page 832 describes the output fields for the show mpls cspf command. Output fields are listed in the approximate order in which they appear. |

Table 195: show mpls cspf Output Fields

| Field Name | Field Description |
|---------------------|---|
| Queue length | Number of LSPs queued for automatic path computation. |
| current | Current queue length. |
| maximum | Maximum queue length (high-water mark). |
| dequeued | Number of aborted computation attempts. |
| Paths | Counters for label-switched path computations. |
| total | Sum of the next four fields. |
| successful | Number of path computations that were successfully completed. |
| no route | Number of path computations that failed because the destination is unreachable. |
| Sys Error | Number of path computations that failed because of lack of memory. |

Table 195: show mpls cspf Output Fields (*continued*)

| Field Name | Field Description |
|---------------------|--|
| CSPFs | Total number of CSPF computations. A single path might require multiple CSPF computations. |
| Time | Time, in seconds, required to perform the label-switched path computation. |
| Total | Total amount of time consumed by the CSPF path computation algorithm. |
| CSPFs | Total number of CSPF computations. |
| Avg per CSPF | Average amount of time required for each CSPF computation. |
| % of rpd | Percentage of routing process CPU used in the CSPF computation. |

Sample Output

show mpls cspf

```

user@host> show mpls cspf
CSPF statistics
Queue length  current      maximum      dequeued
              0           0           0
Paths          total      successful    no route    sys error    CSPFs
              0           0           0           0           0
Time (secs)    total      CSPFs      avg per CSPF    % of rpd
              0.000000    0.000000    0.000000      0.0000

```

show mpls diffserv-te

| | |
|------------------------------------|--|
| Syntax | show mpls diffserve-te <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switches) | show mpls diffserve-te |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches. |
| Description | Display Multiprotocol Label Switching (MPLS) label-switched path (LSP) Differentiated Services (DiffServ) class and preemption priority information. |
| Options | none —Display DiffServ classes and priorities used by MPLS LSPs. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | view |
| List of Sample Output | show mpls diffserv-te on page 835 |
| Output Fields | Table 196 on page 834 describes the output fields for the show mpls diffserv-te command. Output fields are listed in the approximate order in which they appear. |

Table 196: show mpls diffserv-te Output Fields

| Field Name | Field Description |
|------------------------|--|
| Bandwidth model | Bandwidth constraint model supported. The maximum allocation model (MAM) for EXP-inferred LSPs (E-LSPs) is currently supported. |
| TE class | DiffServ traffic engineering class. |
| Traffic class | MPLS class type that corresponds to the DiffServ traffic engineering class: <ul style="list-style-type: none"> • ct0—Best effort • ct1—Assured forwarding • ct2—Expedited forwarding • ct3—Network control |
| Priority | MPLS preemption priority for this class type, a value from 0 through 7 . Interior gateway protocols (IGPs) distribute information about the available bandwidth for each traffic engineering class. |

Sample Output

```
show mpls diffserv-te  user@host> show mpls diffserv-te
Bandwidth model: Maximum Allocation Model with support for E-LSPs.
TE class    Traffic class    Priority
te0         ct0              3
te1         ct1              2
```

show mpls interface

| | |
|------------------------------------|--|
| Syntax | show mpls interface <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switches) | show mpls interface |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches. |
| Description | Display information about Multiprotocol Label Switching (MPLS)-enabled interfaces. |
| Options | none —Display information about MPLS-enabled interfaces. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Additional Information | MPLS is enabled on an interface when the interface is configured with both the set protocol mpls interface <i>interface-name</i> and set interface <i>interface-name</i> unit 0 family mpls statements. |
| Required Privilege Level | view |
| List of Sample Output | show mpls interface on page 837 |
| Output Fields | Table 197 on page 836 describes the output fields for the show mpls interface command. Output fields are listed in the approximate order in which they appear. |

Table 197: show mpls interface Output Fields

| Field Name | Field Description |
|--------------------------------------|---|
| Interface | Name of the interface. |
| State | State of the interface: Up or Dn (down). |
| Administrative groups | Administratively assigned colors of the link. |
| Maximum labels | Maximum number of MPLS labels upon which MPLS can operate on a logical interface. This is configured using the maximum-labels statement at the [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family mpls] or the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family mpls] hierarchy levels. |
| Static protection revert time | Time (in seconds) that a static LSP must wait before traffic reverts from the bypass path to the original path. This is configured using the protection-revert-time statement at the [edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i> static] or the [edit protocols mpls interface <i>interface-name</i> static] hierarchy levels. |

Table 197: show mpls interface Output Fields (*continued*)

| Field Name | Field Description |
|---------------------------------------|--|
| Always mark connection protection tlv | Enabled or Disabled: Enabled indicates that the <code>always-mark-connection-protection-tlv</code> statement is configured at the <code>[edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i> static]</code> or the <code>[edit protocols mpls interface <i>interface-name</i> static]</code> hierarchy levels. When this statement is configured, it marks all OAM traffic transiting this interface in preparation for switching the traffic to an alternate path based on the OAM functionality. To switch traffic to the bypass LSP, the <code>switch-away-lsps</code> statement must be configured. |
| Switch away lsps | Enabled or Disabled: Enabled indicates that the <code>switch-away-lsps</code> statement is configured at the <code>[edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i> static]</code> or the <code>[edit protocols mpls interface <i>interface-name</i> static]</code> hierarchy levels. This enables you to switch an LSP away from a network node using a bypass LSP. This feature can be used in maintenance of active networks when a network device needs to be replaced without interrupting traffic passing through the network. The LSPs can be either static or dynamic. |

Sample Output

```

show mpls interface      user@host> show mpls interface

Interface: ge-0/2/1.57
State: Up
Administrative group: <none>
Maximum labels: 5
Static protection revert time: 5 seconds
Always mark connection protection tlv: Disabled
Switch away lsps : Disabled

```

show mpls lsp

| | |
|------------------------------------|--|
| Syntax | <pre>show mpls lsp <brief detail extensive terse> <bidirectional unidirectional> <bypass> <count-active-routes> <defaults> <descriptions> <down up> <logical-system (all <i>logical-system-name</i>)> <lsp-type> <name <i>name</i>> <p2mp> <statistics> <transit></pre> |
| Syntax (EX Series Switches) | <pre>show mpls lsp <brief detail extensive terse> <bidirectional unidirectional> <bypass> <descriptions> <down up> <lsp-type> <name <i>name</i>> <p2mp> <statistics> <transit></pre> |
| Release Information | Command introduced before Junos OS Release 7.4. defaults option added in Junos OS Release 8.5. Command introduced in Junos OS Release 9.5 for EX Series switches. |
| Description | Display information about configured and active dynamic Multiprotocol Label Switching (MPLS) label-switched paths (LSPs). |
| Options | <p>none—Display standard information about all configured and active dynamic MPLS LSPs.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output. The extensive option displays the same information as the detail option, but covers the most recent 50 events.</p> <p>bidirectional unidirectional—(Optional) Display bidirectional or unidirectional LSP information, respectively.</p> <p>bypass—(Optional) Display LSPs used for protecting other LSPs.</p> <p>count-active-routes—(Optional) Display active routes for LSPs.</p> <p>defaults—(Optional) Display the MPLS LSP default settings.</p> <p>descriptions—(Optional) Display the MPLS label-switched path (LSP) descriptions. To view this information, you must configure the description statement at the [edit</p> |

protocol mpls lsp hierarchy level. Only LSPs with a description are displayed. This command is only valid for the ingress routing device, because the description is not propagated in RSVP messages.

down | up—(Optional) Display only LSPs that are inactive or active, respectively.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

lsp-type—(Optional) Display information about a particular LSP type:

- **bypass**—Sessions for bypass LSPs.
- **egress**—Sessions that terminate on this routing device.
- **ingress**—Sessions that originate from this routing device.
- **transit**—Sessions that pass through this routing device.

name *name*—(Optional) Display information about the specified LSP or group of LSPs.

p2mp—(Optional) Display information about point-to-multipoint LSPs.

statistics—(Optional) (Ingress and transit routers only) Display accounting information about LSPs. Statistics are not available for LSPs on the egress routing device, because the penultimate routing device in the LSP sets the label to 0. Also, as the packet arrives at the egress routing device, the hardware removes its MPLS header and the packet reverts to being an IPv4 packet. Therefore, it is counted as an IPv4 packet, not an MPLS packet.



NOTE: If a bypass LSP is configured for the primary static LSP, display cumulative statistics of packets traversing through the protected LSP and bypass LSP when traffic is re-optimized when the protected LSP link is restored.

When used with the **bypass** option (**show mpls lsp bypass statistics**), display statistics for the traffic that flows only through the bypass LSP.

transit—(Optional) Display LSPs transiting this routing device.

Required Privilege Level

view

Related Documentation

- [clear mpls lsp on page 809](#)

List of Sample Output

[show mpls lsp defaults on page 846](#)
[show mpls lsp descriptions on page 846](#)
[show mpls lsp detail on page 846](#)
[show mpls lsp extensive on page 846](#)
[show mpls lsp ingress extensive on page 847](#)

[show mpls lsp p2mp on page 848](#)

[show mpls lsp p2mp detail on page 848](#)

[show mpls lsp detail count-active-routes on page 849](#)

[show mpls lsp statistics extensive on page 849](#)

Output Fields [Table 198 on page 840](#) describes the output fields for the **show mpls lsp** command. Output fields are listed in the approximate order in which they appear.

Table 198: show mpls lsp Output Fields

| Field Name | Field Description | Level of Output |
|--------------------------|---|-------------------------|
| Ingress LSP | Information about LSPs on the ingress routing device. Each session has one line of output. | All levels |
| Egress LSP | Information about the LSPs on the egress routing device. MPLS learns this information by querying RSVP, which holds all the transit and egress session information. Each session has one line of output. | All levels |
| Transit LSP | Number of LSPs on the transit routing devices and the state of these paths. MPLS learns this information by querying RSVP, which holds all the transit and egress session information. | All levels |
| P2MP name | Name of the point-to-multipoint LSP. Dynamically generated P2MP LSPs used for VPLS flooding use dynamically generated P2MP LSP names. The name uses the format <i>identifier:vpls:router-id:routing-instance-name</i> . The <i>identifier</i> is automatically generated by Junos OS. | All levels |
| P2MP branch count | Number of destination LSPs the point-to-multipoint LSP is transmitting to. | All levels |
| P | An asterisk (*) under this heading indicates that the LSP is a primary path. | All levels |
| address | (detail and extensive) Destination (egress routing device) of the LSP. | detail extensive |
| To | Destination (egress routing device) of the session. | brief |
| From | Source (ingress routing device) of the session. | brief detail |
| State | State of the LSP handled by this RSVP session: Up , Dn (down), or Restart . | brief detail |
| Active Route | Number of active routes (prefixes) installed in the forwarding table. For ingress LSPs, the forwarding table is the primary IPv4 table (inet.0). For transit and egress RSVP sessions, the forwarding table is the primary MPLS table (mpls.0). | detail extensive |
| Rt | Number of active routes (prefixes) installed in the routing table. For ingress RSVP sessions, the routing table is the primary IPv4 table (inet.0). For transit and egress RSVP sessions, the routing table is the primary MPLS table (mpls.0). | brief |
| P | Path. An asterisk (*) underneath this column indicates that the LSP is a primary path. | brief |
| ActivePath | (Ingress LSP) Name of the active path: Primary or Secondary . | detail extensive |
| LSPname | Name of the LSP. | brief detail |

Table 198: show mpls lsp Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------------------|---|-------------------------|
| Statistics | Displays the number of packets and the number of bytes transmitted over the LSP. These counters are reset to zero whenever the LSP path is optimized (for example, during an automatic bandwidth allocation). | extensive |
| Aggregate statistics | Displays the number of packets and the number of bytes transmitted over the LSP. These counters continue to iterate even if the LSP path is optimized. You can reset these counters to zero using the clear mpls lsp statistics command. | extensive |
| Packets | Displays the number of packets transmitted over the LSP. | brief extensive |
| Bytes | Displays the number of bytes transmitted over the LSP. | brief extensive |
| DiffServInfo | Type of LSP: multiclass LSP (multiclass diffServ-TE LSP) or Differentiated-Services-aware traffic engineering LSP (diffServ-TE LSP). | detail |
| LSPtype | Type of LSP: static Static configured or dynamic Dynamic configured . Also indicates if the LSP is a Penultimate hop popping LSP or an Ultimate hop popping LSP. | detail extensive |
| Bypass | (Bypass LSP) Destination address (egress routing device) for the bypass LSP. | All levels |
| LSPpath | Indicates whether the RSVP session is for the primary or secondary LSP path. LSPpath can be either primary or secondary and can be displayed on the ingress, egress, and transit routing devices. | detail |
| Bidir | (GMPLS) The LSP allows data to travel in both directions between GMPLS devices. | All levels |
| Bidirectional | (GMPLS) The LSP allows data to travel both ways between GMPLS devices. | All levels |
| FastReroute desired | Fast reroute has been requested by the ingress routing device. | detail |
| Link protection desired | Link protection has been requested by the ingress routing device. | detail |
| LoadBalance | (Ingress LSP) CSPF load-balancing rule that was configured to select the LSP's path among equal-cost paths: Most-fill , Least-fill , or Random . | detail extensive |
| Signal type | Signal type for GMPLS LSPs. The signal type determines the peak data rate for the LSP: DS0 , DS3 , STS-1 , STM-1 , or STM-4 . | All levels |
| Encoding type | LSP encoding type: Packet , Ethernet , PDH , SDH/SONET , Lambda , or Fiber . | All levels |
| Switching type | Type of switching on the links needed for the LSP: Fiber , Lambda , Packet , TDM , or PSC-1 . | All levels |
| GPID | Generalized Payload Identifier (identifier of the payload carried by an LSP): HDLC , Ethernet , IPv4 , PPP , or Unknown . | All levels |

Table 198: show mpls lsp Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---|--|-------------------------|
| Protection | Configured protection capability desired for the LSP: Extra, Enhanced, none, One plus one, One to one , or Shared . | All levels |
| Upstream label in | (Bidirectional LSPs) Incoming label for reverse direction traffic for this LSP. | All levels |
| Upstream label out | (Bidirectional LSPs) Outgoing label for reverse direction traffic for this LSP. | All levels |
| Suggested label received | (Bidirectional LSPs) Label the upstream node suggests to use in the Resv message that is sent. | All levels |
| Suggested label sent | (Bidirectional LSPs) Label the downstream node suggests to use in the Resv message that is returned. | All levels |
| Autobandwidth | (Ingress LSP) The LSP is performing autobandwidth allocation. | detail extensive |
| MinBW | (Ingress LSP) Configured minimum value of the LSP, in bps. | detail extensive |
| MaxBW | (Ingress LSP) Configured maximum value of the LSP, in bps. | detail extensive |
| AdjustTimer | (Ingress LSP) Configured value of the bandwidth adjustment timer, indicating the total amount of time allowed before bandwidth adjustment will take place, in seconds. | detail extensive |
| MaxAvgBW util | (Ingress LSP) Current value of the actual maximum average bandwidth utilization, in bps. | detail extensive |
| Overflow limit | (Ingress LSP) Configured value of the threshold overflow limit. | detail extensive |
| Overflow sample count | (Ingress LSP) Current value for the overflow sample count. | detail extensive |
| Bandwidth Adjustment in <i>nnn</i> second(s) | (Ingress LSP) Current value of the bandwidth adjustment timer, indicating the amount of time remaining until the bandwidth adjustment will take place, in seconds. | detail extensive |
| Underflow limit | (Ingress LSP) Configured value of the threshold underflow limit. | detail extensive |
| Underflow sample count | (Ingress LSP) Current value for the underflow sample count. | detail extensive |
| Underflow Max AvgBW | (Ingress LSP) The highest sample bandwidth among the underflow samples recorded currently. This is the signaling bandwidth if an adjustment occurs because of an underflow. | detail extensive |
| Active path indicator | (Ingress LSP) A value of * indicates that the path is active. The absence of * indicates that the path is not active. In the following example, "long" is the active path. *Primary long Standby short | detail extensive |

Table 198: show mpls lsp Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---|---|-------------------------|
| Primary | (Ingress LSP) Name of the primary path. | detail extensive |
| Secondary | (Ingress LSP) Name of the secondary path. | detail extensive |
| Standby | (Ingress LSP) Name of the path in standby mode. | detail extensive |
| State | (Ingress LSP) State of the path: Up or Dn (down). | detail extensive |
| COS | (Ingress LSP) Class-of-service value. | detail extensive |
| Bandwidth per class | (Ingress LSP) Active bandwidth for the LSP path for each MPLS class type, in bps. | detail extensive |
| Priorities | (Ingress LSP) Configured value of the setup priority and the hold priority respectively (the setup priority is displayed first), where 0 is the highest priority and 7 is the lowest priority. If you have not explicitly configured these values, the default values are displayed (7 for the setup priority and 0 for the hold priority). | detail extensive |
| OptimizeTimer | (Ingress LSP) Configured value of the optimize timer, indicating the total amount of time allowed before path reoptimization, in seconds. | detail extensive |
| SmartOptimizeTimer | (Ingress LSP) Configured value of the smart optimize timer, indicating the total amount of time allowed before path reoptimization, in seconds. | detail extensive |
| Reoptimization in xxx seconds | (Ingress LSP) Current value of the optimize timer, indicating the amount of time remaining until the path will be reoptimized, in seconds. | detail extensive |
| Computed ERO (S [L] denotes strict [loose] hops) | (Ingress LSP) Computed explicit route. A series of hops, each with an address followed by a hop indicator. The value of the hop indicator can be strict (S) or loose (L). | detail extensive |
| CSPF metric | (Ingress LSP) Constrained Shortest Path First metric for this path. | detail extensive |

Table 198: show mpls lsp Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------|---|-------------------------------|
| Received RRO | <p>(Ingress LSP) Received record route. A series of hops, each with an address followed by a flag. (In most cases, the received record route is the same as the computed explicit route. If Received RRO is different from Computed ERO, there is a topology change in the network, and the route is taking a detour.) The following flags identify the protection capability and status of the downstream node:</p> <ul style="list-style-type: none"> • 0x01—Local protection available. The link downstream from this node is protected by a local repair mechanism. This flag can be set only if the Local protection flag was set in the SESSION_ATTRIBUTE object of the corresponding Path message. • 0x02—Local protection in use. A local repair mechanism is in use to maintain this tunnel (usually because of an outage of the link it was routed over previously). • 0x03—Combination of 0x01 and 0x02. • 0x04—Bandwidth protection. The downstream routing device has a backup path providing the same bandwidth guarantee as the protected LSP for the protected section. • 0x08—Node protection. The downstream routing device has a backup path providing protection against link and node failure on the corresponding path section. If the downstream routing device can set up only a link-protection backup path, the Local protection available bit is set but the Node protection bit is cleared. • 0x09—Detour is established. Combination of 0x01 and 0x08. • 0x10—Preemption pending. The preempting node sets this flag if a pending preemption is in progress for the traffic engine LSP. This flag indicates to the ingress legacy edge router (LER) of this LSP that it should be rerouted. • 0xb—Detour is in use. Combination of 0x01, 0x02, and 0x08. | detail extensive |
| Index number | (Ingress LSP) Log entry number of each LSP path event. The numbers are in chronological descending order, with a maximum of 50 index numbers displayed. | extensive |
| Date | (Ingress LSP) Date of the LSP event. | extensive |
| Time | (Ingress LSP) Time of the LSP event. | extensive |
| Event | (Ingress LSP) Description of the LSP event. | extensive |
| Created | (Ingress LSP) Date and time the LSP was created. | extensive |
| Resv style | (Bypass) RSVP reservation style. This field consists of two parts. The first is the number of active reservations. The second is the reservation style, which can be FF (fixed filter), SE (shared explicit), or WF (wildcard filter). | brief detail extensive |
| Labelin | Incoming label for this LSP. | brief detail |
| Labelout | Outgoing label for this LSP. | brief detail |
| LSPname | Name of the LSP. | brief detail |

Table 198: show mpls lsp Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------------------|---|-----------------|
| Time left | Number of seconds remaining in the lifetime of the reservation. | detail |
| Since | Date and time when the RSVP session was initiated. | detail |
| Tspec | Sender's traffic specification, which describes the sender's traffic parameters. | detail |
| Port number | Protocol ID and sender or receiver port used in this RSVP session. | detail |
| PATH rcvfrom | Address of the previous-hop (upstream) routing device or client, interface the neighbor used to reach this router, and number of packets received from the upstream neighbor. | detail |
| PATH sentto | Address of the next-hop (downstream) routing device or client, interface used to reach this neighbor, and number of packets sent to the downstream routing device. | detail |
| RESV rcvfrom | Address of the previous-hop (upstream) routing device or client, interface the neighbor used to reach this routing device, and number of packets received from the upstream neighbor. The output in this field, which is consistent with that in the PATH rcvfrom field, indicates that the RSVP negotiation is complete. | detail |
| Record route | Recorded route for the session, taken from the record route object. | detail |
| Soft preempt | Number of soft preemptions that occurred on a path and when the last soft preemption occurred. Only successful soft preemptions are counted (those that actually resulted in a new path being used). | detail |
| Soft preemption pending | Path is in the process of being soft preempted. This display is removed once the ingress router has calculated a new path. | detail |
| MPLS-TE LSP Defaults | Default settings for MPLS traffic engineered LSPs: <ul style="list-style-type: none"> • LSP Holding Priority—Determines the degree to which an LSP holds on to its session reservation after the LSP has been set up successfully. • LSP Setup Priority—Determines whether a new LSP that preempts an existing LSP can be established. • Hop Limit—Specifies the maximum number of routers the LSP can traverse (including the ingress and egress). • Bandwidth—Specifies the bandwidth in bits per second for the LSP. • LSP Retry Timer—Length of time in seconds that the ingress router waits between attempts to establish the primary path. | defaults |

The XML tag name of the **bandwidth** tag under the **auto-bandwidth** tag has been updated to **maximum-average-bandwidth**. You can see the new tag when you issue the **show mpls lsp extensive** command with the **| display xml** pipe option. If you have any scripts that use the **bandwidth** tag, ensure that they are updated to **maximum-average-bandwidth**.

Sample Output

show mpls lsp defaults

```
user@host> show mpls lsp defaults
MPLS-TE LSP Defaults
  LSP Holding Priority      0
  LSP Setup Priority       7
  Hop Limit                255
  Bandwidth                0
  LSP Retry Timer          30 seconds
```

show mpls lsp descriptions

```
user@host> show mpls lsp descriptions
Ingress LSP: 3 sessions
To          LSP name          Description
10.0.0.195  to-sanjose                 to-sanjose-desc
10.0.0.195  to-sanjose-other-desc      other-desc
Total 2 displayed, Up 2, Down 0
```

show mpls lsp detail

```
user@host> show mpls lsp detail
Ingress LSP: 1 sessions

192.168.0.4
  From: 192.168.0.5, State: Up, ActiveRoute: 0, LSPname: E-D
  ActivePath: (primary)
  LSPtype: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
  10.0.0.18 S 10.0.0.22 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
  20=Node-ID):
      10.0.0.18 10.0.0.22
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

192.168.0.5
  From: 192.168.0.4, LSPstate: Up, ActiveRoute: 0
  LSPname: E-D, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left: 157, Since: Wed Jul 18 17:55:12 2012
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 46128 protocol 0
  PATH rcvfrom: 10.0.0.18 (lt-1/2/0.17) 3 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.0.0.22 10.0.0.18 <self>
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```


**show mpls lsp
extensive**

```
user@host> show mpls lsp extensive
Ingress LSP: 1 sessions
```

```
192.168.0.4
  From: 192.168.0.5, State: Up, ActiveRoute: 0, LSPname: E-D
  ActivePath: (primary)
  LSPtype: Static Configured, Ultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
10.0.0.18 S 10.0.0.22 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
    10.0.0.18 10.0.0.22
  11 Sep 20 15:54:35.032 Make-before-break: Switched to new instance
  10 Sep 20 15:54:34.029 Record Route: 10.0.0.18 10.0.0.22
  9 Sep 20 15:54:34.029 Up
  8 Sep 20 15:54:20.271 Originate make-before-break call
  7 Sep 20 15:54:20.271 CSPF: computation result accepted 10.0.0.18 10.0.0.22

  6 Sep 20 15:52:10.247 Selected as active path
  5 Sep 20 15:52:10.246 Record Route: 10.0.0.18 10.0.0.22
  4 Sep 20 15:52:10.243 Up
  3 Sep 20 15:52:09.745 Originate Call
  2 Sep 20 15:52:09.745 CSPF: computation result accepted 10.0.0.18 10.0.0.22

  1 Sep 20 15:51:39.903 CSPF failed: no route toward 192.168.0.4
  Created: Thu Sep 20 15:51:08 2012
Total 1 displayed, Up 1, Down 0
```

```
Egress LSP: 1 sessions
```

```
192.168.0.5
  From: 192.168.0.4, LSPstate: Up, ActiveRoute: 0
  LSPname: E-D, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left: 148, Since: Thu Sep 20 15:52:10 2012
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 49601 protocol 0
  PATH rcvfrom: 10.0.0.18 (lt-1/2/0.17) 27 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.0.0.22 10.0.0.18 <self>
Total 1 displayed, Up 1, Down 0
```

```
Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

**show mpls lsp ingress
extensive**

```
user@host> show mpls lsp ingress extensive
Ingress LSP: 1 sessions
```

```
50.0.0.1
  From: 10.0.0.1, State: Up, ActiveRoute: 0, LSPname: test
  ActivePath: (primary)
  LSPtype: Static Configured
```

```

LoadBalance: Random
Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary                               State: Up
  Priorities: 7 0
  OptimizeTimer: 300
  SmartOptimizeTimer: 180
  Reoptimization in 240 second(s).
  Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
1.1.1.2 S 4.4.4.1 S 5.5.5.2 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
    1.1.1.2 4.4.4.1 5.5.5.2
17 Aug 3 13:17:33.601 CSPF: computation result ignored, new path less avail
bw[3 times]
16 Aug 3 13:02:51.283 CSPF: computation result ignored, new path no benefit[2
times]
15 Aug 3 12:54:36.678 Selected as active path
14 Aug 3 12:54:36.676 Record Route: 1.1.1.2 4.4.4.1 5.5.5.2
13 Aug 3 12:54:36.676 Up
12 Aug 3 12:54:33.924 Deselected as active
11 Aug 3 12:54:33.924 Originate Call
10 Aug 3 12:54:33.923 Clear Call
9 Aug 3 12:54:33.923 CSPF: computation result accepted 1.1.1.2 4.4.4.1
5.5.5.2
8 Aug 3 12:54:33.922 2.2.2.2: No Route toward dest
7 Aug 3 12:54:28.177 CSPF: computation result ignored, new path no benefit[4
times]
6 Aug 3 12:35:03.830 Selected as active path
5 Aug 3 12:35:03.828 Record Route: 2.2.2.2 3.3.3.2
4 Aug 3 12:35:03.827 Up
3 Aug 3 12:35:03.814 Originate Call
2 Aug 3 12:35:03.814 CSPF: computation result accepted 2.2.2.2 3.3.3.2
1 Aug 3 12:34:34.921 CSPF failed: no route toward 50.0.0.1
Created: Tue Aug 3 12:34:35 2010
Total 1 displayed, Up 1, Down 0

```

show mpls lsp p2mp

```

user@host> show mpls lsp p2mp
Ingress LSP: 2 sessions
P2MP name: p2mp-lsp1, P2MP branch count: 1
To          From          State Rt P ActivePath      LSPname
10.255.245.51 10.255.245.50 Up    0 * path1         p2mp-branch-1
P2MP name: p2mp-lsp2, P2MP branch count: 1
To          From          State Rt P ActivePath      LSPname
10.255.245.51 10.255.245.50 Up    0 * path1         p2mp-st-br1
Total 2 displayed, Up 2, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

show mpls lsp p2mp detail

```

user@host> show mpls lsp p2mp detail
Ingress LSP: 2 sessions
P2MP name: p2mp-lsp1, P2MP branch count: 1

10.255.245.51
  From: 10.255.245.50, State: Up, ActiveRoute: 0, LSPname: p2mp-branch-1
  ActivePath: path1 (primary)
  P2MP name: p2mp-lsp1

```

```

LoadBalance: Random
Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary path1 State: Up
  Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 25)
192.168.208.17 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):
    192.168.208.17
P2MP name: p2mp-lsp2, P2MP branch count: 1

10.255.245.51
  From: 10.255.245.50, State: Up, ActiveRoute: 0, LSPname: p2mp-st-br1
  ActivePath: path1 (primary)
  P2MP name: p2mp-lsp2
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary path1 State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 25)
192.168.208.17 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):
      192.168.208.17
Total 2 displayed, Up 2, Down 0

```

show mpls lsp detail count-active-routes

```

user@host> show mpls lsp detail count-active-routes
Ingress LSP: 1 sessions

213.119.192.2
  From: 156.154.162.128, State: Up, ActiveRoute: 1, LSPname: to-lahore
  ActivePath: (primary)
  LSPtype: Static Configured
  LoadBalance: Random
  Autobandwidth
  MinBW: 5Mbps MaxBW: 250Mbps
  AdjustTimer: 300 secs
  Max AvgBW util: 0bps, Bandwidth Adjustment in 102 second(s).
  Overflow limit: 0, Overflow sample count: 0
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary State: Up
    Priorities: 7 0
    Bandwidth: 5Mbps
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 4)
10.252.0.177 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
    10.252.0.177
Total 1 displayed, Up 1, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

show mpls lsp statistics extensive

```

user@host> show mpls lsp statistics extensive
Ingress LSP: 1 sessions

192.168.0.4
  From: 192.168.0.5, State: Up, ActiveRoute: 0, LSPname: E-D

```

```
Statistics: Packets 302, Bytes 28992
Aggregate statistics: Packets 302, Bytes 28992
ActivePath: (primary)
LSPtype: Static Configured, Penultimate hop popping
LoadBalance: Random
Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary                               State: Up
  Priorities: 7 0
  SmartOptimizeTimer: 180
  Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
10.0.0.18 S 10.0.0.22 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
    10.0.0.18 10.0.0.22
    6 Oct  3 11:18:28.281 Selected as active path
    5 Oct  3 11:18:28.281 Record Route:  10.0.0.18 10.0.0.22
    4 Oct  3 11:18:28.280 Up
    3 Oct  3 11:18:27.995 Originate Call
    2 Oct  3 11:18:27.995 CSPF: computation result accepted  10.0.0.18 10.0.0.22

    1 Oct  3 11:17:59.118 CSPF failed: no route toward 192.168.0.4[2 times]
  Created: Wed Oct  3 11:17:01 2012
Total 1 displayed, Up 1, Down 0
```

show mpls path

| | |
|------------------------------------|--|
| Syntax | show mpls path <logical-system (all <i>logical-system-name</i>)> < <i>path-name</i> > |
| Syntax (EX Series Switches) | show mpls path < <i>path-name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches. |
| Description | Display dynamic Multiprotocol Label Switching (MPLS) label-switched paths (LSPs). |
| Options | <p>none—Display standard information about all MPLS LSPs.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><i>path-name</i>—(Optional) Display information about the specified LSP only.</p> |
| Required Privilege Level | view |
| List of Sample Output | show mpls path on page 851 |
| Output Fields | Table 199 on page 851 describes the output fields for the show mpls path command. Output fields are listed in the approximate order in which they appear. |

Table 199: show mpls path Output Fields

| Field Name | Field Description |
|----------------------|---|
| Path name | Information about ingress LSPs. Each path has one line of output. |
| Address | Addresses of the routing devices that form the LSP. |
| Strict/loose address | Whether the address is configured as a strict or loose address. |

Sample Output

show mpls path

```

user@host> show mpls path
Path name      Address          Strict/loose address
p1             123.456.55.6    Strict
               123.456.1.6     Loose
p2             191.456.1.4     Strict

```

show mpls srlg

| | |
|---------------------------------|--|
| Syntax | <code>show mpls srlg</code> <code><logical-systems (all <i>logical-system-name</i>)></code> |
| Release Information | Command introduced before Junos OS Release 11.4. |
| Description | Display Shared Risk Link Group (SRLG) cost and value configuration information. |
| Options | <code>logical-system (all <i>logical-system-name</i>)</code> —(Optional) View SRLG configuration information for all logical systems or a particular logical system. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">Example: Configuring SRLG |
| Output Fields | Table 200 on page 852 lists the output fields for the show mpls srlg command. Output fields are listed in the approximate order in which they appear. |

Table 200: show mpls srlg Output Fields

| Field Name | Field Description |
|------------|--|
| SRLG | Name of the SRLG. |
| Value | A group ID for the SRLG ranging from 1 through 4294967295. |
| Cost | A cost for the Shared Risk Link Group (SRLG) ranging from 1 through 65535. |

Sample Output

```
user@host> show mpls srlg
```

| SRLG | Value | Cost |
|--------|-------|------|
| srlg-a | 101 | 10 |

show mpls static-lsp

Syntax show mpls static-lsp
 <brief | detail | extensive | terse>
 <bypass>
 <descriptions>
 <down | up>
 <ingress>
 <logical-system (all | *logical-system-name*)>
 <lsp-type>
 <name *name*>
 <statistics>
 <transit>

Release Information Command introduced in Junos OS Release 10.1.

Description Display information about configured and active static Multiprotocol Label Switching (MPLS) label-switched paths (LSPs).

Options **none**—Display standard information about all configured and active static MPLS LSPs.

brief | detail | extensive | terse—(Optional) Display the specified level of output. The **extensive** option displays the same information as the **detail** option, but covers the most recent 50 events.

bypass—(Optional) Display LSPs used for protecting other static LSPs.

descriptions—(Optional) Display the MPLS static LSP descriptions. To view this information, you must configure the description statement at the **[edit protocols mpls static-label-switched-path *path-name* bypass]**, **[edit protocols mpls static-label-switched-path *path-name* ingress]**, or **[edit protocols mpls static-label-switched-path *path-name* transit *incoming-label*]** hierarchy levels. Only static LSPs with a description are displayed.

down | up—(Optional) Display only static LSPs that are inactive or active, respectively.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

lsp-type—(Optional) Display information about a particular LSP type:

- **bypass**—Sessions for bypass LSPs.
- **ingress**—Sessions that originate from this routing device.
- **transit**—Sessions that pass through this routing device.

name *name*—(Optional) Display information about the specified static LSP or group of LSPs.

statistics—(Optional) Display accounting information about static LSPs.

transit—(Optional) Display static LSPs transiting this routing device.

Required Privilege Level view

List of Sample Output [show mpls static-lsp extensive on page 855](#)
[show mpls static-lsp statistics ingress on page 855](#)

Output Fields [Table 201 on page 854](#) describes the output fields for the **show mpls static-lsp** command. Output fields are listed in the approximate order in which they appear.

Table 201: show mpls static-lsp Output Fields

| Field Name | Field Description | Level of Output |
|--------------------------------|---|--------------------------|
| Ingress LSPs | Information about the static LSPs on the ingress routing device. Each session has one line of output. | All levels |
| Transit LSPs | Number of static LSPs on the transit routing devices and the state of these paths. MPLS learns this information by querying RSVP, which holds all the transit and egress session information. | All levels |
| Bypass LSPs | Information about the bypass LSPs configured on the routing device. Each session has one line of output. | All levels |
| LSPname | Name of the static LSP. | All levels |
| To | Destination (egress routing device) of the session. | All levels |
| State | State of the static LSP handled by this RSVP session: Up , Dn (down), or Restart . | All levels |
| Packets | Number of packet transiting the static LSP (statistics option only). | All levels |
| Bytes | Number of bytes transiting the static LSP (statistics option only). | All levels |
| Nexthop | IP address for the next-hop router for the static LSP. | detail, extensive |
| Bypass | (Bypass LSP) Destination address (egress routing device) for the bypass LSP. | All levels |
| Link protection desired | Link protection has been requested by the ingress routing device. | detail, extensive |
| LabelOperation | Label operation to perform: Push , Pop , Swap . | detail, extensive |
| Outgoing-label | Outgoing label to use for the MPLS packet in either push or swap label operations. | detail, extensive |
| Created | (Ingress LSP) Date and time the static LSP was created. | extensive |
| Bandwidth | Bandwidth configured for the static LSP. | detail, extensive |
| Resv style | (Bypass) RSVP reservation style. This field consists of two parts: the number of active reservations and the reservation style, which can be FF (fixed filter), SE (shared explicit), or WF (wildcard filter). | All levels |

Sample Output

**show mpls static-lsp
extensive**

```
user@host> show mpls static-lsp extensive
Ingress LSPs:
LSPname: alpha-to-beta, To: 192.168.14.1
  State: Dn
  Nexthop: 192.168.10.1
  LabelOperation: Push, Outgoing-label: 1000001
  Created: Thu Jan 14 16:44:43 2010
  Bandwidth: 0 bps
Total 1, displayed 1, Up 0, Down 1

Transit LSPs:
Total 0, displayed 0, Up 0, Down 0

Bypass LSPs:
Total 0, displayed 0, Up 0, Down 0
```

**show mpls static-lsp
statistics ingress**

```
user@host> show mpls static-lsp statistics ingress
Ingress LSPs:
LSPname                To           State   Packets   Bytes
alpha-to-beta          192.168.14.1 Dn      NA        NA
Total 1, displayed 1, Up 0, Down 1
```

show ted database

| | |
|------------------------------------|---|
| Syntax | show ted database <brief detail extensive> <logical-system (all <i>logical-system-name</i>)> < <i>system-name</i> > |
| Syntax (EX Series Switches) | show ted database <brief detail extensive> < <i>system-name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches. |
| Description | Display the entries in the Multiprotocol Label Switching (MPLS) traffic engineering database. |
| Options | <p>none—Display standard information about all entries in the traffic engineering database.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><i>system-name</i>—(Optional) Display traffic engineering database information for a particular system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show ted database brief on page 859 show ted database detail system-name on page 859 show ted database extensive on page 859 |
| Output Fields | Table 202 on page 856 describes the output fields for the show ted database command. Output fields are listed in the approximate order in which they appear. |

Table 202: show ted database Output Fields

| Field Name | Field Description | Level of Output |
|---------------------|---|------------------|
| TED database | Number of nodes and pseudonodes participating in IS-IS and OSPF domain routing. | All levels |
| ID | Hostname and address of the node that the link is coming from. An address of .00 indicates that the node is the routing device itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode. If the node contains a router ID, it is displayed in parentheses. | brief |
| NodeID | Hostname and address of the node that the link is coming from. An address of .00 indicates that the node is the routing device itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode. | extensive |

Table 202: show ted database Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------------------|--|-------------------------|
| Type | Type of node. It can be either Rtr (router) or Net (pseudonode). | All levels |
| Age(s) | How long since the node was last refreshed, in seconds. | All levels |
| LnkIn | Number of nodes pointing toward this node. | All levels |
| LnkOut | Number of nodes to which this node points. | All levels |
| Protocol | Protocol that reported the node information: <ul style="list-style-type: none"> • IS-IS(1)—IS-IS Level 1. • IS-IS(2)—IS-IS Level 2. • OSPF (area-number)—OSPF from the specified area. | All levels |
| To | Address on the far end of a link. | detail extensive |
| Local | Address of the local interface being used to reach the remote node. | detail extensive |
| Remote | Address of the interface on the remote node. | detail extensive |
| Metric | Configured traffic engineering metric. | extensive |
| Static BW | Total interface bandwidth in bps. | extensive |
| Reservable bandwidth | Subscription factor for the interface, which is the percentage of the link bandwidth that can be used for the RSVP reservation process. You configure this by including the subscription statement when configuring RSVP. | extensive |
| Available BW [priority] | (Must include diffserv-te statement when configuring LSPs) Amount of bandwidth actually reserved by RSVP for each priority level. The bandwidth shown is for the entire interface, not for each individual LSP. | extensive |
| Diffserv-TE BW Model | Bandwidth constraint model used by the LSPs. | extensive |
| Available BW [TE-class] | (Must include the diffserv-te statement when configuring LSPs) Amount of bandwidth actually reserved by RSVP for each traffic engineering class. | extensive |
| Static BW [CT-class] | Total interface bandwidth used by an MPLS traffic class, in bps. | extensive |

Table 202: show ted database Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---|---|------------------|
| Interface Switching Capability Descriptor (<i>n</i>) | <p>Information about the interface switching capability descriptor, which is a subtype length value (TLV) of the link TLV. <i>n</i> is the index number.</p> <ul style="list-style-type: none"> • Switching type—Type of switching to be performed on a particular link: <ul style="list-style-type: none"> • PSC-1—Packet switch-capable 1 • PSC-2—Packet switch-capable 2 • PSC-3—Packet switch-capable 3 • PSC-4—Packet switch-capable 4 • L2SC—Layer-2-switch-capable • TDM—Time-division-multiplexing-capable • LSC—Lambda switch-capable • FSC—Fiber switch-capable • Encoding type—Encoding of the LSP being requested: <ul style="list-style-type: none"> • Packet • Ethernet • ANSI/ETSI PDH • Reserved • SDH /SONET • Digital Wrapper • Lambda (photonic) • Fiber • FiberSDH/SONET • Maximum LSP BW [priority] bps—Maximum LSP bandwidth information. Amount of bandwidth actually reserved for each priority level. The bandwidth shown is for the entire interface. <ul style="list-style-type: none"> • [<i>n</i>]—Priority level. The range is from 0 (high) through 7 (low). • <i>n</i> Mbps—Amount of the maximum bandwidth. • Minimum LSP BW—Minimum LSP bandwidth in Mbps. Amount of bandwidth actually reserved for each priority level. The bandwidth shown is for the entire interface. Minimum LSP BW is displayed only when switching type is PSC-1 or TDM. • Interface MTU—Displayed only when switching type is TDM. • Interface supports standard SONET/SDH—Displayed only when switching type is TDM. | extensive |

Sample Output

**show ted database
brief**

```
user@host> show ted database brief
TED database: 6 ISIS nodes 6 INET nodes
ID                               Type Age(s) LnkIn LnkOut Protocol
cheviot.00(123.456.1.10)         Rtr   383     1     1 IS-IS(2) IS-IS(1)
corriedale.00(123.456.1.11)      Rtr    36     2     0 IS-IS(2) IS-IS(1)
wolff.00(123.456.1.12)          Rtr   399     0     0 IS-IS(2) IS-IS(1)
perendale.00(123.456.1.13)       Rtr   385     2     0 IS-IS(2) IS-IS(1)
merino.00(123.456.1.14)          Rtr   379     1     3 IS-IS(2) IS-IS(1)
romney.00(123.456.1.15)          Rtr   427     0     2 IS-IS(2) IS-IS(1)
```

**show ted database
detail system-name**

```
user@host> show ted database detail merino
TED database: 6 ISIS nodes 6 INET nodes
NodeID: merino.00(123.456.1.14)
  Type: Rtr, Age: 507 secs, LinkIn: 1, LinkOut: 3
  Protocol: IS-IS(2)
    To: corriedale.00(123.456.1.11), Local: 123.456.8.206, Remote: 123.456.8.207

    To: perendale.00(123.456.1.13), Local: 123.456.8.204, Remote: 123.456.8.205
    To: cheviot.00(123.456.1.10), Local: 123.456.10.65, Remote: 123.456.10.66
  Protocol: IS-IS(1)
    To: corriedale.00(123.456.1.11), Local: 123.456.8.206, Remote: 123.456.8.207

    To: perendale.00(123.456.1.13), Local: 123.456.8.204, Remote: 123.456.8.205
    To: cheviot.00(123.456.1.10), Local: 123.456.10.65, Remote: 123.456.10.66
```

**show ted database
extensive**

```
user@host> show ted database extensive
TED database: 0 ISIS nodes 2 INET nodes
NodeID: 10.255.245.196
  Type: Rtr, Age: 46 secs, LinkIn: 1, LinkOut: 1
  Protocol: OSPF(0.0.0.0)
    To: 10.255.245.24, Local: 4.4.4.4, Remote: 5.5.5.5
    Metric: 1
    Static BW: 155.52Mbps
    Reservable BW: 155.52Mbps
    Available BW [TE-class] bps:
      [te0] 155.52Mbps [te1] 155.52Mbps [te2] 155.52Mbps [te3] 155.52Mbps
      [te4] 155.52Mbps [te5] 155.52Mbps [te6] 155.52Mbps [te7] 155.52Mbps

    Diffserv-TE BW model: Maximum allocation model
    Static BW [CT-class] bps:
      [ct0] 155.52Mbps [ct1] 155.52Mbps [ct2] 155.52Mbps [ct3] 155.52Mbps

    Interface Switching Capability Descriptor(1):
      Switching type: PSC-1
      Encoding type: SDH/SONET
      Maximum LSP BW [priority] bps:
        [0] 155.52Mbps [1] 155.52Mbps [2] 155.52Mbps [3] 155.52Mbps
        [4] 155.52Mbps [5] 155.52Mbps [6] 155.52Mbps [7] 155.52Mbps
      Minimum LSP BW: 155.52Mbps
      Interface MTU: 1285
    Interface Switching Capability Descriptor(2):
      Switching type: TDM
      Encoding type: SDH/SONET
      Maximum LSP BW [priority] bps:
        [0] 155.52Mbps [1] 155.52Mbps [2] 155.52Mbps [3] 155.52Mbps
```

[4] 155.52Mbps [5] 155.52Mbps [6] 155.52Mbps [7] 155.52Mbps
Minimum LSP BW: 155.52Mbps
Interface supports standard SONET/SDH

show ted link

| | |
|------------------------------------|---|
| Syntax | show ted link <brief detail> <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switches) | show ted link <brief detail> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches. |
| Description | Display Multiprotocol Label Switching (MPLS) traffic engineering database link information. |
| Options | none —Display standard information about traffic engineering database link information. brief detail —(Optional) Display the specified level of output. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | view |
| List of Sample Output | show ted link brief on page 862 show ted link detail on page 862 |
| Output Fields | Table 203 on page 861 describes the output fields for the show ted link command. Output fields are listed in the approximate order in which they appear. |

Table 203: show ted link Output Fields

| Field Name | Field Description | Level of Output |
|-----------------|---|-----------------|
| ID | Hostname and address of the node that the link is coming from. An address of .00 indicates that the node is the routing device itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode. | brief |
| -->ID | Hostname and address of the node that the link is going to. An address of .00 indicates that the node is the routing device itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode. | brief |
| <i>hostname</i> | Hostname and address of the node that the link is coming from. An address of .00 indicates that the node is the routing device itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode. | detail |
| <i>hostname</i> | Hostname and address of the node that the link is going to. An address of .00 indicates that the node is the routing device itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode. | detail |
| Local Path | Number of paths CSPF on the local routing device has placed on the link. | All levels |

Table 203: show ted link Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-----------------|--|-----------------|
| Local BW | Amount of bandwidth the local routing device has placed on the link. | All levels |

Sample Output

show ted link brief

```

user@host> show ted link brief
TED link:
ID                                ->ID                                LocalPath LocalBW
cheviot.00(123.456.1.10)         merino.00(123.456.1.14)           0 0bps
merino.00(123.456.1.14)         corriedale.00(123.456.1.11)       0 0bps
merino.00(123.456.1.14)         perendale.00(123.456.1.13)        0 0bps
merino.00(123.456.1.14)         cheviot.00(123.456.1.10)          0 0bps
romney.00(123.456.1.15)         corriedale.00(123.456.1.11)       0 0bps
romney.00(123.456.1.15)         perendale.00(123.456.1.13)        0 0bps

```

show ted link detail

```

user@host> show ted link detail
TED link:
cheviot.00(123.456.1.10)->merino.00(123.456.1.14), LocalPath 0
  localBW [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
  localBW [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps
merino.00(123.456.1.14)->corriedale.00(123.456.1.11), LocalPath 0
  localBW [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
  localBW [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps
merino.00(123.456.1.14)->perendale.00(123.456.1.13), LocalPath 0
  localBW [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
  localBW [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps
merino.00(123.456.1.14)->cheviot.00(123.456.1.10), LocalPath 0
  localBW [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
  localBW [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps
romney.00(123.456.1.15)->corriedale.00(123.456.1.11), LocalPath 0
  localBW [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
  localBW [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps
romney.00(123.456.1.15)->perendale.00(123.456.1.13), LocalPath 0
  localBW [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
  localBW [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps

```


show ted protocol

| | |
|------------------------------------|---|
| Syntax | show ted protocol <brief detail> <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switches) | show ted protocol <brief detail> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches. |
| Description | Display information about the protocols from which the Multiprotocol Label Switching (MPLS) traffic engineering database learned about its nodes. |
| Options | <p>none—Display standard information about the protocols from which the traffic engineering database learned about its nodes.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show ted protocol on page 864 |
| Output Fields | Table 204 on page 863 describes the output fields for the show ted protocol command. Output fields are listed in the approximate order in which they appear. |

Table 204: show ted protocol Output Fields

| Field Name | Field Description |
|----------------------|--|
| Protocol name | Protocol that reported the node information: <ul style="list-style-type: none"> IS-IS(1)—IS-IS Level 1. IS-IS(2)—IS-IS Level 2. OSPF (<i>area-number</i>)—OSPF from the specified area. |
| Credibility | If the protocols provide conflicting information about a node, the protocol with the highest credibility value is the one that the traffic engineering database uses. |
| Self node | Address the protocol uses as the local address. |

Sample Output

`show ted protocol`

```
user@host> show ted protocol
Protocol name      Credibility Self node
IS-IS(2)           2 (highest) corriedale.00(123.456.1.11)
IS-IS(1)           1           corriedale.00(123.456.1.11)
```

RSVP Operational Mode Commands

Table 205 on page 865 summarizes the command-line interface (CLI) commands you can use to monitor Resource Reservation Protocol (RSVP) sessions. Commands are listed in alphabetical order.

Table 205: RSVP Operational Mode Commands

| Task | Command |
|--|-------------------------------------|
| Clear RSVP sessions and trigger fast reroute optimization. | <code>clear rsvp session</code> |
| Clear RSVP packet and error counters. | <code>clear ripng statistics</code> |
| Display the status of interfaces on which RSVP is running. | <code>show rsvp interface</code> |
| Display RSVP neighbors. | <code>show rsvp neighbor</code> |
| Display currently active RSVP sessions. | <code>show rsvp session</code> |
| Display RSVP packet and error counters. | <code>show rsvp statistics</code> |
| Display RSVP version and configuration information. | <code>show rsvp version</code> |



NOTE: For more RSVP-related commands, such as `show route protocol`, `show route instance`, and `show route table`, see Protocol-Independent Routing Operational Mode Commands.

For information about the `monitor label-switched path` command, used to monitor an RSVP LSP in real time, see the *Junos System Basics and Services Command Reference*.

For information about how to configure RSVP, see the *Junos MPLS Applications Configuration Guide*.

clear rsvp session

| | |
|------------------------------------|--|
| Syntax | <pre>clear rsvp session <connection-destination address> <connection-source address> <gracefully> <logical-system (all logical-system-name)> <lsp-id identifier> <name name> <optimize-fast-reroute> <tunnel-id identifier></pre> |
| Syntax (EX Series Switches) | <pre>clear rsvp session <connection-destination address> <connection-source address> <gracefully> <lsp-id identifier> <name name> <optimize-fast-reroute> <tunnel-id identifier></pre> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches. |
| Description | Reset and restart Resource Reservation Protocol (RSVP) sessions. |
| Options | <p>none—Reset and restart all RSVP sessions for which this routing device is the ingress, transit, or egress routing device.</p> <p>connection-source address—(Optional) Source address for GMPLS and MPLS LSPs from the RSVP sender template.</p> <p>connection-destination address—(Optional) Destination address for GMPLS and MPLS LSPs from the RSVP sender template.</p> <p>gracefully—(Optional) Gracefully reset an RSVP session for a nonpacket LSP in two passes. In the first pass, the Admin-Status object is signaled along the path to the other endpoint of the RSVP session. In the second pass, the path used by the RSVP session is torn down. This option can only be used on the ingress or egress routing device of the RSVP session and is only valid for nonpacket LSPs.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>lsp-id identifier—(Optional) LSP identifier (source port) for the RSVP sender template.</p> <p>name name—(Optional) Reset and restart the specified RSVP session.</p> <p>optimize-fast-reroute—(Optional) Begin fast reroute optimization.</p> <p>tunnel-id identifier—(Optional) Tunnel identifier (destination port) for the RSVP session.</p> |

Required Privilege Level clear

Related Documentation

- [clear mpls lsp on page 809](#)
- [show rsvp session on page 880](#)

List of Sample Output [clear rsvp session on page 867](#)

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

Sample Output

`clear rsvp session` user@host> clear rsvp session

clear rsvp statistics

| | |
|------------------------------------|---|
| Syntax | clear rsvp statistics <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switches) | clear rsvp statistics |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches. |
| Description | Clear Resource Reservation Protocol (RSVP) packet and error statistics. |
| Options | none —Clear RSVP packet and error statistics. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• show rsvp statistics on page 890 |
| List of Sample Output | clear rsvp statistics on page 868 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear rsvp statistics user@host> clear rsvp statistics

show rsvp interface

| | |
|------------------------------------|---|
| Syntax | show rsvp interface <brief detail extensive> <link-management> <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switches) | show rsvp interface <brief detail extensive> <link-management> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches. |
| Description | Display the status of Resource Reservation Protocol (RSVP)-enabled interfaces and packet statistics. |
| Options | <p>none—Display standard information about the status of RSVP-enabled interfaces and packet statistics.</p> <p>brief detail extensive link-management—(Optional) Display the specified level of output.</p> <p>link-management—(Optional) Use the link-management option to display the control peers and corresponding TE-link information created by the Link Management Protocol (LMP).</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show rsvp interface brief on page 873 show rsvp interface detail on page 873 show rsvp interface extensive on page 873 show rsvp interface link-management on page 874 |
| Output Fields | Table 206 on page 869 lists the output fields for the show rsvp interface command. Output fields are listed in the approximate order in which they appear. |

Table 206: show rsvp interface Output Fields

| Field Name | Field Description | Level of Output |
|-----------------------|--|-----------------|
| RSVP interface | Number of interfaces on which RSVP is active. Each interface has one line of output. | All levels |
| Interface | Name of the interface. | All levels |
| Index | Index of the interface. | detail |

Table 206: show rsvp interface Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-------------------------------|--|------------------|
| State | State of the interface. <ul style="list-style-type: none"> • Disabled—No traffic engineering information is displayed. • Down—Interface is not operational. • Enabled—Displays traffic engineering information. • Up—Interface is operational. | All levels |
| NoAuthentication | Interface does not support RSVP authentication. | detail |
| NoAggregate | Interface does not support refresh reduction. | detail |
| NoReliable | Interface does not support refresh reduction message ID extension. | detail |
| NoLinkProtection | Interface does not support link protection. | detail |
| HelloInterval | Frequency at which RSVP hellos are sent on this interface (in seconds). | detail |
| Address | IP address of the local interface. | detail |
| Active control channel | Next-hop link address to transmit messages. | None specified |
| TElink | Traffic-engineered links that are managed by the peer they are associated with. | None specified |
| Active resv | Number of reservations that are actively reserving bandwidth on the interface. | All levels |
| PreemptionCnt | Number of times an RSVP session was preempted on this interface. | detail |
| Update threshold | Percentage change in reserved bandwidth to trigger an IGP update. | detail |
| Subscription | User-configured subscription factor. | All levels |
| bc number | Bandwidth allocated for the specified bandwidth constraint. | extensive |
| ct number | Bandwidth allocated for the specified class type. | extensive |
| Static BW | Total interface bandwidth, in bps. | All levels |
| Available BW | Amount of bandwidth that RSVP is allowed to reserve, in bps. It is equal to (static bandwidth * subscription factor). | all levels |
| Reserved BW | Currently reserved bandwidth, in bps. | All levels |
| SoftPreemptionCnt | Number of times a soft preemption occurred on this interface. This number is not included in the PreemptionCnt value. | detail |
| Overbooked BW | Currently overbooked bandwidth, in bps, by class type (ct0 through ct3). | detail |

Table 206: show rsvp interface Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------------------|---|------------------|
| Highwater mark | Highest bandwidth that has ever been reserved on this interface, in bps. | brief |
| PacketType | Type of RSVP packet. | detail |
| Total Sent | Total number of packets sent. | detail |
| Total Received | Total number of packets received since RSVP was enabled. | detail |
| Last 5 seconds Sent | Number of packets sent in the last 5 seconds. | detail |
| Last 5 seconds Received | Number of packets received in the last 5 seconds. | detail |
| Path | Statistics about Path messages, which are sent from the RSVP sender along the data paths and store path state information in each node along the path. | detail |
| PathErr | Statistics about PathErr messages, which are advisory messages that are sent upstream to the sender. | detail |
| PathTear | Statistics about PathTear messages, which remove path states and dependent reservation states in any routers along a path. | detail |
| Resv | Statistics about Resv messages, which are sent from the RSVP receiver along the data paths and store reservation state information in each node along the path. | detail |
| ResvErr | Statistics about ResvErr messages, which are advisory messages that are sent when an attempt to establish a reservation fails. | detail |
| ResvTear | Statistics about ResvTear messages, which remove reservation states along a path. | detail |
| Hello | Number of RSVP hello packets that have been sent to and received from the neighbor. | detail |
| Ack | Acknowledge message for refresh reductions. | detail |
| Srefresh | Summary refresh messages. | detail |
| EndtoEnd RSVP | Statistics for the number of end-to-end RSVP messages sent. | detail |
| Queue | CoS transmit queue number and its associated forwarding class designation. | extensive |
| TxRate | Configured bandwidth in Mbps and configured bandwidth as a percentage of the specified queue. | extensive |
| Priority | Weight of the queue relative to other configured queues, in percentage. | extensive |

Table 206: show rsvp interface Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-----------------------------|--|-----------------|
| <i>queue-priority-value</i> | Low, High, None, or Exact. None indicates no rate limiting. Exact indicates the queue transmits at the configured rate only. | extensive |

Sample Output

show rsvp interface brief

```
user@host> show rsvp interface brief
RSVP interface: 1 active
```

| Interface | State | Active resv | Subscr- ption | Static BW | Available BW | Reserved BW | Highwater mark |
|-----------|-------|-------------|------------------|--------------|-----------------|----------------|-------------------|
| de0.0 | Up | 1 | 23% | 10Mbps | 989.992kbps | 1.31Mbps | 1.31Mbps |

show rsvp interface detail

```
user@host> show rsvp interface detail
so-0/1/1.0 Index 6, State: Ena/Up
  NoAuthentication, NoAggregate, NoReliable, NoLinkProtection
  HelloInterval 3(second)
  Address 192.168.207.29, 10.255.245.194
  ActiveResv 0, PreemptionCnt 0, Update threshold 10%
  Subscription 100%, StaticBW 155.52Mbps, AvailableBW 155.52Mbps
  ReservedBW [0] 155Mbps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
  SoftPreemptionCnt1
  OverbookedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 155Mbps[5] 0bps[6] 0bps[7] 0bps
  PacketType
```

| | Total | | Last 5 seconds | |
|---------------|-------|----------|----------------|----------|
| | Sent | Received | Sent | Received |
| Path | 16 | 0 | 1 | 0 |
| PathErr | 0 | 0 | 0 | 0 |
| PathTear | 1 | 0 | 0 | 0 |
| Resv | 0 | 11 | 0 | 1 |
| ResvErr | 0 | 0 | 0 | 0 |
| ResvTear | 0 | 0 | 0 | 0 |
| Hello | 66 | 67 | 1 | 1 |
| Ack | 0 | 0 | 0 | 0 |
| Srefresh | 0 | 0 | 0 | 0 |
| EndtoEnd RSVP | 0 | 0 | 0 | 0 |

...

show rsvp interface extensive

```
user@host> show rsvp interface extensive
so-1/0/0.0 Index 72, State Ena/Up
  NoAuthentication, NoAggregate, NoReliable, NoLinkProtection
  HelloInterval 9(second)
  Address 192.168.213.22, 10.255.240.175
  ActiveResv 1, PreemptionCnt 0, Update threshold 10%
  Subscription 100%,
  bc0 = (ct0+ct1+ct2+ct3), StaticBW 622.08Mbps
  bc1 = (ct1+ct2+ct3), StaticBW 466.56Mbps
  bc2 = (ct2+ct3), StaticBW 311.04Mbps
  bc3 = ct3, StaticBW 155.52Mbps
  ct0: StaticBW 155.52Mbps, AvailableBW 522.08Mbps
  ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
  ct1: StaticBW 155.52Mbps, AvailableBW 366.56Mbps
  ReservedBW [0] 100Mbps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps

  ct2: StaticBW 155.52Mbps, AvailableBW 311.04Mbps
  ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
  ct3: StaticBW 155.52Mbps, AvailableBW 155.52Mbps
  ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
  Queue
```

| | TxRate | Priority | Exact |
|---|------------|----------|-------|
| 0 | 155.52Mbps | 25% | Low |
| 1 | 155.52Mbps | 25% | Low |
| 2 | 155.52Mbps | 25% | Low |
| 3 | 155.52Mbps | 25% | Low |

**show rsvp interface
link-management**

```
user@host> show rsvp interface link-management
RSVP interface: 2 active
PEER-C State: Up
Active Control Channel: so-0/1/0.0

  TElink: TElnk1, Link ID: 37811
  ActiveResv 0, PreemptionCnt 0
  StaticBW 155.52Mbps, ReservedBW: 0bps, AvailableBW: 155.52Mbps

  TElink: TElnk2, Link ID: 37808
  ActiveResv 1, PreemptionCnt 0
  StaticBW 155.52Mbps, ReservedBW: 0bps, AvailableBW: 155.52Mbps

PEER-B State: Up
Active Control Channel: so-1/0/0.0

  TElink: TElnkAB1, Link ID: 1598
  ActiveResv 0, PreemptionCnt 0
  StaticBW 622.08Mbps, ReservedBW: 0bps, AvailableBW: 622.08Mbps

  TElink: TElnkAB2, Link ID: 1597
  ActiveResv 0, PreemptionCnt 0
  StaticBW 622.08Mbps, ReservedBW: 0bps, AvailableBW: 622.08Mbps
```

show rsvp neighbor

| | |
|------------------------------------|--|
| Syntax | show rsvp neighbor <brief detail> <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switches) | show rsvp neighbor <brief detail> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches. |
| Description | Display Resource Reservation Protocol (RSVP) neighbors that were discovered dynamically during the exchange of RSVP packets. |
| Options | none —Display standard information about RSVP neighbors. brief detail —(Optional) Display the specified level of output. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | view |
| List of Sample Output | show rsvp neighbor on page 879 show rsvp neighbor detail on page 879 |
| Output Fields | Table 207 on page 875 lists the output fields for the show rsvp neighbor command. Output fields are listed in the approximate order in which they appear. |

Table 207: show rsvp neighbor Output Fields

| Field Name | Field Description | Level of Output |
|----------------------|---|-----------------|
| RSVP neighbor | Number of neighbors that the routing device has learned of. Each neighbor has one line of output. | All levels |
| via | Name of the interface where the neighbor has been detected. In the case of generalized MPLS (GMPLS) LSPs, the name of the peer where the neighbor has been detected. | detail |
| Address | Address of a learned neighbor. | All levels |
| Idle | Length of time the neighbor has been idle, in seconds. | All levels |
| Up/Dn | Number of neighbor up or down transitions detected by RSVP hello packets. If the up count is 1 greater than the down count, the neighbor is currently up. Otherwise, the neighbor is down. Neighbors that do not support RSVP hello packets, such as routers running Junos OS Release 3.2 or earlier, are not reported as up or down. | All levels |

Table 207: show rsvp neighbor Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|----------------------------|---|-----------------|
| Up cnt and Down cnt | Number of neighbor up or down transitions detected by RSVP hello packets. If the up count is 1 greater than the down count, the neighbor is currently up. Otherwise, the neighbor is down. Neighbors that do not support RSVP hello packets, such as routers running Junos OS Release 3.2 or earlier, are not reported as up or down. | detail |
| status | State of the RSVP neighbor: <ul style="list-style-type: none"> • Up—Routing device can detect RSVP Hello messages from the neighbor. • Down—Routing device has received one of the following indications: <ul style="list-style-type: none"> • Communication failure from the neighbor. • Communication from IGP that the neighbor is unavailable. • Change in the sequence numbers in the RSVP Hello messages sent by the neighbor. • Restarting—RSVP neighbor is unavailable and might be restarting. The neighbor remains in this state until it has restarted or is declared dead. This state is possible only when graceful restart is enabled. • Restarted—RSVP neighbor has restarted and is undergoing state recovery (graceful restart) procedures. • Dead—Routing device has lost all communication with the RSVP neighbor. Any RSVP sessions with that neighbor are torn down. | detail |
| LastChange | Time elapsed since the neighbor state changed either from up to down or from down to up. The format is <i>hh:mm:ss</i> . | All levels |
| Last changed time | Time elapsed since the neighbor state changed either from up to down or from down to up. | detail |
| HelloInt | Frequency at which RSVP hellos are sent on this interface (in seconds). | All levels |
| HelloTx/Rx | Number of hello packets sent to and received from the neighbor. | All levels |
| Hello | Number of RSVP hello packets that have been sent to and received from the neighbor. | detail |
| Message received | Number of Path and Resv messages that this routing device has received from the neighbor. | detail |
| Remote Instance | Identification provided by the remote routing device during Hello message exchange. | detail |
| Local Instance | Identification sent to the remote routing device during Hello message exchange. | detail |

Table 207: show rsvp neighbor Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|------------------------------|--|-----------------|
| Refresh reduction | <p>Measure of processing overhead requests of refresh messages. Refresh reduction extensions improve routing device performance by reducing the process overhead, thus increasing the number of LSPs a routing device can support. Refresh reduction can have the following values:</p> <ul style="list-style-type: none"> • operational—All four RSVP refresh reduction extensions—message ack, bundling, summary refresh, and staged refresh timer—are functional between the two neighboring routing devices. For a detailed explanation of these extensions, see RFC 2961. • incomplete—Some RSVP refresh reduction extensions are functional between the two neighboring routing devices. • no operational—Either the refresh reduction feature has been turned off, or the remote routing device cannot support the refresh reduction extensions. | detail |
| Remote end | <p>Neighboring routing device's status with regard to refresh reduction:</p> <ul style="list-style-type: none"> • enabled—Remote routing device has requested refresh reduction during RSVP message exchanges. • disabled—Remote routing device does not require refresh reduction. | detail |
| Ack-extension | <p>An RSVP refresh reduction extension:</p> <ul style="list-style-type: none"> • enabled—Both local and remote routing devices support the ack-extension (RFC 2961). • disabled—Remote routing device does not support the ack-extension. | detail |
| Link protection | <p>Status of the MPLS fast reroute mechanism that protects traffic from link failure:</p> <ul style="list-style-type: none"> • enabled—Link protection feature has been turned on, protecting the neighbor with a bypass LSP. • disabled—No link protection feature has been enabled for this neighbor. | detail |
| LSP name | Name of the bypass LSP. | detail |
| Bypass LSP | <p>Status of the bypass LSP. It can have the following values:</p> <ul style="list-style-type: none"> • does not exist—Bypass LSP is not available. • connecting—Routing device is in the process of establishing a bypass LSP, and the LSP is not available for link protection at the moment. • operational—Bypass LSP is up and running. • down—Bypass LSP has gone down, with the most probable cause a node or a link failure on the bypass path. | detail |
| Backup routes | Number of user LSPs (or routes) that are being protected by a bypass LSP (before link failure). | detail |
| Backup LSPs | Number of LSPs that have been temporarily established to maintain traffic by refreshing the downstream LSPs during link failure (not a one-to-one correspondence). | detail |
| Bypass explicit route | Explicit route object's (ERO) path that is taken by the bypass LSP. | detail |

Table 207: show rsvp neighbor Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|----------------------|--|-----------------|
| Restart time | Length of time a neighbor waits to receive a Hello from the restarting node before declaring the node dead and deleting the states (in milliseconds). | detail |
| Recovery time | Length of time during which the restarting node attempts to recover its lost states with help from its neighbors (in milliseconds). Recovery time is advertised by the restarting node to its neighbors, and applies to nodal faults. The restarting node considers its graceful restart complete after this time has elapsed. | detail |

Sample Output

show rsvp neighbor

```
user@host> show rsvp neighbor
RSVP neighbor: 2 learned
Address      Idle Up/Dn LastChange HelloInt HelloTx/Rx
192.168.207.203  0 3/2    13:01      3   366/349
192.168.207.207  0 1/0    22:49      3   448/448
```

show rsvp neighbor detail

```
user@host> show rsvp neighbor detail
RSVP neighbor: 2 learned
Address: 192.168.207.203 via: ecstasy1 status: Up
  Last changed time: 28:47, Idle: 0 sec, Up cnt: 3, Down cnt: 2
  Message received: 632
  Hello: sent 673, received 656, interval 3 sec
  Remote instance: 0x6432838a, Local instance: 0x74b72e36
  Refresh reduction: operational
    Remote end: enabled, Ack-extension: enabled
  Link protection: enabled
    LSP name: Bypass_to_192.168.207.203
    Bypass LSP: operational, Backup routes: 1, Backup LSPs: 0
    Bypass explicit route: 192.168.207.207 192.168.207.224
  Restart time: 60000 msec, Recovery time: 0 msec
```

show rsvp session

| | |
|------------------------------------|--|
| Syntax | <pre>show rsvp session <brief detail extensive terse> <bidirectional unidirectional> <bypass> <down up> <interface <i>interface-name</i>> <logical-system (all <i>logical-system-name</i>)> <lsp-type> <name <i>session-name</i>> <p2mp> <session-type> <statistics> <te-link <i>te-link</i>></pre> |
| Syntax (EX Series Switches) | <pre>show rsvp session <brief detail extensive terse> <bidirectional unidirectional> <bypass> <down up> <interface <i>interface-name</i>> <lsp-type> <name <i>session-name</i>> <p2mp> <session-type> <statistics> <te-link <i>te-link</i>></pre> |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches. |
| Description | Display information about Resource Reservation Protocol (RSVP) sessions. |
| Options | <p>none—Display standard information about all RSVP sessions.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>bidirectional unidirectional—(Optional) Display information about bidirectional or unidirectional RSVP sessions only, respectively.</p> <p>bypass—(Optional) Display RSVP sessions for bypass LSPs.</p> <p>down up—(Optional) Display only LSPs that are inactive or active, respectively.</p> <p>interface <i>interface-name</i>—(Optional) Display RSVP sessions for the specified interface only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><i>lsp-type</i>—(Optional) Display information about RSVP sessions with regard to LSPs:</p> <ul style="list-style-type: none">• bypass—Sessions used for bypass LSPs. |

- **lsp**—Sessions used to set up LSPs.
- **nolsp**—Sessions not used to set up LSPs.

name session-name—(Optional) Display information about the named session.

p2mp—(Optional) Display point-to-multipoint information.

session-type—(Optional) Display information about a particular session type:

- **egress**—Sessions that terminate on this routing device.
- **ingress**—Sessions that originate from this routing device.
- **transit**—Sessions that transit through this routing device.

statistics—(Optional) Display packet statistics.

te-link te-link—(Optional) Display sessions with reservations on the specified TE link.

Required Privilege Level view

Related Documentation [• clear rsvp session on page 866](#)

List of Sample Output [show rsvp session on page 886](#)
[show rsvp session statistics on page 886](#)
[show rsvp session detail on page 886](#)
[show rsvp session detail \(Path MTU Output Field\) on page 886](#)
[show rsvp session detail \(GMPLS\) on page 887](#)
[show rsvp session extensive on page 887](#)
[show rsvp session p2mp \(Ingress Router\) on page 888](#)
[show rsvp session p2mp \(Transit Router\) on page 888](#)

Output Fields [Table 208 on page 881](#) describes the output fields for the **show rsvp session** command. Output fields are listed in the approximate order in which they appear.

Table 208: show rsvp session Output Fields

| Field Name | Field Description | Level of Output |
|--------------|--|-----------------|
| Ingress RSVP | Information about ingress RSVP sessions. | detail |
| Ingress RSVP | Information about ingress RSVP sessions. Each session has one line of output. | All levels |
| Egress RSVP | Information about egress RSVP sessions. | All levels |
| Transit RSVP | Information about the transit RSVP sessions. | All levels |
| P2MP name | (Appears only when the p2mp option is specified). Name of the point-to-multipoint LSP path. | All levels |

Table 208: show rsvp session Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------------|--|---------------------|
| P2MP branch count | (Appears only when the p2mp option is specified). Number of LSPs receiving packets from the point-to-multipoint LSP. | All levels |
| To | Destination (egress routing device) of the session. | All levels |
| From | Source (ingress routing device) of the session. | All levels |
| State | State of the path: Up , Down , or AdminDn . AdminDn indicates that the LSP is being taken down gracefully. | All levels |
| Address | Destination (egress routing device) of the LSP. | detail |
| From | Source (ingress routing device) of the session. | detail |
| LSPstate | State of the LSP that is being handled by this RSVP session. It can be either Up , Dn (down), or AdminDn . AdminDn indicates that the LSP is being taken down gracefully. | brief detail |
| Rt | Number of active routes (prefixes) that have been installed in the routing table. For ingress RSVP sessions, the routing table is the primary IPv4 table (inet.0). For transit and egress RSVP sessions, the routing table is the primary MPLS table (mpls.0). | brief |
| Active Route | Number of active routes (prefixes) that have been installed in the forwarding table. For ingress RSVP sessions, the forwarding table is the primary IPv4 table (inet.0). For transit and egress RSVP sessions, the forwarding table is the primary MPLS table (mpls.0). | detail |
| LSPname | Name of the LSP. | brief detail |
| LSPpath | Indicates whether the RSVP session is for the primary or secondary LSP path. LSPpath can be either primary or secondary and can be displayed on the ingress, egress, and transit routing devices. LSPpath can also indicate when a graceful LSP deletion has been triggered. | detail |
| Bypass | (Egress routing device) Destination address for the bypass LSP. | detail |
| Bidir | (When LSP is bidirectional) LSP will allow data to travel in both directions between GMPLS devices. | detail |
| Bidirectional | (When LSP is bidirectional) LSP will allow data to travel both ways between GMPLS devices. | detail |
| Upstream label in | (When LSP is bidirectional) Incoming label for reverse direction traffic for this LSP. | detail |
| Upstream label out | (When LSP is bidirectional) Outgoing label for reverse direction traffic for this LSP. | detail |

Table 208: show rsvp session Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------------------|--|-------------------------|
| Recovery label received | (When LSP is bidirectional) Label the upstream node suggests for use in the Resv message that is sent. | detail |
| Recovery label sent | (When LSP is bidirectional) Label the downstream node suggests for use in its Resv messages that is returned. | detail |
| Suggested label received | (When LSP is bidirectional) Label the upstream node suggests for use in the Resv message that is sent. | detail |
| Suggested label sent | (When LSP is bidirectional) Label the downstream node suggests for use in its Resv message that is returned. | detail |
| Resv style or Style | RSVP reservation style. This field consists of two parts. The first is the number of active reservations. The second is the reservation style, which can be FF (fixed filter), SE (shared explicit), or WF (wildcard filter). | brief detail |
| Label in | Incoming label for this LSP. | brief detail |
| Label out | Outgoing label for this LSP. | brief detail |
| Time left | Number of seconds remaining in the lifetime of the reservation. | brief detail |
| Since | Date and time when the RSVP session was initiated. | detail |
| Tspec | Sender's traffic specification, which describes the sender's traffic parameters. | detail |
| DiffServ info | Indicates whether the LSP is a multiclass LSP (multiclass diffServ-TE LSP) or a Differentiated-Services-aware traffic engineering LSP (diffServ-TE LSP). | detail |
| bandwidth | Bandwidth for each class type (ct0 , ct1 , ct2 , or ct3). | detail |
| Port number | Protocol ID and sender/receiver port used in this RSVP session. | detail |
| Attrib flags | Non-PHP indicates that ultimate hop popping has been requested by the LSP using this RSVP session | extensive |
| FastReroute desired | Fast reroute has been requested by the ingress routing device. | detail |
| Soft preemption desired | Soft preemption has been requested by the ingress routing device. | detail |
| FastReroute desired | (Data [not a bypass or backup] LSP when the protection scheme has been requested) Fast reroute (one-to-one backup) has been requested by the ingress routing device. | detail extensive |
| Link protection desired | (Data [not a bypass or backup] LSP when the protection scheme has been requested) Link protection (many-to-one backup) has been requested by the ingress routing device. | detail extensive |

Table 208: show rsvp session Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---|---|-------------------------|
| Node/Link protection desired | (Data [not a bypass or backup] LSP when the protection scheme has been requested) Node and link protection (many-to-one backup) has been requested by the ingress routing device. | detail extensive |
| Type | <p>LSP type:</p> <ul style="list-style-type: none"> • Link protected LSP—LSP has been protected by link protection at the outgoing interface. The name of the bypass used is also listed here (extensive). • Node/Link protected LSP—LSP has been protected by node and link protection at the outgoing interface. The name of the bypass used is also listed here (extensive). • Protection down—LSP is not currently protected. • Bypass LSP—LSP that is used to protect one or more user LSPs in case of link failure. • Backup LSP at Point-of-Local-Repair (PLR)—LSP that has been temporarily established to protect a user LSP at the ingress of a failed link. • Backup LSP at Merge Point (MP)—LSP that has been temporarily established to protect a user LSP at the egress of a failed link. | detail extensive |
| New bypass | New bypass (the bypass name is also displayed) has been activated to protect the LSP. | extensive |
| Link protection up, using <i>bypass-name</i> | Link protection (the bypass name is also displayed) has been activated for the LSP. | extensive |
| Creating backup LSP, link down | A link down event occurred, and traffic is being switched over to the bypass LSP. | extensive |
| Deleting backup LSP, protected LSP restored | Link has come back up and the LSP has been restored. Because the backup LSP is no longer needed, it is deleted. | extensive |
| Path mtu | Displays the value of the path MTU received from the network (through signaling) and the value used for forwarding. This value is only displayed on ingress routing devices with the allow-fragmentation statement configured at the [edit protocols mpls path-mtu] hierarchy level. If there is a detour LSP, the path MTU for the detour is also displayed. | detail |
| PATH rcvfrom | Address of the previous-hop (upstream) routing device or client, interface the neighbor used to reach this routing device, and number of packets received from the upstream neighbor. | detail |
| Adspec | MTU signaled from the ingress routing device to the egress routing device by means of the adspec object. | detail |
| PATH sentto | Address of the next-hop (downstream) routing device or client, interface used to reach this neighbor (or peer-name in the GMPLS LSP case), and number of packets sent to the downstream routing device. | detail |

Table 208: show rsvp session Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------|---|-----------------|
| Explct route | Explicit route for the session. Normally this value will be the same as that of record route. Differences indicate that path rerouting has occurred, typically during fast reroute. | detail |
| Record route | Recorded route for the session, taken from the record route object. Normally this value will be the same as that of explct route. Differences indicate that path rerouting has occurred, typically during fast reroute. | detail |

Sample Output

show rsvp session

```

user@host> show rsvp session
Ingress RSVP: 1 sessions
To          From          State  Rt Style Labelin Labelout LSPName
10.255.245.214 10.255.245.212 AdminDn 0 1 FF      -    22293 LSP Bidir
Total 1 displayed, Up 1, Down 0

Egress RSVP: 2 sessions
To          From          State  Rt Style Labelin Labelout LSPName
10.255.245.194 10.255.245.195 Up      0 1 FF    39811      - Gpro3-ba Bidir
10.255.245.194 10.255.245.195 Up      0 1 FF      3      - pro3-ba
Total 2 displayed, Up 2, Down 0

Transit RSVP: 1 sessions
To          From          State  Rt Style Labelin Labelout LSPName
10.255.245.198 10.255.245.197 Up      0 1 SE   100000      3 pro3-de
Total 1 displayed, Up 1, Down 0

```

show rsvp session statistics

```

user@host> show rsvp session statistics
Ingress RSVP: 2 sessions
To          From          State  Packets  Bytes  LSPName
10.255.245.24 10.255.245.22 Up        0        0   pro3-bd
10.255.245.24 10.255.245.22 Up    44868  2333136 pro3-bd-2
Total 2 displayed, Up 2, Down 0

Egress RSVP: 2 sessions
To          From          State  Packets  Bytes  LSPName
10.255.245.22 10.255.245.24 Up        0        0   pro3-db
10.255.245.22 10.255.245.24 Up        0        0 pro3-db-2
Total 2 displayed, Up 2, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

show rsvp session detail

```

user@host> show rsvp session detail
Ingress RSVP: 1 sessions
1.1.1.1
  From: 2.2.2.2, LSPstate: Up, ActiveRoute: 0
  LSPName: to-a, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: -, Label out: 3
  Time left: -, Since: Fri Mar 26 18:42:42 2004
  Tspec: rate 300kbps size 300kbps peak Infbps m 20 M 1500
  DiffServ info: diffServ-TE LSP, bandwidth: <ct1 300kbps>
  Port number: sender 1 receiver 15876 protocol 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  PATH sentto: 192.168.37.16 (t1-0/2/1.0) 1 pkt

```

show rsvp session detail (Path MTU Output Field)

```

user@host> show rsvp session detail
Ingress RSVP: 1 sessions
10.255.245.3
  From: 10.255.245.5, LSPstate: Up, ActiveRoute: 3
  LSPName: to-c, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 100432
  Resv style: 1 FF, Label in: -, Label out: 100432

```



```

Time left:    -, Since: Mon Aug 16 17:54:40 2006
Tspec: rate 0bps size 0bps peak Infbps m 20 M 9192
Port number: sender 1 receiver 57843 protocol 0
FastReroute desired
PATH rcvfrom: localclient
Adspec: sent MTU 4470
Path mtu: received 4470, using 4458 for forwarding
PATH sentto: 192.168.37.89 (so-0/2/3.0) 11 pkts
RESV rcvfrom: 192.168.37.89 (so-0/2/3.0) 10 pkts
Explct route: 192.168.37.89
Record route: <self> 192.168.37.89 192.168.37.87
Detour is Up
Detour Tspec: rate 0bps size 0bps peak Infbps m 20 M 9192
Detour adspec: sent MTU 1512
Path mtu: received 1512, using 1500 for forwarding

```

show rsvp session detail (GMPLS)

```

user@host> show rsvp session detail
Ingress RSVP: 1 sessions
192.168.4.1
  From: 192.168.1.1, LSPstate: Dn, ActiveRoute: 0
  LSPname: gmp1s-r1-to-r3, LSPpath: Primary
  Bidirectional, Upstream label in: 21253, Upstream label out: -
  Suggested label received: -, Suggested label sent: 21253
  Recovery label received: -, Recovery label sent: -
  Resv style: 0 -, Label in: -, Label out: -
  Time left:    -, Since: Mon Aug 16 17:54:40 2006
  Tspec: rate 0bps size 0bps peak 155.52Mbps m 20 M 1500
  Port number: sender 2 receiver 46115 protocol 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  PATH MTU: received 0
  PATH sentto: 10.35.1.5 (so-0/2/3.0) 11 pkts
  Explct route: 100.100.100.100 93.93.93.93
  Record route: <self> 100.100.100.100 93.93.93.93
  Total 1 displayed, Up 0, Down 1
Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

show rsvp session extensive

```

user@host> show rsvp session extensive
Ingress RSVP: 1 sessions
192.168.0.4
  From: 192.168.0.5, LSPstate: Up, ActiveRoute: 0
  LSPname: E-D, LSPpath: Primary
  LSPtype: Static Configured
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 299808
  Resv style: 1 FF, Label in: -, Label out: 299808
  Time left:    -, Since: Thu Sep 20 15:54:20 2012
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 2 receiver 61576 protocol 0
  Attrib flags: Non-PHP
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.0.18 (lt-1/2/0.17) 41 pkts
  RESV rcvfrom: 10.0.0.18 (lt-1/2/0.17) 40 pkts
  Explct route: 10.0.0.18 10.0.0.22

```

```
Record route: <self> 10.0.0.18 10.0.0.22
Total 1 displayed, Up 1, Down 0
```

```
Egress RSVP: 1 sessions
```

```
192.168.0.5
```

```
From: 192.168.0.4, LSPstate: Up, ActiveRoute: 0
LSPname: E-D, LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: -
Resv style: 1 FF, Label in: 3, Label out: -
Time left: 140, Since: Thu Sep 20 15:52:10 2012
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 1 receiver 49601 protocol 0
PATH rcvfrom: 10.0.0.18 (lt-1/2/0.17) 44 pkts
Adspec: received MTU 1500
PATH sentto: localclient
RESV rcvfrom: localclient
Record route: 10.0.0.22 10.0.0.18 <self>
Total 1 displayed, Up 1, Down 0
```

```
Transit RSVP: 0 sessions
```

```
Total 0 displayed, Up 0, Down 0
```

show rsvp session p2mp (Ingress Router)

```
user@host> show rsvp session p2mp
```

```
Ingress RSVP: 3 sessions
```

```
P2MP name: test, P2MP branch count: 1
```

| To | From | State | Rt | Style | Labelin | Labelout | LSPname |
|--------------|-------------|-------|----|-------|---------|----------|---------|
| 10.255.10.95 | 10.255.10.2 | Up | 0 | 1 SE | - | 3 | to-pe1 |

```
P2MP name: test2, P2MP branch count: 2
```

| To | From | State | Rt | Style | Labelin | Labelout | LSPname |
|--------------|-------------|-------|----|-------|---------|----------|---------|
| 10.255.10.23 | 10.255.10.2 | Up | 0 | 1 SE | - | 299776 | to-pe3 |
| 10.255.10.16 | 10.255.10.2 | Up | 0 | 1 SE | - | 299776 | to-pe4 |

```
Total 3 displayed, Up 3, Down 0
```

```
Egress RSVP: 0 sessions
```

```
Total 0 displayed, Up 0, Down 0
```

```
Transit RSVP: 0 sessions
```

```
Total 0 displayed, Up 0, Down 0
```

show rsvp session p2mp (Transit Router)

```
user@host> show rsvp session p2mp
```

```
Ingress RSVP: 1 sessions
```

```
P2MP name: test, P2MP branch count: 1
```

| To | From | State | Rt | Style | Labelin | Labelout | LSPname |
|--------------|--------------|-------|----|-------|---------|----------|---------|
| 10.255.10.23 | 10.255.10.95 | Up | 0 | 1 SE | - | 299792 | to-pe2 |

```
Total 1 displayed, Up 1, Down 0
```

```
Egress RSVP: 1 sessions
```

```
P2MP name: test, P2MP branch count: 1
```

| To | From | State | Rt | Style | Labelin | Labelout | LSPname |
|--------------|-------------|-------|----|-------|---------|----------|---------|
| 10.255.10.95 | 10.255.10.2 | Up | 0 | 1 SE | 3 | - | to-pe1 |

```
Total 1 displayed, Up 1, Down 0
```

```
Transit RSVP: 2 sessions
```

```
P2MP name: test2, P2MP branch count: 2
```

| To | From | State | Rt | Style | Labelin | Labelout | LSPname |
|--------------|-------------|-------|----|-------|---------|----------|---------|
| 10.255.10.23 | 10.255.10.2 | Up | 0 | 1 SE | 299776 | 299808 | to-pe3 |
| 10.255.10.16 | 10.255.10.2 | Up | 0 | 1 SE | 299776 | 299856 | to-pe4 |

```
Total 2 displayed, Up 2, Down 0
```


show rsvp statistics

| | |
|------------------------------------|---|
| Syntax | show rsvp statistics <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switches) | show rsvp statistics |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches. |
| Description | Display Resource Reservation Protocol (RSVP) packet and error statistics. |
| Options | none —Display RSVP packet and error statistics. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> clear rsvp statistics on page 868 |
| List of Sample Output | show rsvp statistics on page 893 |
| Output Fields | Table 209 on page 890 describes the output fields for the show rsvp statistics command. Output fields are listed in the approximate order in which they appear. |

Table 209: show rsvp statistics Output Fields

| Field Name | Field Description |
|--------------------------------|--|
| Packet Type | Statistics about different RSVP messages. |
| Total Sent | Total number of packets sent since RSVP was enabled. |
| Total Received | Total number of packets received since RSVP was enabled. |
| Last 5 seconds Sent | Total number of packets sent in the last 5 seconds. |
| Last 5 seconds Received | Number of packets received in the last 5 seconds. |
| Path | Statistics about Path messages, which are sent from the RSVP sender along the data paths and which store path state information in each node along the path. |
| PathErr | Statistics about PathErr messages, which are advisory messages that are sent upstream to the sender. |
| PathTear | Statistics about PathTear messages, which remove path states and dependent reservation states in any routing devices along a path. |

Table 209: show rsvp statistics Output Fields (*continued*)

| Field Name | Field Description |
|-------------------------------|--|
| Resv FF | Statistics about fixed-filter reservation style messages, which consist of distinct reservations among explicit senders. |
| Resv WF | Statistics about wildcard-filter reservation style messages, which consist of shared reservations among wildcard senders. |
| Res SE | Statistics about shared-explicit reservation style messages, which consist of shared reservations among explicit senders. |
| ResvErr | Statistics about ResvErr messages, which are advisory messages that are sent when an attempt to establish a reservation fails. |
| ResvTear | Statistics about ResvTear messages, which remove reservation states along a path. |
| ResvConf | Statistics about ResvConfirm messages, which are responses to confirm a reservation request. |
| Ack | Acknowledge message for refresh reductions. |
| SRefresh | Summary refresh messages. |
| Hello | Number of RSVP hello packets that have been sent to and received from the neighbor. |
| EndtoEnd RSVP | Statistics for the number of End-to-end RSVP messages. |
| Errors | Statistics about errored RSVP packets. |
| Rcv pkt bad length | The packet was not processed because its length is inappropriate. |
| Rcv pkt unknown type | The packet is not one of the well-known RSVP types, as defined in RFC 2205, <i>Resource ReSerVation Protocol (RSVP)</i> . |
| Rcv pkt bad version | The packet is not an RSVP version 1 packet. |
| Rcv pkt auth fail | The packet failed authentication checks. |
| Rcv pkt bad checksum | The RSVP checksum check failed. |
| Rcv pkt bad format | General packet processing failed because the packet was badly formed. |
| Memory allocation fail | An internal resource failure occurred. |
| No path information | A reservation was received, but no sender is active. |
| Resv style conflict | The same session contains inconsistent reservation styles. |
| Port conflict | There were inconsistent port numbers for the same session. |
| Resv no interface | An interface for the receive reservation packets cannot be located. |

Table 209: show rsvp statistics Output Fields (*continued*)

| Field Name | Field Description |
|------------------------------------|---|
| PathErr to client | Number of PathErr packets delivered to the local client. |
| ResvErr to client | Number of ResvErr packets delivered to the local client. |
| Path timeout | Number of times the sender timed out because the path was removed. |
| Resv timeout | Number of times the receiver timed out because the reservation was removed. |
| Message out-of-order | Records the number of RSVP incoming messages that are considered out of order. This is detected from the message ID object's sequence number. |
| Unknown ack msg | A neighboring routing device replies with an ACK object that contains an unknown message ID. This can indicate a message ID handshake problem. For example, a router receives an ACK for message IDs 1, 2, and 3. However, it only has state for message IDs 1 and 3. The router increments the unknown ack counter by 1. |
| Recv nack | If a neighboring router receives an unknown message ID in an RSVP refresh message, the router sends a Resv nack message back to the sender. This can happen if that neighbor has been rebooted. For this case, the router sends a regular RSVP refresh message to recover the state and start the message-ID handshake process again. |
| Recv duplicated msg-id | Number of times the same message ID is used by two different RSVP messages. This duplication is usually caused when a neighboring routing device restarts. |
| No TE-link to recv Hop | Counter of packets discarded because a TE link was not found. |
| Rcv pkt disabled interface | Number of RSVP packets received on an interface that is not enabled for RSVP. |
| Transmit buffer full | Number of times the buffer for assembling an outgoing RSVP message was not large enough. |
| Transmit failure | Number of times the RSVP task failed to send out a packet. |
| Receive failure | Number of times the RSVP task failed to read an incoming packet. |
| P2MP RESV discarded by appl | Number of Resv messages discarded because the MPLS label is not valid for the P2MP LSP application. |
| Rate limit | Number of RSVP packets dropped due to rate limiting. |
| Err msg loop detected | Number of RSVP error messages that have looped back to their originator. This is detected by checking the error node address in the ERROR_SPEC object. |

Sample Output

show rsvp statistics

```

user@host> show rsvp statistics

```

| PacketType | Total | | Last 5 seconds | |
|---------------|--------|----------|----------------|----------|
| | Sent | Received | Sent | Received |
| Path | 355 | 408 | 0 | 0 |
| PathErr | 2 | 13 | 0 | 0 |
| PathTear | 101 | 139 | 0 | 0 |
| Resv FF | 0 | 0 | 0 | 0 |
| Resv WF | 0 | 0 | 0 | 0 |
| Resv SE | 419 | 225 | 0 | 0 |
| ResvErr | 0 | 0 | 0 | 0 |
| ResvTear | 0 | 13 | 0 | 0 |
| ResvConf | 0 | 0 | 0 | 0 |
| Ack | 682 | 1414 | 0 | 0 |
| SRefresh | 395198 | 236030 | 5 | 2 |
| Hello | 578809 | 578221 | 4 | 4 |
| EndtoEnd RSVP | 0 | 0 | 0 | 0 |

| Errors | Total | Last 5 seconds |
|-----------------------------|-------|----------------|
| Rcv pkt bad length | 0 | 0 |
| Rcv pkt unknown type | 0 | 0 |
| Rcv pkt bad version | 0 | 0 |
| Rcv pkt auth fail | 0 | 0 |
| Rcv pkt bad checksum | 0 | 0 |
| Rcv pkt bad format | 0 | 0 |
| Memory allocation fail | 0 | 0 |
| No path information | 10 | 0 |
| Resv style conflict | 0 | 0 |
| Port conflict | 0 | 0 |
| Resv no interface | 0 | 0 |
| PathErr to client | 38 | 0 |
| ResvErr to client | 0 | 0 |
| Path timeout | 8 | 0 |
| Resv timeout | 57 | 0 |
| Message out-of-order | 0 | 0 |
| Unknown ack msg | 2978 | 0 |
| Recv nack | 86 | 0 |
| Recv duplicated msg-id | 5 | 0 |
| No TE-link to recv Hop | 0 | 0 |
| Rcv pkt disabled interface | 0 | 0 |
| Transmit buffer full | 0 | 0 |
| Transmit failure | 0 | 0 |
| Receive failure | 0 | 0 |
| P2MP RESV discarded by appl | 0 | 0 |
| Rate limit | 306 | 0 |
| Err msg loop detected | 0 | 0 |

show rsvp version

| | |
|------------------------------------|--|
| Syntax | show rsvp version <logical-system (all <i>logical-system-name</i>)> |
| Syntax (EX Series Switches) | show rsvp version |
| Release Information | Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches. |
| Description | Display information about the Resource Reservation Protocol (RSVP) protocol settings, such as the version of the RSVP software, the refresh timer and keep multiplier, and local RSVP graceful restart capabilities on a routing device. |
| Options | none —Display RSVP protocol settings. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | view |
| List of Sample Output | show rsvp version on page 896 |
| Output Fields | Table 210 on page 894 describes the output fields for the show rsvp version command. Output fields are listed in the approximate order in which they appear. |

Table 210: show rsvp version Output Fields

| Field Name | Field Description |
|--|---|
| Resource ReSerVation Protocol, version | RSVP software version. |
| RSVP protocol | Status of RSVP: Enabled or Disabled . |
| R(refresh timer) | Configured time interval used to generate periodic RSVP messages. |
| K(keep multiplier) | Number of RSVP messages that can be lost before an RSVP state is declared stale. |
| Preemption | Currently configured preemption capability: Aggressive , Disabled , or Normal . The default is Normal . |
| Soft-preemption cleanup | Time, in seconds, that an LSP is kept after it has been soft preempted. This is a global property of the RSVP protocol. |
| Graceful deleting timeout | Currently configured value for the graceful-deletion-timeout statement. The router that initiates the graceful deletion procedure for an RSVP session waits for the graceful deletion timeout interval to ensure that all routers along the path (especially the ingress and egress routers) have prepared for the LSP to be taken down. |

Table 210: show rsvp version Output Fields (*continued*)

| Field Name | Field Description |
|--------------------------------------|--|
| NSR Mode | Status of the nonstop active routing feature for RSVP on the restarting device: Disabled , Enabled/Master , or Enabled/Standby . |
| NSR State | <p>State of the nonstop active routing feature for RSVP on the restarting device.</p> <p>Possible values are:</p> <ul style="list-style-type: none"> • Idle • TE-link sync complete • Neighbor sync complete • Path state sync complete • Resv state sync complete • Bypass sync complete • Init sync complete |
| Setup protection | Status of point-to-point and point-to-multipoint LSP setup protection configuration on the device: Enabled or Disabled |
| Graceful restart | Status of the graceful restart feature for RSVP on the restarting routing device: Enabled or Disabled . |
| Restart helper mode | Status of the helper mode feature: Enabled or Disabled . When this feature is enabled, the restarting routing device can help the neighbor with its RSVP restart procedures. |
| Maximum helper restart time | Number of milliseconds (ms) configured for the maximum helper restart time. The maximum helper restart time is the length of time the routing device waits before declaring that an RSVP neighbor attempting to restart gracefully is down. |
| Maximum helper recovery time | Number of milliseconds configured for the maximum helper recovery time. The maximum helper recovery time is the amount of time the routing device maintains the state of an RSVP neighbor attempting to restart gracefully. |
| Restart time | Number of milliseconds that a neighbor waits to receive a Hello message from the restarting node before declaring the node dead and deleting the states. |
| Recovery time | Number of milliseconds during which the restarting node attempts to recover its lost states with help from its neighbors. Recovery time is advertised by the restarting node to its neighbors, and applies to nodal faults. The restarting node considers its graceful restart complete after this time has elapsed. |
| P2p transit LSP nexthop mode | Point-to-point transit LSP nexthop mode on PTX Series devices. The possible values are Chained or Unchained |
| P2mp transit LSP nexthop mode | Point-to-multipoint transit LSP nexthop mode on PTX Series devices. The possible values are Chained or Unchained |

Sample Output

show rsvp version

```
user@host> show rsvp version
Resource ReSerVation Protocol, version 1. rfc2205
  RSVP protocol:           Enabled
  R(refresh timer):        30 seconds
  K(keep multiplier):      3
  Preemption:              Normal
  Soft-preemption cleanup:  30 seconds
  Graceful deletion timeout: 30 seconds
  NSR mode:                Enabled/Master
  NSR state:                Init sync complete
  Setup protection:        Disabled
  Graceful restart:        Disabled
  Restart helper mode:      Enabled
  Maximum helper restart time: 20000 msec
  Maximum helper recovery time: 180000 msec
  Restart time:             0 msec
  P2p transit LSP nexthop mode: Unchained
  P2mp transit LSP nexthop mode: Unchained
```

PART 4

Layer 2 Bridging and Switching Operational Mode Commands

- [Layer 2 Bridging and Switching Operational Mode Commands on page 899](#)
- [Spanning Tree Operational Mode Commands on page 927](#)

Layer 2 Bridging and Switching

Operational Mode Commands

Table 211 on page 899 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot Layer 2 bridging and switching. Commands are listed in alphabetical order.

Table 211: Layer 2 Bridging and Switching Operational Mode Commands

| Task | Command |
|--|--|
| Clear learned Layer 2 address information from the media access control (MAC) address table. | <code>clear bridge mac-table</code> |
| Clear bridge protocol data unit (BPDU) error on interface due to possible bridge spanning tree protocol (STP) loop. | <code>clear error bpdu</code> |
| Clear a MAC rewrite error condition for Layer 2 protocol tunneling. | <code>clear error mac-rewrite</code> |
| Display bridge domain information. | <code>show bridge domain</code> |
| Display bridging flooding information. | <code>show bridge flood</code> |
| Display learned Layer 2 MAC address information. | <code>show bridge mac-table</code> |
| Display bridge statistics. | <code>show bridge statistics</code> |
| Display Layer 2 learning process-related information. | <code>show l2-learning global-information</code> |
| (MX Series routers only) Display the total number of dynamic and static MAC addresses learned for the entire router. | <code>show l2-learning global-mac-count</code> |
| Display configured Layer 2 routing instances. | <code>show l2-learning instance</code> |
| Display configured Layer 2 interfaces. | <code>show l2-learning interface</code> |
| Display Layer 2 interfaces. | <code>show mac-rewrite interface</code> |

clear bridge mac-table

| | |
|---------------------------------|--|
| Syntax | <code>clear bridge mac-table <bridge-domain (all <i>bridge-domain-name</i>)> <instance <i>instance-name</i>> <interface <i>interface-name</i>> <learning-vlan id (all-vlan <i>learning-vlan-id</i>)> <<i>mac-address</i>></code> |
| Release Information | Command introduced in Junos OS Release 8.4. |
| Description | (MX Series routers only) Clear learned Layer 2 address information from the media access control (MAC) address table. |
| Options | <p>none—Clear all learned Layer 2 address information from the MAC address table.</p> <p>bridge-domain (all <i>bridge-domain-name</i>)—(Optional) Clear learned Layer 2 MAC addresses for all bridging domains or for the specified bridging domain.</p> <p>instance <i>instance-name</i>—(Optional) Clear learned Layer 2 MAC addresses for the specified routing instance.</p> <p>interface <i>interface-name</i>—(Optional) Clear learned Layer 2 MAC addresses for the specified interface.</p> <p>learning-vlan-id (all-vlan <i>learning-vlan-id</i>)—(Optional) Clears learned Layer 2 MAC addresses for all VLANs or for the specified VLAN.</p> <p><i>mac-address</i>—(Optional) Clear the specified learned Layer 2 address from the MAC address table.</p> |
| Required Privilege Level | clear |
| List of Sample Output | clear bridge mac-table on page 900 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear bridge mac-table user@host> clear bridge mac-table

clear error bpdu

| | |
|---------------------------------|--|
| Syntax | <code>clear error bpdu</code> <code><interface <i>interface-name</i>></code> |
| Release Information | Command introduced in Junos OS Release 9.4. |
| Description | (MX Series routers only) Clear a bridge protocol data unit (BPDU) error condition caused by the detection of a possible bridging loop from Spanning Tree Protocol (STP) operation. |
| Options | <code>interface <i>interface-name</i></code> —(Optional) Clear the BPDU error condition for the specified interface. |
| Required Privilege Level | clear |
| List of Sample Output | clear error bpdu interface on page 901 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

| | |
|---|--|
| <code>clear error bpdu</code> <code>interface</code> | <code>user@host> clear error bpdu interface ge-1/1/1</code> |
|---|--|

clear error mac-rewrite

| | |
|---------------------------------|--|
| Syntax | <code>clear error mac-rewrite</code> <code><interface <i>interface-name</i>></code> |
| Release Information | Command introduced in Junos OS Release 9.1. |
| Description | (MX Series routers only) Clear a MAC rewrite error condition caused by the reception of tunneled Cisco Discovery Protocol (CDP), Spanning Tree Protocol (STP), or VLAN Trunk Protocol (VTP) packets on an interface with Layer 2 protocol tunneling enabled. |
| Options | <code>interface <i>interface-name</i></code> —(Optional) Clear the MAC rewrite error condition for the specified interface. |
| Required Privilege Level | clear |
| List of Sample Output | clear error mac-rewrite interface on page 902 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

```
clear error mac-rewrite interface user@host> clear error mac-rewrite interface ge-1/0/1
```


show bridge domain

| | |
|---------------------------------|--|
| Syntax | <code>show bridge domain</code> <code><brief detail extensive></code> <code><bridge-domain (all <i>domain-name</i>)></code> <code><instance <i>instance-name</i>></code> <code><operational></code> |
| Release Information | Command introduced in Junos OS Release 8.4. |
| Description | (MX Series routers only) Display bridge domain information. |
| Options | <p>none—Display information for all bridge domains.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>bridge-domain (all <i>domain-name</i>)— (Optional) Display information about all bridge domains or the specified bridge domain.</p> <p>instance <i>instance-name</i>—(Optional) Display information for the specified routing instance.</p> <p>operational—(Optional) Display information for the operational routing instances.</p> |
| Required Privilege Level | view |
| List of Sample Output | show bridge domain on page 904 show bridge domain brief on page 904 show bridge domain detail on page 904 |

Sample Output

show bridge domain

```
user@host> show bridge domain
Instance      Primary Table  Bridging Domain  Type      Active
vs1           bridge.0      vlan100          bridge     2
vs1           bridge.0      vlan200          bridge     0
```

show bridge domain brief

```
user@host> show bridge domain brief
Instance      Primary Table  Bridging Domain  Type      Active
vs1           bridge.0      vlan100          bridge     2
vs1           bridge.0      vlan200          bridge     0
```

show bridge domain detail

```
user@host> show bridge domain detail
Routing Instance: vs1
  Bridging Domain: vlan100
  Router ID: 0.0.0.0
  Type: bridge          State: Active
  Interfaces:
    ge-11/0/3.0
    ge-11/1/4.100
    ge-11/1/1.100
    ge-11/1/0.100
    xe-10/2/0.100
    xe-10/0/0.100
  Tables:
    bridge.0            : 2 macs (2 active)
Routing Instance: vs1
  Bridging Domain: vlan200
  Router ID: 0.0.0.0
  Type: bridge          State: Active
  Interfaces:
    ge-11/1/0.200
    ge-11/1/1.200
    ge-11/1/4.200
    xe-10/0/0.200
    xe-10/2/0.200
  Tables:
    bridge.0            : 0 macs (0 active)
```

show bridge flood

| | |
|---------------------------------|--|
| Syntax | <pre>show bridge flood <brief detail extensive> <bridge-domain <i>domain-name</i>> <event-queue> <instance <i>instance-name</i>> <route (all-ce-flood all ve-flood alt-root-flood bd-flood mlp-flood re-flood)></pre> |
| Release Information | Command introduced in Junos OS Release 8.4. |
| Description | (MX Series routers only) Display bridging flooding information. |
| Options | <p>none—Display all bridging flooding information for all bridging domains.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>bridge-domain <i>domain-name</i>—(Optional) Display bridging flooding information for the specified bridge domain.</p> <p>event-queue—(Optional) Display the queue of pending bridge flood events.</p> <p>instance <i>instance-name</i>—(Optional) Display bridging flooding information for the specified routing instance.</p> <p>route (all-ce-flood all ve-flood alt-root-flood bd-flood mlp-flood re-flood)—(Optional) Display the following:</p> <ul style="list-style-type: none"> all-ce-flood—Display the route for flooding traffic to all customer edge routers if no-local-switching is enabled. all-ve-flood—Display the route for flooding traffic to all VPLS edge routers if no-local-switching is enabled. alt-root-flood—Display the Spanning Tree Protocol (STP) alt-root flooding route used for the interface. bd-flood—Display the route for flooding traffic of a bridge domain if no-local-switching is not enabled. mlp-flood—Display the route for flooding traffic to MAC learning chips. re-flood—Display the route for Routing Engine flooding to all interfaces. |
| Required Privilege Level | view |
| List of Sample Output | show bridge flood on page 906 show bridge flood brief on page 906 show bridge flood detail on page 906 show bridge flood extensive on page 907 |
| Output Fields | to be provided |

Sample Output

show bridge flood

```

user@host> show bridge flood
Name: __juniper_private1__
CEs: 0
VEs: 0
Flood Routes:
  Prefix  Type      Owner      NhType      NhIndex
  0x36/16  MLP_FLOOD  __vs1+vlan100__  flood      426
  0x3a/16  MLP_FLOOD  __vs1+vlan200__  flood      428
Name: vs1::vlan100
CEs: 6
VEs: 0
Flood Routes:
  Prefix  Type      Owner      NhType      NhIndex
  0x35/16  ALL_FLOOD  __vs1+vlan100__  flood      425
  0x35/16  RE_FLOOD   __vs1+vlan100__  flood      425
  0x3780/17 ALT_ROOT_RT ge-11/0/3.0      flood      425
  0x3b80/17 ALT_ROOT_RT ge-11/1/4.100    flood      425
  0x3c80/17 ALT_ROOT_RT ge-11/1/1.100    flood      425
  0x3d80/17 ALT_ROOT_RT ge-11/1/0.100    flood      425
  0x3e80/17 ALT_ROOT_RT xe-10/2/0.100    flood      425
  0x3f80/17 ALT_ROOT_RT xe-10/0/0.100    flood      425
Name: vs1::vlan200
CEs: 5
VEs: 0
Flood Routes:
  Prefix  Type      Owner      NhType      NhIndex
  0x39/16  ALL_FLOOD  __vs1+vlan200__  flood      427
  0x39/16  RE_FLOOD   __vs1+vlan200__  flood      427
  0x4180/17 ALT_ROOT_RT ge-11/1/0.200    flood      427
  0x4080/17 ALT_ROOT_RT ge-11/1/1.200    flood      427
  0x4280/17 ALT_ROOT_RT ge-11/1/4.200    flood      427
  0x4480/17 ALT_ROOT_RT xe-10/0/0.200    flood      427
  0x4380/17 ALT_ROOT_RT xe-10/2/0.200    flood      427

```

show bridge flood brief

```

user@host> show bridge flood brief
Name      Active CEs      Active VEs
__juniper_private1__  0                0
vs1::vlan100          6                0
vs1::vlan200          5                0

```

show bridge flood detail

```

user@host> show bridge flood detail
Name: __juniper_private1__
CEs: 0
VEs: 0
Flood Routes:
  Prefix  Type      Owner      NhType      NhIndex
  0x36/16  MLP_FLOOD  __vs1+vlan100__  flood      426
  0x3a/16  MLP_FLOOD  __vs1+vlan200__  flood      428
Name: vs1::vlan100
CEs: 6
VEs: 0
Flood Routes:
  Prefix  Type      Owner      NhType      NhIndex
  0x35/16  ALL_FLOOD  __vs1+vlan100__  flood      425
  0x35/16  RE_FLOOD   __vs1+vlan100__  flood      425
  0x3780/17 ALT_ROOT_RT ge-11/0/3.0      flood      425

```

```

0x3b80/17 ALT_ROOT_RT ge-11/1/4.100 flood 425
0x3c80/17 ALT_ROOT_RT ge-11/1/1.100 flood 425
0x3d80/17 ALT_ROOT_RT ge-11/1/0.100 flood 425
0x3e80/17 ALT_ROOT_RT xe-10/2/0.100 flood 425
0x3f80/17 ALT_ROOT_RT xe-10/0/0.100 flood 425
Name: vs1::vlan200
CEs: 5
VEs: 0
Flood Routes:
  Prefix    Type      Owner                NhType    NhIndex
  0x39/16   ALL_FLOOD __vs1+vlan200__     flood     427
  0x39/16   RE_FLOOD  __vs1+vlan200__     flood     427
  0x4180/17 ALT_ROOT_RT ge-11/1/0.200     flood     427
  0x4080/17 ALT_ROOT_RT ge-11/1/1.200     flood     427
  0x4280/17 ALT_ROOT_RT ge-11/1/4.200     flood     427
  0x4480/17 ALT_ROOT_RT xe-10/0/0.200     flood     427
  0x4380/17 ALT_ROOT_RT xe-10/2/0.200     flood     427

```

show bridge flood extensive

```

user@host> show bridge flood extensive
Name: __juniper_private1__
CEs: 0
VEs: 0
Flood route prefix: 0x36/16
Flood route type: MLP_FLOOD
Flood route owner: __vs1+vlan100__
Nexthop type: flood
Nexthop index: 426
  Interfaces Flooding to:
    Name                Type      NhType    Index
    1c-11/0/0.32769     LC
    1c-10/2/0.32769     LC
    1c-10/0/0.32769     LC
    1c-11/1/0.32769     LC

Flood route prefix: 0x3a/16
Flood route type: MLP_FLOOD
Flood route owner: __vs1+vlan200__
Nexthop type: flood
Nexthop index: 428
  Interfaces Flooding to:
    Name                Type      NhType    Index
    1c-10/0/0.32769     LC
    1c-10/2/0.32769     LC
    1c-11/1/0.32769     LC

Name: vs1::vlan100
CEs: 6
VEs: 0

Flood route prefix: 0x35/16
Flood route type: ALL_FLOOD
Flood route owner: __vs1+vlan100__
Nexthop type: flood
Nexthop index: 425
  Interfaces Flooding to:
    Name                Type      NhType    Index
    ge-11/0/3.0         CE
    ge-11/1/4.100       CE
    ge-11/1/1.100       CE
    ge-11/1/0.100       CE
    xe-10/2/0.100       CE
    xe-10/0/0.100       CE

```

Flood route prefix: 0x35/16
 Flood route type: RE_FLOOD
 Flood route owner: __vs1+vlan100__
 Nexthop type: flood
 Nexthop index: 425

Interfaces Flooding to:

| Name | Type | NhType | Index |
|---------------|------|--------|-------|
| ge-11/0/3.0 | CE | | |
| ge-11/1/4.100 | CE | | |
| ge-11/1/1.100 | CE | | |
| ge-11/1/0.100 | CE | | |
| xe-10/2/0.100 | CE | | |
| xe-10/0/0.100 | CE | | |

Flood route prefix: 0x3780/17
 Flood route type: ALT_ROOT_RT
 Flood route owner: ge-11/0/3.0
 Nexthop type: flood
 Nexthop index: 425

Interfaces Flooding to:

| Name | Type | NhType | Index |
|---------------|------|--------|-------|
| ge-11/0/3.0 | CE | | |
| ge-11/1/4.100 | CE | | |
| ge-11/1/1.100 | CE | | |
| ge-11/1/0.100 | CE | | |
| xe-10/2/0.100 | CE | | |
| xe-10/0/0.100 | CE | | |

Flood route prefix: 0x3b80/17
 Flood route type: ALT_ROOT_RT
 Flood route owner: ge-11/1/4.100
 Nexthop type: flood
 Nexthop index: 425

Interfaces Flooding to:

| Name | Type | NhType | Index |
|---------------|------|--------|-------|
| ge-11/0/3.0 | CE | | |
| ge-11/1/4.100 | CE | | |
| ge-11/1/1.100 | CE | | |
| ge-11/1/0.100 | CE | | |
| xe-10/2/0.100 | CE | | |
| xe-10/0/0.100 | CE | | |

Flood route prefix: 0x3c80/17
 Flood route type: ALT_ROOT_RT
 Flood route owner: ge-11/1/1.100
 Nexthop type: flood
 Nexthop index: 425

Interfaces Flooding to:

| Name | Type | NhType | Index |
|---------------|------|--------|-------|
| ge-11/0/3.0 | CE | | |
| ge-11/1/4.100 | CE | | |
| ge-11/1/1.100 | CE | | |
| ge-11/1/0.100 | CE | | |
| xe-10/2/0.100 | CE | | |
| xe-10/0/0.100 | CE | | |

Flood route prefix: 0x3d80/17
 Flood route type: ALT_ROOT_RT
 Flood route owner: ge-11/1/0.100
 Nexthop type: flood

```

Nexthop index: 425
  Interfaces Flooding to:
    Name          Type          NhType          Index
    ge-11/0/3.0    CE
    ge-11/1/4.100  CE
    ge-11/1/1.100  CE
    ge-11/1/0.100  CE
    xe-10/2/0.100  CE
    xe-10/0/0.100  CE

```

```

Flood route prefix: 0x3e80/17
Flood route type: ALT_ROOT_RT
Flood route owner: xe-10/2/0.100
Nexthop type: flood
Nexthop index: 425

```

```

  Interfaces Flooding to:
    Name          Type          NhType          Index
    ge-11/0/3.0    CE
    ge-11/1/4.100  CE
    ge-11/1/1.100  CE
    ge-11/1/0.100  CE
    xe-10/2/0.100  CE
    xe-10/0/0.100  CE

```

```

Flood route prefix: 0x3f80/17
Flood route type: ALT_ROOT_RT
Flood route owner: xe-10/0/0.100
Nexthop type: flood
Nexthop index: 425

```

```

  Interfaces Flooding to:
    Name          Type          NhType          Index
    ge-11/0/3.0    CE
    ge-11/1/4.100  CE
    ge-11/1/1.100  CE
    ge-11/1/0.100  CE
    xe-10/2/0.100  CE
    xe-10/0/0.100  CE

```

```

Name: vs1::vlan200
CEs: 5
VEs: 0

```

```

Flood route prefix: 0x39/16
Flood route type: ALL_FLOOD
Flood route owner: __vs1+vlan200__
Nexthop type: flood
Nexthop index: 427

```

```

  Interfaces Flooding to:
    Name          Type          NhType          Index
    ge-11/1/0.200  CE
    ge-11/1/1.200  CE
    ge-11/1/4.200  CE
    xe-10/0/0.200  CE
    xe-10/2/0.200  CE

```

```

Flood route prefix: 0x39/16
Flood route type: RE_FLOOD
Flood route owner: __vs1+vlan200__
Nexthop type: flood
Nexthop index: 427

```

```

  Interfaces Flooding to:
    Name          Type          NhType          Index

```

```

ge-11/1/0.200    CE
ge-11/1/1.200    CE
ge-11/1/4.200    CE
xe-10/0/0.200    CE
xe-10/2/0.200    CE

```

```

Flood route prefix: 0x4180/17
Flood route type: ALT_ROOT_RT
Flood route owner: ge-11/1/0.200
Nexthop type: flood
Nexthop index: 427

```

```

  Interfaces Flooding to:
  Name                Type                NhType                Index
  ge-11/1/0.200        CE
  ge-11/1/1.200        CE
  ge-11/1/4.200        CE
  xe-10/0/0.200        CE
  xe-10/2/0.200        CE

```

```

Flood route prefix: 0x4080/17
Flood route type: ALT_ROOT_RT
Flood route owner: ge-11/1/1.200
Nexthop type: flood
Nexthop index: 427

```

```

  Interfaces Flooding to:
  Name                Type                NhType                Index
  ge-11/1/0.200        CE
  ge-11/1/1.200        CE
  ge-11/1/4.200        CE
  xe-10/0/0.200        CE
  xe-10/2/0.200        CE

```

```

Flood route prefix: 0x4280/17
Flood route type: ALT_ROOT_RT
Flood route owner: ge-11/1/4.200
Nexthop type: flood
Nexthop index: 427

```

```

  Interfaces Flooding to:
  Name                Type                NhType                Index
  ge-11/1/0.200        CE
  ge-11/1/1.200        CE
  ge-11/1/4.200        CE
  xe-10/0/0.200        CE
  xe-10/2/0.200        CE

```

```

Flood route prefix: 0x4480/17
Flood route type: ALT_ROOT_RT
Flood route owner: xe-10/0/0.200
Nexthop type: flood
Nexthop index: 427

```

```

  Interfaces Flooding to:
  Name                Type                NhType                Index
  ge-11/1/0.200        CE
  ge-11/1/1.200        CE
  ge-11/1/4.200        CE
  xe-10/0/0.200        CE
  xe-10/2/0.200        CE

```

```

Flood route prefix: 0x4380/17
Flood route type: ALT_ROOT_RT
Flood route owner: xe-10/2/0.200

```


Nexthop type: flood

Nexthop index: 427

Interfaces Flooding to:

| Name | Type | NhType | Index |
|---------------|------|--------|-------|
| ge-11/1/0.200 | CE | | |
| ge-11/1/1.200 | CE | | |
| ge-11/1/4.200 | CE | | |
| xe-10/0/0.200 | CE | | |
| xe-10/2/0.200 | CE | | |

show bridge mac-table

| | |
|---------------------------------|--|
| Syntax | <code>show bridge mac-table</code> <code><brief count detail extensive></code> <code><bridge-domain (all <i>bridge-domain-name</i>)></code> <code><global-count></code> <code><interface <i>interface-name</i>></code> <code><mac-address></code> <code><vlan-id (all-vlan <i>vlan-id</i>)></code> |
| Release Information | Command introduced in Junos OS Release 8.4. |
| Description | (MX Series routers only) Display Layer 2 MAC address information. |
| Options | <p>none—Display all learned Layer 2 MAC address information.</p> <p>brief count detail extensive—(Optional) Display the specified level of output.</p> <p>bridge-domain (all <i>bridge-domain-name</i>)—(Optional) Display learned Layer 2 MAC addresses for all bridging domains or for the specified bridging domain.</p> <p>global-count—(Optional) Display the total number of learned Layer 2 MAC addresses on the system.</p> <p>instance <i>instance-name</i>—(Optional) Display learned Layer 2 MAC addresses for the specified routing instance.</p> <p>interface <i>interface-name</i>—(Optional) Display learned Layer 2 MAC addresses for the specified interface.</p> <p>mac-address—(Optional) Display the specified learned Layer 2 MAC address information.</p> <p>vlan-id (all-vlan <i>vlan-id</i>)—(Optional) Display learned Layer 2 MAC addresses for all VLANs or for the specified VLAN.</p> |
| Additional Information | When Layer 2 protocol tunneling is enabled, the tunneling MAC address 01:00:0c:cd:cd:d0 is installed in the MAC table. When the Cisco Discovery Protocol (CDP), Spanning Tree Protocol (STP), or VLAN Trunk Protocol (VTP) is configured for Layer 2 protocol tunneling on an interface, the corresponding protocol MAC address is installed in the MAC table. |
| Required Privilege Level | view |
| List of Sample Output | show bridge mac-table on page 914 show bridge mac-table brief on page 914 show brief mac-table count on page 914 show bridge mac-table detail on page 914 |
| Output Fields | Table 212 on page 913 describes the output fields for the show bridge mac-table command. Output fields are listed in the approximate order in which they appear. |

Table 212: show bridge mac-table Output fields

| Field Name | Field Description |
|---------------------------|--|
| Routing instance | Name of the routing instance. |
| Bridging domain | Name of the bridging domain. |
| MAC address | MAC address or addresses learned on a logical interface. |
| MAC flags | Status of MAC address learning properties for each interface: <ul style="list-style-type: none"> • S—Static MAC address is configured. • D—Dynamic MAC address is configured. • SE—MAC accounting is enabled. • NM—Non-configured MAC. |
| Logical interface | Name of the logical interface. |
| MAC count | Number of MAC addresses learned on the specific routing instance or interface. |
| Learning interface | Name of the logical interface on which the MAC address was learned. |
| Learning VLAN | VLAN ID of the routing instance or bridge domain in which the MAC address was learned. |
| Layer 2 flags | Debugging flags signifying that the MAC address is present in various lists. |
| Epoch | Spanning Tree Protocol epoch number identifying when the MAC address was learned. Used for debugging. |
| Sequence number | Sequence number assigned to this MAC address. Used for debugging. |
| Learning mask | Mask of the Packet Forwarding Engines where this MAC address was learned. Used for debugging. |
| IPC generation | Creation time of the logical interface when this MAC address was learned. Used for debugging. |

Sample Output

```

user@host> show bridge mac-table
MAC flags (S -static MAC, D -dynamic MAC,
           SE -Statistics enabled, NM -Non configured MAC)

Routing instance : vs1
Bridging domain : vlan100, VLAN : 100
  Learning  MAC          MAC          Logical
  VLAN      address      flags      interface
           00:00:00:19:1c:db  D        ge-11/0/3.0
           00:00:00:59:3a:2f  D        xe-10/2/0.100

```

```

user@host> show bridge mac-table brief
MAC flags (S -static MAC, D -dynamic MAC,
           SE -Statistics enabled, NM -Non configured MAC)

Routing instance : vs1
Bridging domain : vlan100, VLAN : 100
  Learning  MAC          MAC          Logical
  VLAN      address      flags      interface
           00:00:00:19:1c:db  D        ge-11/0/3.0
           00:00:00:59:3a:2f  D        xe-10/2/0.100

```

```

user@host> show bridge mac-table count
2 MAC address learned in routing instance vs1 bridge domain vlan100

MAC address count per interface within routing instance:
  Logical interface      MAC count
  ge-11/0/3.0            1
  ge-11/1/4.100          0
  ge-11/1/1.100          0
  ge-11/1/0.100          0
  xe-10/2/0.100          1
  xe-10/0/0.100          0

MAC address count per learn VLAN within routing instance:
  Learn VLAN ID          MAC count
  0                      2

0 MAC address learned in routing instance vs1 bridge domain vlan200

MAC address count per interface within routing instance:
  Logical interface      MAC count
  ge-11/1/0.200          0
  ge-11/1/1.200          0
  ge-11/1/4.200          0
  xe-10/0/0.200          0
  xe-10/2/0.200          0

MAC address count per learn VLAN within routing instance:
  Learn VLAN ID          MAC count
  0                      0

```

```

user@host> show bridge mac-table detail
MAC address: 00:00:00:19:1c:db
Routing instance: vs1

```

```
Bridging domain: vlan100
Learning interface: ge-11/0/3.0    Learning VLAN: 0
Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
Epoch: 4                          Sequence number: 0
Learning mask: 0x800               IPC generation: 0

MAC address: 00:00:00:59:3a:2f
Routing instance: vs1
Bridging domain: vlan100
Learning interface: xe-10/2/0.100  Learning VLAN: 0
Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
Epoch: 7                          Sequence number: 0
Learning mask: 0x400               IPC generation: 0
```

show bridge statistics

| | |
|---------------------------------|---|
| Syntax | <code>show bridge statistics</code> <code><bridge-domain <i>domain-name</i>></code> <code><instance <i>instance-name</i>></code> |
| Release Information | Command introduced in Junos OS Release 8.4. |
| Description | (MX Series routers only) Display bridge statistics. |
| Options | none —Display bridge statistics for all bridge domains in all routing instances. bridge-domain <i>domain-name</i> —(Optional) Display statistics for the specified bridge domain. instance <i>instance-name</i> —(Optional) Display statistics for the specified routing instance. |
| Required Privilege Level | view |
| List of Sample Output | show bridge statistics on page 917 |

Sample Output

show bridge statistics

```

user@host> show bridge statistics
Information for routing instance:

Routing instance : __juniper_private1__
  Index: 1                      Sequence number: 0
  MAC limit: 5000                MACs learned: 0
  Static MACs learned: 0         Non config Static MACs learned: 0
  Handle: 0x829e800

Information for routing instance:

Routing instance : vs1
  Bridging domain : vlan100
  Index: 3                      Sequence number: 0
  MAC limit: 5120                MACs learned: 2
  Static MACs learned: 0         Non config Static MACs learned: 0
  Handle: 0x829e400
  Flags: Bridge instance, Config defined, VLAN : 100
  Local interface: ge-11/0/3.0, Index: 79
    Broadcast packets: 1
    Broadcast bytes : 65
    Multicast packets: 0
    Multicast bytes : 0
    Flooded packets : 0
    Flooded bytes : 0
    Unicast packets : 358624489
    Unicast bytes : 23310592305
    Current MAC count: 1 (Limit 1024)
  Local interface: ge-11/1/4.100, Index: 84
    Broadcast packets: 0
    Broadcast bytes : 0
    Multicast packets: 0
    Multicast bytes : 0
    Flooded packets : 0
    Flooded bytes : 0
    Unicast packets : 0
    Unicast bytes : 0
    Current MAC count: 0 (Limit 1024)
  Local interface: ge-11/1/1.100, Index: 86
    Broadcast packets: 0
    Broadcast bytes : 0
    Multicast packets: 0
    Multicast bytes : 0
    Flooded packets : 0
    Flooded bytes : 0
    Unicast packets : 0
    Unicast bytes : 0
    Current MAC count: 0 (Limit 1024)
  Local interface: ge-11/1/0.100, Index: 87
    Broadcast packets: 0
    Broadcast bytes : 0
    Multicast packets: 0
    Multicast bytes : 0
    Flooded packets : 0
    Flooded bytes : 0
    Unicast packets : 0
    Unicast bytes : 0
    Current MAC count: 0 (Limit 1024)

```

```

Local interface: xe-10/2/0.100, Index: 88
  Broadcast packets: 0
  Broadcast bytes : 0
  Multicast packets: 0
  Multicast bytes : 0
  Flooded packets : 0
  Flooded bytes : 0
  Unicast packets : 358627393
  Unicast bytes : 23310781065
  Current MAC count: 1 (Limit 1024)
Local interface: xe-10/0/0.100, Index: 89
  Broadcast packets: 0
  Broadcast bytes : 0
  Multicast packets: 0
  Multicast bytes : 0
  Flooded packets : 0
  Flooded bytes : 0
  Unicast packets : 0
  Unicast bytes : 0
  Current MAC count: 0 (Limit 1024)

Information for routing instance:

Routing instance : vs1
  Bridging domain : vlan200
    Index: 4
    Sequence number: 0
    MAC limit: 5120
    MACs learned: 0
    Static MACs learned: 0
    Non config Static MACs learned: 0
    Handle: 0x829e600
    Flags: Bridge instance, Config defined, VLAN : 200
  Local interface: ge-11/1/0.200, Index: 90
    Broadcast packets: 0
    Broadcast bytes : 0
    Multicast packets: 0
    Multicast bytes : 0
    Flooded packets : 0
    Flooded bytes : 0
    Unicast packets : 0
    Unicast bytes : 0
    Current MAC count: 0 (Limit 1024)
  Local interface: ge-11/1/1.200, Index: 91
    Broadcast packets: 0
    Broadcast bytes : 0
    Multicast packets: 0
    Multicast bytes : 0
    Flooded packets : 0
    Flooded bytes : 0
    Unicast packets : 0
    Unicast bytes : 0
    Current MAC count: 0 (Limit 1024)
  Local interface: ge-11/1/4.200, Index: 92
    Broadcast packets: 0
    Broadcast bytes : 0
    Multicast packets: 0
    Multicast bytes : 0
    Flooded packets : 0
    Flooded bytes : 0
    Unicast packets : 0
    Unicast bytes : 0
    Current MAC count: 0 (Limit 1024)
  Local interface: xe-10/0/0.200, Index: 93

```



```
Broadcast packets:          0
Broadcast bytes  :          0
Multicast packets:          0
Multicast bytes  :          0
Flooded packets  :          0
Flooded bytes   :          0
Unicast packets  :          0
Unicast bytes   :          0
Current MAC count:          0 (Limit 1024)
Local interface: xe-10/2/0.200, Index: 94
Broadcast packets:          4
Broadcast bytes  :        260
Multicast packets:          0
Multicast bytes  :          0
Flooded packets  :          0
Flooded bytes   :          0
Unicast packets  :          0
Unicast bytes   :          0
Current MAC count:          0 (Limit 1024)
```

show l2-learning global-information

| | |
|---------------------------------|---|
| Syntax | show l2-learning global-information |
| Release Information | Command introduced in Junos OS Release 8.4. |
| Description | (MX Series routers only) Display Layer 2 learning process-related information for the entire router. |
| Options | This command has no options. |
| Required Privilege Level | view |
| List of Sample Output | show l2-learning global-information on page 920 |
| Output Fields | Table 213 on page 920 describes the output fields for the show l2-learning global-information command. Output fields are listed in the approximate order in which they appear. |

Table 213: show l2-learning global-information Output Fields

| Field Name | Field Description |
|-------------------------------|--|
| MAC aging interval | Configured timeout interval, in seconds, for all MAC table entries. |
| MAC learning | Status of MAC learning: Enabled or Disabled . |
| MAC statistics | Status of MAC accounting: Enabled or Disabled . |
| MAC limit Count | Configured maximum limit on the number of MAC addresses that can be learned. |
| MAC limit hit flag | Status of the learned MAC limit hit flag: Enabled (the learned MAC exceeds the global MAC limit) or Disabled (the learned MAC does not exceed the global MAC limit). |
| MAC packet action drop | Status of action to drop packets after the configured MAC address limit is reached: Enabled (packets are dropped) or Disabled (packets are forwarded). |

Sample Output

show l2-learning global-information

```

user@host> show l2-learning global-information
Global Configuration:

MAC aging interval      : 300
MAC learning            : Enabled
MAC statistics          : Disabled
MAC limit Count         : 393215
MAC limit hit flag      : Disabled
MAC packet action drop  : Disabled

```

show l2-learning global-mac-count

| | |
|---------------------------------|--|
| Syntax | show l2-learning global-mac-count |
| Release Information | Command introduced in Junos OS Release 9.3. |
| Description | (MX Series routers only) Display the total number of dynamic and static MAC addresses learned for the entire router. |
| Options | This command has no options. |
| Required Privilege Level | view |
| List of Sample Output | show l2-learning global-mac-count on page 921 |
| Output Fields | Displays the total number of dynamic and static MAC addresses learned for the entire router. |

Sample Output

| | |
|--|---|
| <code>show l2-learning global-mac-count</code> | <code>user@host> show l2-learning global-mac-count</code> 100 dynamic and static MAC addresses learned globally |
|--|---|

show l2-learning instance

| | |
|---------------------------------|---|
| Syntax | show l2-learning instance |
| Release Information | (MX Series routers only) Command introduced in Junos OS Release 8.4. |
| Description | Display Layer 2 learning properties for all the configured routing instances. |
| Options | This command has no options. |
| Required Privilege Level | view |
| List of Sample Output | show l2-learning instance on page 923 |
| Output Fields | Table 214 on page 922 describes the output fields for the show l2-learning instance command. Output fields are listed in the approximate order in which they appear. |

Table 214: show l2-learning instance Output Fields

| Field Name | Field Description |
|-------------------------------|--|
| Routing Instance | Name of routing instance. |
| Bridging Domain | Name of bridging domain. On MX Series routers you can use the show l2-learning instance <extensive> command option to display the Bridge Service-id information which includes the Config Service ID and the Active Service ID. |
| Index | Number associated with the routing instance or bridging domain. |
| Logical System | Name of logical system or Default if no logical system is configured. |
| Routing instance flags | Status of Layer 2 learning properties for each routing instance: <ul style="list-style-type: none"> • DL—MAC learning is disabled. • SE—MAC accounting is enabled. • AD—Packets are dropped after MAC address limit is reached. • LH—The maximum number of MAC addresses has been learned on the routing instance. The routing instance is not able to learn any additional MAC addresses. |
| MAC limit | Maximum number of MAC addresses that can be learned from each interface in the routing instance or bridging domain. |

Sample Output

`show l2-learning
instance`

```
user@host> show l2-learning instance
```

Information for routing instance:

Routing Instance flags (DL -disable learning, SE -stats enabled,
AD -packet action drop, LH -mac limit hit)

| Routing Instance | Bridging Domain | Index | Logical System | Routing flags | MAC limit |
|----------------------|--------------------|-------|-------------------|------------------|--------------|
| __juniper_private1__ | | 1 | Default | | 5000 |
| vs1 | vlan100 | 3 | Default | | 5120 |
| vs1 | vlan200 | 4 | Default | | 5120 |

show l2-learning interface

| | |
|---------------------------------|--|
| Syntax | show l2-learning interface |
| Release Information | Command introduced in Junos OS Release 8.4. |
| Description | (MX Series routers only) Display Layer 2 learning information for all the interfaces. |
| Options | This command has no options. |
| Required Privilege Level | view |
| List of Sample Output | show l2-learning interface on page 925 |
| Output Fields | Table 215 on page 924 describes the output fields for the show l2-learning interface command. Output fields are listed in the approximate order in which they appear. |

Table 215: show l2-learning interfaceOutput Fields

| Field Name | Field Description |
|--------------------------------|---|
| Logical interface | Name of the logical interface. |
| Index | Index of the interface. |
| Routing Instance | Number of the routing instance to which the interface belongs. |
| Interface device | Value of the order in which the Junos OS finds and initializes the interface. |
| Logical interface flags | Status of Layer 2 learning properties for each interface: <ul style="list-style-type: none">• DL—MAC learning is disabled.• SE—MAC accounting is enabled.• AD—Packets are dropped after the MAC interface limit is reached.• MAC limit—Maximum number of MAC addresses that can be learned from the interface. |

Sample Output

**show l2-learning
interface**

user@host> show l2-learning interface

Information for interface family:

Logical Interface flags (DL -disable learning, SE -stats enabled,
AD -packet action drop, LH -mac limit hit)

| Logical interface | Index | Routing instance | Interface device | Logical Interface flags | MAC limit |
|----------------------|-------|---------------------|---------------------|----------------------------|--------------|
| ge-11/0/3.0 | 79 | 3 | 136 | | 1024 |
| ge-11/1/4.100 | 84 | 3 | 150 | | 1024 |
| ge-11/1/1.100 | 86 | 3 | 147 | | 1024 |
| ge-11/1/0.100 | 87 | 3 | 146 | | 1024 |
| xe-10/2/0.100 | 88 | 3 | 144 | | 1024 |
| xe-10/0/0.100 | 89 | 3 | 129 | | 1024 |
| ge-11/1/0.200 | 90 | 4 | 146 | | 1024 |
| ge-11/1/1.200 | 91 | 4 | 147 | | 1024 |
| ge-11/1/4.200 | 92 | 4 | 150 | | 1024 |
| xe-10/0/0.200 | 93 | 4 | 129 | | 1024 |
| xe-10/2/0.200 | 94 | 4 | 144 | | 1024 |

show mac-rewrite interface

| | |
|---------------------------------|--|
| Syntax | show mac-rewrite interface <brief detail> <interface-name> |
| Release Information | Command introduced in Junos OS Release 9.1. |
| Description | (MX Series routers only) Display Layer 2 protocol tunneling information. |
| Options | brief detail —(Optional) Display the specified level of output. interface <i>interface-name</i> —(Optional) Display Layer 2 protocol tunneling information for the specified interface. |
| Required Privilege Level | view |
| List of Sample Output | show mac-rewrite interface on page 926 |
| Output Fields | Table 216 on page 926 lists the output fields for the show mac-rewrite interface command. Output fields are listed in the approximate order in which they appear. |

Table 216: show mac-rewrite interface Output Fields

| Field Name | Field Description | Level of Output |
|------------------|--|---------------------|
| Interface | Name of the interface that has Layer 2 protocol tunneling configured on it. | brief detail |
| Protocols | Layer 2 protocols being tunneled on this interface: Cisco Discovery Protocol (CDP), Spanning Tree Protocol (STP), or VLAN Trunk Protocol (VTP) | brief detail |

Sample Output

```

show mac-rewrite interface  user@host> show mac-rewrite interface
                             Interface      Protocols
                             ge-1/0/1      STP VTP CDP

```


Spanning Tree Operational Mode Commands

Table 217 on page 927 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot the Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), and Multiple Spanning Tree Protocol (MSTP). Commands are listed in alphabetical order.

Table 217: STP Operational Mode Commands

| Task | Command |
|---|---|
| Clear STP protocol. | <code>clear spanning-tree protocol-migration</code> |
| Clear STP statistics. | <code>clear spanning-tree statistics</code> |
| Display STP bridge domain configuration and status. | <code>show spanning-tree bridge</code> |
| Display STP interface configuration and status. | <code>show spanning-tree interface</code> |
| Display MSTP configuration and status. | <code>show spanning-tree mstp configuration</code> |
| Display STP statistics. | <code>show spanning-tree statistics</code> |



NOTE: For more STP-related interface commands, such as `show interface`, see the *Junos Interfaces Command Reference*.

For more STP-related bridging commands, such as `clear bridge`, `show bridge`, and `show l2-learning`, see Layer 2 Bridging and Switching Operational Mode Commands.

For information about how to configure STP, see the *Junos Routing Protocols Configuration Guide*.


clear spanning-tree protocol-migration

| | |
|---------------------------------|--|
| Syntax | <code>clear spanning-tree protocol-migration</code> <code><interface <i>interface-name</i>></code> <code><routing-instance <i>routing-instance-name</i>></code> |
| Release Information | Command introduced in Junos OS Release 9.0. |
| Description | Revert from the original IEEE 802.1D Spanning Tree Protocol (STP) back to the Rapid Spanning Tree Protocol after the force-version statement has been removed from the configuration. |
| Options | none —Reset the STP protocol for all interfaces and all routing instances. interface <i>interface-name</i> —(Optional) Reset the STP protocol for the specified interface only. routing-instance <i>routing-instance-name</i> —(Optional) Reset the STP protocol for a particular routing instance. |
| Additional Information | For information about the force-version statement, see the <i>Junos Routing Protocols Configuration Guide</i> . |
| Required Privilege Level | clear |

Sample Output

| | |
|---|---|
| <code>clear spanning-tree protocol-migration</code> | <code>user@host> clear spanning-tree protocol-migration</code> |
|---|---|

clear spanning-tree statistics

| | |
|---|---|
| Syntax | clear spanning-tree statistics <interface <i>interface-name</i> > <logical-system <i>logical-system-name</i> > |
| Syntax (EX Series Switches and the QFX Series) | clear spanning-tree statistics <interface <i>interface-name</i> > |
| Release Information | Command introduced in Junos OS Release 8.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. |
| Description | Clear Spanning Tree Protocol statistics. |
| Options | <p>none—Reset STP counters for all interfaces for all routing instances.</p> <p>interface <i>interface-name</i>—(Optional) Clear STP statistics for the specified interface only.</p> <p>logical-system <i>logical-system-name</i>—(Optional) Clear STP statistics on a particular logical system.</p> |
| | <div>  <p>NOTE: The logical-system option is not available on QFabric systems.</p> </div> |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none"> • show spanning-tree statistics on page 943 |
| List of Sample Output | clear stp statistics on page 929 |

Sample Output

clear stp statistics user@host> clear stp statistics

show spanning-tree bridge

| | |
|---------------------------------|--|
| Syntax | show spanning-tree bridge <brief detail> <msti <i>msti-id</i> > <routing-instance <i>routing-instance-name</i> > <vlan-id <i>vlan-id</i> > |
| Syntax (QFX Series) | show spanning-tree bridge <brief detail> <msti <i>msti-id</i> > <vlan-id <i>vlan-id</i> > |
| Release Information | Command introduced in Junos OS Release 8.4. Command introduced in Junos OS Release 11.1 for the QFX Series. |
| Description | Display the configured or calculated Spanning Tree Protocol (STP) parameters. |
| Options | <p>none—(Optional) Display brief STP bridge information for all multiple spanning-tree instances (MSTIs).</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>msti <i>msti-id</i>—(Optional) Display STP bridge information for the specified MSTI.</p> <p>routing-instance <i>routing-instance-name</i>—(Optional) Display STP bridge information for the specified routing instance.</p> <p>vlan-id <i>vlan-id</i>—(Optional) Display STP bridge information for the specified VLAN.</p> |
| Required Privilege Level | view |
| List of Sample Output | show spanning-tree bridge routing-instance on page 932 show spanning-tree bridge msti on page 932 show spanning-tree bridge vlan-id (MSTP) on page 933 show spanning-tree bridge (RSTP) on page 933 show spanning-tree bridge vlan-id (RSTP) on page 934 |
| Output Fields | Table 218 on page 930 lists the output fields for the show spanning-tree bridge command. Output fields are listed in the approximate order in which they appear. |

Table 218: show spanning-tree bridge Output Fields

| Field Name | Field Description |
|-----------------------|---|
| Routing instance name | Name of the routing instance under which the bridge is configured. |
| Enabled protocol | Spanning Tree Protocol type enabled. |
| Root ID | Bridge ID of the elected spanning-tree root bridge. The bridge ID consists of a configurable bridge priority and the MAC address of the bridge. |

Table 218: show spanning-tree bridge Output Fields (*continued*)

| Field Name | Field Description |
|--|---|
| Root cost | Calculated cost to reach the root bridge from the bridge where the command is entered. |
| Root port | Interface that is the current elected root port for this bridge. |
| CIST regional root | Bridge ID of the elected MSTP regional root bridge. |
| CIST internal root cost | Calculated cost to reach the regional root bridge from the bridge where the command is entered. |
| Hello time | Configured number of seconds between transmissions of configuration bridge protocol data units (BPDUs). |
| Maximum age | Configured maximum expected arrival time of hello bridge protocol data units (BPDUs). |
| Forward delay | How long an STP bridge port remains in the listening and learning states before transitioning to the forwarding state. |
| Hop count | Configured maximum number of hops a BPDU can be forwarded in the MSTP region. |
| Message age | Number of elapsed seconds since the most recent BPDU was received. |
| Number of topology changes | Total number of STP topology changes detected since the routing device last booted. |
| Time since last topology change | Number of elapsed seconds since the most recent topology change. |
| Bridge ID (Local) | Locally configured bridge ID. The bridge ID consists of a configurable bridge priority and the MAC address of the bridge. |
| Extended system ID | System identifier. |
| MSTI regional root | Bridge ID of the elected MSTP regional root bridge. |

Sample Output

```

show spanning-tree bridge routing-instance user@host> show spanning-tree bridge routing-instance vs1 detail
bridge routing-instance STP bridge parameters
Routing instance name      : vs1
Enabled protocol           : MSTP

STP bridge parameters for CIST
Root ID                    : 32768.00:13:c3:9e:c8:80
Root cost                   : 0
Root port                  : ge-10/2/0
CIST regional root         : 32768.00:13:c3:9e:c8:80
CIST internal root cost    : 22000
Hello time                 : 2 seconds
Maximum age                : 20 seconds
Forward delay              : 15 seconds
Hop count                  : 18
Message age                : 0
Number of topology changes : 1
Time since last topology change : 1191 seconds
Local parameters
  Bridge ID                : 32768.00:90:69:0b:7f:d1
  Extended system ID       : 1

STP bridge parameters for MSTI 1
MSTI regional root         : 32769.00:13:c3:9e:c8:80
Root cost                   : 22000
Root port                  : ge-10/2/0
Hello time                 : 2 seconds
Maximum age                : 20 seconds
Forward delay              : 15 seconds
Hop count                  : 18
Number of topology changes : 1
Time since last topology change : 1191 seconds
Local parameters
  Bridge ID                : 32769.00:90:69:0b:7f:d1
  Extended system ID       : 1

STP bridge parameters for MSTI 2
MSTI regional root         : 32770.00:13:c3:9e:c8:80
Root cost                   : 22000
Root port                  : ge-10/2/0
Hello time                 : 2 seconds
Maximum age                : 20 seconds
Forward delay              : 15 seconds
Hop count                  : 18
Number of topology changes : 1
Time since last topology change : 1191 seconds
Local parameters
  Bridge ID                : 32770.00:90:69:0b:7f:d1
  Extended system ID       : 1

show spanning-tree bridge msti user@host> show spanning-tree bridge msti 1 routing-instance vs1 detail
bridge msti STP bridge parameters
Routing instance name      : vs1
Enabled protocol           : MSTP

STP bridge parameters for MSTI 1
MSTI regional root         : 32769.00:13:c3:9e:c8:80

```

```

Root cost                : 22000
Root port                : xe-10/2/0
Hello time               : 2 seconds
Maximum age              : 20 seconds
Forward delay            : 15 seconds
Hop count                : 18
Number of topology changes : 1
Time since last topology change : 1191 seconds
Local parameters
  Bridge ID              : 32769.00:90:69:0b:7f:d1
  Extended system ID     : 1

```

show spanning-tree bridge vlan-id (MSTP)

user@host> show spanning-tree bridge vlan-id 1101 routing-instance vs1 detail

```

STP bridge parameters
Routing instance name    : vs1
Enabled protocol        : MSTP

STP bridge parameters for CIST
Root ID                 : 32768.00:13:c3:9e:c8:80
Root cost               : 0
Root port              : xe-10/2/0
CIST regional root     : 32768.00:13:c3:9e:c8:80
CIST internal root cost : 22000
Hello time              : 2 seconds
Maximum age            : 20 seconds
Forward delay          : 15 seconds
Hop count              : 18
Message age            : 0
Number of topology changes : 0
Local parameters
  Bridge ID             : 32768.00:90:69:0b:7f:d1
  Extended system ID    : 1
  Hello time            : 2 seconds
  Maximum age           : 20 seconds
  Forward delay         : 15 seconds
  Path cost method      : 32 bit
  Maximum hop count     : 20

```

show spanning-tree bridge (RSTP)

user@host> show spanning-tree bridge

```

STP bridge parameters
Routing instance name    : GLOBAL
Enabled protocol        : RSTP
Root ID                 : 28672.00:90:69:0b:3f:d0
Hello time              : 2 seconds
Maximum age            : 20 seconds
Forward delay          : 15 seconds
Message age            : 0
Number of topology changes : 58
Time since last topology change : 14127 seconds
Local parameters
  Bridge ID             : 28672.00:90:69:0b:3f:d0
  Extended system ID    : 0

STP bridge parameters for bridge VLAN 10
Root ID                 : 28672.00:90:69:0b:3f:d0
Hello time              : 2 seconds
Maximum age            : 20 seconds
Forward delay          : 15 seconds
Message age            : 0
Number of topology changes : 58

```

```
Time since last topology change : 14127 seconds
Local parameters
  Bridge ID                : 28672.00:90:69:0b:3f:d0
  Extended system ID       : 0

STP bridge parameters for bridge VLAN 20
  Root ID                  : 28672.00:90:69:0b:3f:d0
  Hello time               : 2 seconds
  Maximum age              : 20 seconds
  Forward delay            : 15 seconds
  Message age              : 0
  Number of topology changes : 58
  Time since last topology change : 14127 seconds
Local parameters
  Bridge ID                : 28672.00:90:69:0b:3f:d0
  Extended system ID       : 0
```

**show spanning-tree
bridge vlan-id (RSTP)**

```
user@host> show spanning-tree bridge vlan-id 10
STP bridge parameters
Routing instance name      : GLOBAL
Enabled protocol          : RSTP

STP bridge parameters for VLAN 10
  Root ID                  : 28672.00:90:69:0b:3f:d0
  Hello time               : 2 seconds
  Maximum age              : 20 seconds
  Forward delay            : 15 seconds
  Message age              : 0
  Number of topology changes : 58
  Time since last topology change : 14127 seconds
Local parameters
  Bridge ID                : 28672.00:90:69:0b:3f:d0
  Extended system ID       : 0
```


show spanning-tree interface

| | |
|---|--|
| Syntax | show spanning-tree interface <brief detail> <msti <i>msti-id</i> > <routing-instance <i>routing-instance-name</i> > <vlan-id <i>vlan-id</i> > |
| Syntax (EX Series Switches and the QFX Series) | show spanning-tree interface <brief detail> <msti <i>msti-id</i> > <vlan-id <i>vlan-id</i> > |
| Release Information | Command introduced in Junos OS Release 8.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. |
| Description | Display the configured or calculated interface-level STP parameters. |
| Options | <p>none—Display brief STP interface information.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>msti <i>msti-id</i>—(Optional) Display STP interface information for the specified MST instance.</p> <p>routing-instance <i>routing-instance-name</i>—(Optional) Display STP interface information for the specified routing instance.</p> <p>vlan-id <i>vlan-id</i>—(Optional) Display STP interface information for the specified VLAN.</p> |
| Required Privilege Level | view |
| List of Sample Output | show spanning-tree interface on page 937 show spanning-tree interface (QFX Series) on page 937 show spanning-tree interface detail on page 938 show spanning-tree interface msti on page 939 show spanning-tree interface vlan-id on page 940 show spanning-tree interface (VSTP) on page 940 show spanning-tree interface vlan-id (VSTP) on page 940 |
| Output Fields | Table 219 on page 935 lists the output fields for the show spanning-tree interface command. Output fields are listed in the approximate order in which they appear. |

Table 219: show spanning-tree Interface Output Fields

| Field Name | Field Description |
|-----------------------|---|
| Interface name | Interface configured to participate in the STP, RSTP, VSTP, or MSTP instance. |

Table 219: show spanning-tree Interface Output Fields (*continued*)

| Field Name | Field Description |
|-----------------------------|---|
| Port ID | Logical interface identifier configured to participate in the MSTP or VSTP instance. |
| Designated port ID | Port ID of the designated port for the LAN segment to which this interface is attached. |
| Designated bridge ID | Bridge ID of the designated bridge for the LAN segment to which this interface is attached. |
| Port Cost | Configured cost for the interface. |
| Port State | STP port state: forwarding (FWD), blocking (BLK), listening, learning, or disabled. |
| Port Role | MSTP, VSTP, or RSTP port role: designated (DESG), backup (BKUP), alternate (ALT), (ROOT), or Root Prevented (Root-Prev). |
| Link type | MSTP, VSTP, or RSTP link type. Shared or point-to-point (pt-pt) and edge or nonedge. |
| Alternate | Identifies the interface as an MSTP, VSTP, or RSTP alternate root port (Yes) or nonalternate root port (No). |
| Boundary Port | Identifies the interface as an MSTP regional boundary port (Yes) or nonboundary port (No). |

Sample Output

show spanning-tree interface

```
user@host> show spanning-tree interface routing-instance vs1 detail
Spanning tree interface parameters for instance 0
```

| Interface | Port ID | Designated port ID | Designated bridge ID | Port Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------|-------|------|
| ae1 | 128:1 | 128:1 | 32768.0090690b47d1 | 1000 | FWD | DESG |
| ge-2/1/2 | 128:2 | 128:2 | 32768.0090690b47d1 | 20000 | FWD | DESG |
| ge-2/1/5 | 128:3 | 128:3 | 32768.0090690b47d1 | 29999 | FWD | DESG |
| ge-2/2/1 | 128:4 | 128:26 | 32768.0013c39ec880 | 20000 | FWD | ROOT |
| xe-9/2/0 | 128:5 | 128:5 | 32768.0090690b47d1 | 2000 | FWD | DESG |
| xe-9/3/0 | 128:6 | 128:6 | 32768.0090690b47d1 | 2000 | FWD | DESG |

Spanning tree interface parameters for instance 1

| Interface | Port ID | Designated port ID | Designated bridge ID | Port Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------|-------|------|
| ae1 | 128:1 | 128:1 | 32769.0090690b47d1 | 1000 | FWD | DESG |
| ge-2/1/2 | 128:2 | 128:2 | 32769.0090690b47d1 | 20000 | FWD | DESG |
| ge-2/1/5 | 128:3 | 128:3 | 32769.0090690b47d1 | 29999 | FWD | DESG |
| ge-2/2/1 | 128:4 | 128:26 | 32769.0013c39ec880 | 20000 | FWD | ROOT |
| xe-9/2/0 | 128:5 | 128:5 | 32769.0090690b47d1 | 2000 | FWD | DESG |
| xe-9/3/0 | 128:6 | 128:6 | 32769.0090690b47d1 | 2000 | FWD | DESG |

Spanning tree interface parameters for instance 2

| Interface | Port ID | Designated port ID | Designated bridge ID | Port Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------|-------|------|
| ae1 | 128:1 | 128:1 | 32770.0090690b47d1 | 1000 | FWD | DESG |
| ge-2/1/2 | 128:2 | 128:2 | 32770.0090690b47d1 | 20000 | FWD | DESG |
| ge-2/1/5 | 128:3 | 128:3 | 32770.0090690b47d1 | 29999 | FWD | DESG |
| ge-2/2/1 | 128:4 | 128:26 | 32770.0013c39ec880 | 20000 | FWD | ROOT |
| xe-9/2/0 | 128:5 | 128:5 | 32770.0090690b47d1 | 2000 | FWD | DESG |
| xe-9/3/0 | 128:6 | 128:6 | 32770.0090690b47d1 | 2000 | FWD | DESG |

show spanning-tree interface (QFX Series)

```
user@1f0> show spanning-tree interface routing-instance vs1 detail
Spanning tree interface parameters for instance 0
```

| Interface | Port ID | Designated port ID | Designated bridge ID | Port Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------|-------|------|
| ae1 | 128:1 | 128:1 | 32768.0090690b47d1 | 1000 | FWD | DESG |
| ge-2/1/2 | 128:2 | 128:2 | 32768.0090690b47d1 | 20000 | FWD | DESG |
| ge-2/1/5 | 128:3 | 128:3 | 32768.0090690b47d1 | 29999 | FWD | DESG |
| ge-2/2/1 | 128:4 | 128:26 | 32768.0013c39ec880 | 20000 | FWD | ROOT |
| xe-9/2/0 | 128:5 | 128:5 | 32768.0090690b47d1 | 2000 | FWD | DESG |
| xe-9/3/0 | 128:6 | 128:6 | 32768.0090690b47d1 | 2000 | FWD | DESG |

Spanning tree interface parameters for instance 1

| Interface | Port ID | Designated port ID | Designated bridge ID | Port Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------|-------|------|
| ae1 | 128:1 | 128:1 | 32769.0090690b47d1 | 1000 | FWD | DESG |
| ge-2/1/2 | 128:2 | 128:2 | 32769.0090690b47d1 | 20000 | FWD | DESG |
| ge-2/1/5 | 128:3 | 128:3 | 32769.0090690b47d1 | 29999 | FWD | DESG |
| ge-2/2/1 | 128:4 | 128:26 | 32769.0013c39ec880 | 20000 | FWD | ROOT |
| xe-9/2/0 | 128:5 | 128:5 | 32769.0090690b47d1 | 2000 | FWD | DESG |
| xe-9/3/0 | 128:6 | 128:6 | 32769.0090690b47d1 | 2000 | FWD | DESG |

Spanning tree interface parameters for instance 2

| Interface | Port ID | Designated port ID | Designated bridge ID | Port Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------|-------|------|
| ae1 | 128:1 | 128:1 | 32770.0090690b47d1 | 1000 | FWD | DESG |
| ge-2/1/2 | 128:2 | 128:2 | 32770.0090690b47d1 | 20000 | FWD | DESG |
| ge-2/1/5 | 128:3 | 128:3 | 32770.0090690b47d1 | 29999 | FWD | DESG |
| ge-2/2/1 | 128:4 | 128:26 | 32770.0013c39ec880 | 20000 | FWD | ROOT |
| xe-9/2/0 | 128:5 | 128:5 | 32770.0090690b47d1 | 2000 | FWD | DESG |
| xe-9/3/0 | 128:6 | 128:6 | 32770.0090690b47d1 | 2000 | FWD | DESG |

**show spanning-tree
interface detail**user@host> **show spanning-tree interface routing-instance vs1 detail**

Spanning tree interface parameters for instance 0

```
Interface name           : ae1
Port identifier          : 128.1
Designated port ID      : 128.1
Port cost                : 1000
Port state               : Forwarding
Designated bridge ID     : 32768.00:90:69:0b:47:d1
Port role                : Designated
Link type                : Pt-Pt/NONEDGE
Boundary port            : No
```

```
Interface name           : ge-2/1/2
Port identifier          : 128.2
Designated port ID      : 128.2
Port cost                : 20000
Port state               : Forwarding
Designated bridge ID     : 32768.00:90:69:0b:47:d1
Port role                : Designated
Link type                : Pt-Pt/NONEDGE
Boundary port            : No
```

```
Interface name           : ge-2/1/5
Port identifier          : 128.3
Designated port ID      : 128.3
Port cost                : 29999
Port state               : Forwarding
Designated bridge ID     : 32768.00:90:69:0b:47:d1
Port role                : Designated
Link type                : Pt-Pt/NONEDGE
Boundary port            : No
```

```
Interface name           : ge-2/2/1
Port identifier          : 128.4
Designated port ID      : 128.26
Port cost                : 20000
Port state               : Forwarding
Designated bridge ID     : 32768.00:13:c3:9e:c8:80
Port role                : Root
Link type                : Pt-Pt/NONEDGE
Boundary port            : No
```

```
Interface name           : xe-9/2/0
Port identifier          : 128.5
Designated port ID      : 128.5
Port cost                : 2000
Port state               : Forwarding
Designated bridge ID     : 32768.00:90:69:0b:47:d1
Port role                : Designated
```

```

Link type           : Pt-Pt/NONEDGE
Boundary port       : No

Interface name       : xe-9/3/0
Port identifier      : 128.6
Designated port ID   : 128.6
Port cost            : 2000
Port state           : Forwarding
Designated bridge ID : 32768.00:90:69:0b:47:d1
Port role            : Designated
Link type           : Pt-Pt/NONEDGE
Boundary port       : No

```

Spanning tree interface parameters for instance 1

```

Interface name       : ae1
Port identifier      : 128.1
Designated port ID   : 128.1
Port cost            : 1000
Port state           : Forwarding
Designated bridge ID : 32768.00:90:69:0b:47:d1
Port role            : Designated
Link type           : Pt-Pt/NONEDGE
Boundary port       : No

```

```

Interface name       : ge-2/1/2
Port identifier      : 128.2
Designated port ID   : 128.2
Port cost            : 20000
Port state           : Forwarding
Designated bridge ID : 32768.00:90:69:0b:47:d1
Port role            : Designated
Link type           : Pt-Pt/NONEDGE
Boundary port       : No

```

```

Interface name       : ge-2/1/5
Port identifier      : 128.3
Designated port ID   : 128.3
Port cost            : 29999
Port state           : Forwarding
Designated bridge ID : 32768.00:90:69:0b:47:d1
Port role            : Designated
Link type           : Pt-Pt/NONEDGE
Boundary port       : No

```

```

Interface name       : ge-2/2/1
Port identifier      : 128.4
Designated port ID   : 128.26
Port cost            : 20000
Port state           : Forwarding
Designated bridge ID : 32768.00:13:c3:9e:c8:80
Port role            : Root
Link type           : Pt-Pt/NONEDGE
Boundary port       : No

```

...

show spanning-tree
interface msti

```

user@host> show spanning-tree interface msti 1 routing-instance vs1 detail
Spanning tree interface parameters for instance 1

```

| Interface | Port ID | Designated port ID | Designated bridge ID | Port Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------|-------|------|
| xe-7/0/0 | 128:1 | 128:1 | 32769.0090690b4fd1 | 2000 | FWD | DESG |
| ge-5/1/0 | 128:2 | 128:2 | 32769.0090690b4fd1 | 20000 | FWD | DESG |
| ge-5/1/1 | 128:3 | 128:3 | 32769.0090690b4fd1 | 20000 | FWD | DESG |
| ae1 | 128:4 | 128:1 | 32769.0090690b47d1 | 10000 | BLK | ALT |
| ge-5/1/4 | 128:5 | 128:3 | 32769.0090690b47d1 | 20000 | BLK | ALT |
| xe-7/2/0 | 128:6 | 128:6 | 32769.0090690b47d1 | 2000 | FWD | ROOT |

show spanning-tree interface vlan-id

user@host> show spanning-tree interface vlan-id 101 routing-instance vs1 detail
Spanning tree interface parameters for instance 0

| Interface | Port ID | Designated port ID | Designated bridge ID | Port Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-11/0/5 | 128:1 | 128:1 | 32768.0090690b7fd1 | 20000 | FWD | DESG |
| ge-11/0/6 | 128:2 | 128:1 | 32768.0090690b7fd1 | 20000 | BLK | BKUP |
| ge-11/1/0 | 128:3 | 128:2 | 32768.0090690b4fd1 | 20000 | BLK | ALT |
| ge-11/1/1 | 128:4 | 128:3 | 32768.0090690b4fd1 | 20000 | BLK | ALT |
| ge-11/1/4 | 128:5 | 128:1 | 32768.0090690b47d1 | 20000 | BLK | ALT |
| xe-10/0/0 | 128:6 | 128:5 | 32768.0090690b4fd1 | 2000 | BLK | ALT |
| xe-10/2/0 | 128:7 | 128:4 | 32768.0090690b47d1 | 2000 | FWD | ROOT |

show spanning-tree interface (VSTP)

user@host> show spanning-tree interface
Spanning tree interface parameters for instance 0

| Interface | Port ID | Designated port ID | Designated bridge ID | Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|-------|-------|------|
| ge-1/0/1 | 128:1 | 128:1 | 28672.0090690b3fe0 | 20000 | FWD | DESG |
| ge-1/0/2 | 128:2 | 128:2 | 28672.0090690b3fe0 | 20000 | FWD | DESG |

Spanning tree interface parameters for VLAN 10

| Interface | Port ID | Designated port ID | Designated bridge ID | Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|-------|-------|------|
| ge-1/0/1 | 128:1 | 128:1 | 28672.0090690b3fe0 | 20000 | FWD | DESG |
| ge-1/0/2 | 128:2 | 128:2 | 28672.0090690b3fe0 | 20000 | FWD | DESG |

Spanning tree interface parameters for VLAN 20

| Interface | Port ID | Designated port ID | Designated bridge ID | Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|-------|-------|------|
| ge-1/0/1 | 128:1 | 128:1 | 28672.0090690b3fe0 | 20000 | FWD | DESG |
| ge-1/0/2 | 128:2 | 128:2 | 28672.0090690b3fe0 | 20000 | FWD | DESG |

show spanning-tree interface vlan-id (VSTP)

user@host> show spanning-tree interface vlan-id 10
Spanning tree interface parameters for VLAN 10

| Interface | Port ID | Designated port ID | Designated bridge ID | Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|-------|-------|------|
| ge-1/0/1 | 128:1 | 128:1 | 28672.0090690b3fe0 | 20000 | FWD | DESG |
| ge-1/0/2 | 128:2 | 128:2 | 28672.0090690b3fe0 | 20000 | FWD | DESG |

show spanning-tree mstp configuration

| | |
|---|---|
| Syntax | show spanning-tree mstp configuration <brief detail> <routing-instance <i>routing-instance-name</i> > |
| Syntax (EX Series Switch and the QFX Series) | show spanning-tree mstp configuration <brief detail> |
| Release Information | Command introduced in Junos OS Release 8.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. |
| Description | Display the MSTP configuration. |
| Options | none —Display MSTP configuration information. brief detail —(Optional) Display the specified level of output. routing-instance <i>routing-instance-name</i> —(Optional) Display MSTP configuration information for the specified routing instance. |
| Required Privilege Level | view |
| List of Sample Output | show spanning-tree mstp configuration detail on page 942 show spanning-tree mstp configuration detail (QFX Series) on page 942 |
| Output Fields | Table 220 on page 941 lists the output fields for the show spanning-tree mstp configuration command. Output fields are listed in the approximate order in which they appear. |

Table 220: show spanning-tree mstp configuration Output Fields

| Field Name | Field Description |
|-----------------------------|--|
| Context id | Internally generated identifier. |
| Region name | MSTP region name carried in the MSTP BPDUs. |
| Revision | Revision number of the MSTP configuration. |
| Configuration digest | Numerical value derived from the VLAN-to-instance mapping table. |
| MSTI | MST instance identifier. |
| Member VLANs | VLAN identifiers associated with the MSTI. |

Sample Output

**show spanning-tree
mstp configuration
detail**

```
user@host> show spanning-tree mstp configuration routing-instance vs1 detail
MSTP configuration information
Context identifier      : 1
Region name             : henry
Revision                : 3
Configuration digest    : 0x6da4b5c4fd587757eef35675365e1
```

```

MSTI      Member VLANs
  0 0-99,101-199,201-4094
  1 100
  2 200
```

**show spanning-tree
mstp configuration
detail (QFX Series)**

```
user@1f0> show spanning-tree mstp configuration routing-instance vs1 detail
MSTP configuration information
Context identifier      : 1
Region name             : henry
Revision                : 3
Configuration digest    : 0x6da4b5c4fd587757eef35675365e1
```

```

MSTI      Member VLANs
  0 0-99,101-199,201-4094
  1 100
  2 200
```


show spanning-tree statistics

| | |
|---|--|
| Syntax | show spanning-tree statistics <brief detail> <interface <i>interface-name</i> > <routing-instance <i>routing-instance-name</i> > |
| Syntax (EX Series Switch and the QFX Series) | show spanning-tree statistics <brief detail> <interface <i>interface-name</i> vlan <i>vlan-id</i> > |
| Release Information | Command introduced in Junos OS Release 8.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for QFX Series switches. |
| Description | Display STP statistics. |
| Options | <p>none—Display brief STP statistics.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>interface <i>interface-name</i>—(Optional) Display STP statistics for the specified interface.</p> <p>routing-instance <i>routing-instance-name</i>—(Optional) Display STP statistics for the specified routing instance.</p> |
| Required Privilege Level | view |
| List of Sample Output | show spanning-tree statistics routing-instance on page 944 show spanning-tree statistics interface routing-instance detail on page 944 |
| Output Fields | Table 221 on page 943 lists the output fields for the show spanning-tree statistics command. Output fields are listed in the approximate order in which they appear. |

Table 221: show spanning-tree statistics Output Fields

| Field Name | Field Description |
|--------------------------------------|--|
| Message type | Type of message being counted. |
| BPDUs sent | Total number of BPDUs sent. |
| BPDUs received | Total number of BPDUs received. |
| BPDUs sent in last 5 secs | Number of BPDUs sent in the most recent 5-second period. |
| BPDUs received in last 5 secs | Number of BPDUs received in the most recent 5-second period. |
| Interface | Interface for which the statistics are being displayed. |

Table 221: show spanning-tree statistics Output Fields (*continued*)

| Field Name | Field Description |
|------------------------|--|
| Next BPDU transmission | Number of seconds until the next BPDU is scheduled to be sent. |

Sample Output

**show spanning-tree
statistics
routing-instance**

```
user@host> show spanning-tree statistics routing-instance vs1 detail
Routing instance level STP statistics
Message type           : bpdus
BPDUs sent              : 121
BPDUs received          : 537
BPDUs sent in last 5 secs : 5
BPDUs received in last 5 secs : 27
```

**show spanning-tree
statistics interface
routing-instance detail**

```
user@host> show spanning-tree statistics interface ge-11/1/4 routing-instance vs1 detail
Interface  BPDUs sent  BPDUs received  Next BPDU
              transmission
ge-11/1/4      7           190             0
```

PART 5

VPNs

- [VPN Operational Mode Commands on page 947](#)

VPN Operational Mode Commands

Table 222 on page 947 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot Layer 2 circuits, Layer 2 virtual private networks (VPNs), virtual private LAN service (VPLS), and Layer 3 VPNs. Commands are listed in alphabetical order.

Table 222: Layer 2 Circuit, Layer 2 VPN, and VPLS Operational Mode Commands

| Task | Command |
|---|--|
| Clear MAC address entries from the VPLS table. | <code>clear vpls mac-address</code> |
| Clear MAC addresses from the VPLS table. | <code>clear vpls mac-table</code> |
| Manually trigger a switch from the active pseudowire to the redundant pseudowire. | <code>request l2circuit-switchover</code> |
| Display Layer 3 dynamic tunnel database information. | <code>show dynamic-tunnels database</code> |
| Display Host fast reroute (HFRR) profile information. | <code>show hfr profiles</code> |
| Display ingress replication provider tunnel information. | <code>show ingress-replication mvpn</code> |
| Display Layer 2 circuit information. | <code>show l2circuit connections</code> |
| Display Layer 2 VPN information. | <code>show l2vpn connections</code> |
| Display multicast VPN c-multicast route information. | <code>show mvpn c-multicast</code> |
| Display multicast VPN instance information. | <code>show mvpn instance</code> |
| Display multicast VPN neighbor information. | <code>show mvpn neighbor</code> |
| Display virtual private LAN service (VPLS) information. | <code>show vpls connections</code> |
| Display the pending events in the level 2 address learning process (l2ald) routing socket code (rtsock) update queue. | <code>show vpls flood event-queue</code> |

Table 222: Layer 2 Circuit, Layer 2 VPN, and VPLS Operational Mode Commands (*continued*)

| Task | Command |
|--|--|
| Display VPLS information related to the level 2 address learning process for the specified routing instance. | show vpls flood instance |
| Display VPLS route information related to the level 2 address learning process. | show vpls flood route |
| Display learned VPLS MAC address information. | show vpls mac-table |
| Display VPLS statistics. | show vpls statistics |



NOTE: For information about how to configure Layer 2 circuits, Layer 2 VPNs, VPLS, and Layer 3 VPNs, see the *Junos VPNs Configuration Guide*.

clear vpls mac-address

| | |
|---------------------------------|--|
| Syntax | clear vpls mac-address <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> < <i>mac-address</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | (T Series and M Series routers, except for the M160 router) Clear media access control (MAC) address entries from the virtual private LAN service (VPLS) table. |
| Options | <p>none—Clear all MAC address entries from the VPLS table for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Clear all MAC address entries for a VPLS instance from the VPLS table.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><i>mac-address</i>—(Optional) Clear a specific MAC address in a VPLS instance from the VPLS table.</p> |
| Required Privilege Level | maintenance |
| List of Sample Output | clear vpls mac-address on page 949 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear vpls mac-address user@host> clear vpls mac-address

clear vpls mac-table

| | |
|---------------------------------|---|
| Syntax | <code>clear vpls mac-table</code> <code><instance <i>instance-name</i>></code> <code><interface <i>interface-name</i>></code> <code><logical-system (all <i>logical-system-name</i>)></code> <code><mac-address></code> <code><vlan-id></code> |
| Release Information | Command introduced before Junos OS Release 9.5. |
| Description | (MX Series routers) Clear media access control (MAC) addresses from the virtual private LAN service (VPLS) MAC table. |
| Options | <p>none—Clear all MAC addresses from the VPLS table for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Clear all MAC addresses for a VPLS instance from the VPLS table.</p> <p>interface <i>interface-name</i>—(Optional) Clear all MAC addresses for a VPLS interface from the VPLS table.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>mac-address—(Optional) Clear a specific MAC address in a VPLS instance from the VPLS table.</p> <p>vlan-id—(Optional) Clear MAC addresses on a specified VLAN (0 through 4095).</p> |
| Required Privilege Level | maintenance |
| List of Sample Output | clear vpls mac-table on page 950 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear vpls mac-table user@host> clear vpls mac-table

request l2circuit-switchover

| | |
|---------------------------------|---|
| Syntax | request l2circuit-switchover <logical-system (all logical-system-name) > <neighbor <i>address</i> > <virtual-circuit-id <i>identifier</i> > |
| Release Information | Command introduced in Junos OS Release 9.2. |
| Description | Manually trigger a switch from the active pseudowire to the redundant pseudowire. This command can be useful when performing network maintenance. |
| Options | <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>neighbor <i>address</i>—(Optional) Trigger a switch of all of the active pseudowire connections with the specified neighbor to their respective redundant pseudowires.</p> <p>virtual-circuit-id <i>identifier</i>—(Optional) Trigger a switch from the active pseudowire connection of the specified Layer 2 circuit to its redundant pseudowire.</p> |
| Required Privilege Level | maintenance |
| List of Sample Output | request l2circuit-switchover virtual-circuit-id on page 951 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

```
request
l2circuit-switchover
virtual-circuit-id
```

```
user@host>request l2circuit-switchover virtual-circuit-id 12
```

show dynamic-tunnels database

| | |
|---------------------------------|--|
| Syntax | show dynamic-tunnels database <destination> <logical-system (all <i>logical-system-name</i>) > <table <i>routing-table-name</i> > |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display dynamic tunnel database information. |
| Options | <p>none—Display dynamic tunnel database information for all destinations and routing tables.</p> <p>destination—(Optional) Display database entries for the specified IP address (with optional destination prefix length) only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>table <i>routing-table-name</i>—(Optional) Display database entries for the specified table only.</p> |
| Required Privilege Level | view |
| List of Sample Output | show dynamic-tunnels database (Tunnel Is Up) on page 953 show dynamic-tunnels database (No Tunnel PIC) on page 953 show dynamic-tunnels database (Tunnel Is Expiring) on page 953 show dynamic-tunnels database (Destination Specified) on page 953 |
| Output Fields | Table 223 on page 952 lists the output fields for the show dynamic-tunnels database command. Output fields are listed in the approximate order in which they appear. |

Table 223: show dynamic-tunnels database Output Fields

| Field Name | Field Description |
|----------------------------|---|
| Table | Name of the routing table (for example, inet.0). |
| Destination-network | Destination IP address and subnet. |
| Tunnel to | Destination IP address and prefix of the tunnel. |
| State | State of the tunnel: Up , Up (expires in <i>nn:nn:nn</i>seconds) , or Dn (down). |
| Reference count | Number of routes across the dynamic tunnel that are currently being resolved. |
| Next-hop type | Type of tunnel: GRE. |
| Source address | Source IP address of the tunnel. |

Table 223: show dynamic-tunnels database Output Fields (*continued*)

| Field Name | Field Description |
|-----------------|--|
| Next-hop | IP address of the destination interface. |
| State | State of the destination interface: Up, Dn, or Dn (no tunnel pic). |

Sample Output

show dynamic-tunnels database (Tunnel Is Up)

```
user@host> show dynamic-tunnels database
Table: inet.3

Destination-network: 10.255.120.94/32
Tunnel to: 10.255.120.94/32 State: Up
Reference count: 2
Next-hop type: gre
Source address: 10.255.120.92
Next hop: gr-4/3/0.32769
State: Up
```

show dynamic-tunnels database (No Tunnel PIC)

```
user@host> show dynamic-tunnels database
Table: inet.3

Destination-network: 10.255.120.94/32
Tunnel to: 10.255.120.94/32 State: Dn
Reference count: 2
Next-hop type: gre
Source address: 10.255.120.92
State: Dn (no tunnel pic)
```

show dynamic-tunnels database (Tunnel Is Expiring)

```
user@host> show dynamic-tunnels database
Table: inet.3

Destination-network: 10.255.120.94/32
Tunnel to: 10.255.120.94/32 State: Up (expires in 00:14:56 seconds)
Reference count: 0
Next-hop type: gre
Source address: 10.255.120.92
Next hop: gr-4/3/0.32769
State: Up
```

show dynamic-tunnels database (Destination Specified)

```
user@host> show dynamic-tunnels database 10.255.120.94
Table: inet.3

Destination-network: 10.255.120.94/32
Tunnel to: 10.255.120.94/32 State: Up
Reference count: 2
Next-hop type: gre
Source address: 10.255.120.92
Next hop: gr-4/3/0.32769
State: Up
```

show hfrf profiles

| | |
|---------------------------------|---|
| Syntax | show hfrf profiles <brief extensive> |
| Release Information | Command introduced in Junos OS Release 12.2. |
| Description | <p>Display Host fast reroute (HFRR) profile information.</p> <p>HFRR adds a precomputed protection path into the Packet Forwarding Engine (PFE), such that if a link between a provider edge device and a server farm becomes unusable for forwarding, the PFE can use another path without having to wait for the router or the protocols to provide updated forwarding information.</p> |
| Options | <p>none—Display information about HFRR profiles.</p> <p>brief extensive—(Optional) Display the specified level of output.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> Example: Configuring Host Fast Reroute |
| List of Sample Output | show hfrf profiles on page 955 |
| Output Fields | <p>Table 224 on page 954 describes the output fields for the show hfrf profiles command. Output fields are listed in the approximate order in which they appear.</p> |

Table 224: show hfrf profiles Output Fields

| Field Name | Field Description |
|---|---|
| HFRR pointer | |
| HFRR current state | Status of the HFRR profile: HFRR_ACTIVE , HFRR_INACTIVE , HFRR_IFLH-NOT-CONF , and so on. |
| HFRR Prefix limit blackout timer expiry (in secs) | Time interval between a HFRR profile becoming inactive on exceeding the ARP prefix limit, and the profile starting the SYNC process. |
| HFRR prefix limit hit count | Number of times that a HFRR profile becomes inactive on exceeding the ARP prefix limit. |
| HFRR protected IFL name | Interface configured for the HFRR feature. |
| HFRR protected IFL handle | |
| HFRR routing instance name | The routing instance in which the HFRR interface is configured. |
| HFRR routing instance handle | |

Table 224: show hfrr profiles Output Fields (*continued*)

| Field Name | Field Description |
|---|--|
| HFRR sync BG scheduled | |
| HFRR RTS filter on | |
| HFRR delete BG scheduled | |
| HFRR ARP prefix limit | Configured ARP prefix limit. |
| HFRR ARP supplementary blackout timeout (in mins) | Supplementary time-out value configured for profile to be inactive when it hits ARP prefix limit |
| HFRR number of ARP routes learned | Number of ARP routes learned on the configured interface. |
| HFRR number of FRR routes created | Number of ARP routes learned on the configured interface. |

Sample Output

show hfrr profiles

```

user@host> show hfrr profiles
HFRR pointer: 0x9254000
HFRR current state: HFRR_ACTIVE
HFRR Prefix limit blackout timer expiry (in secs): 0
HFRR prefix limit hit count: 0
HFRR protected IFL name: ge-4/1/0.0
HFRR protected IFL handle: 0x9248738
HFRR routing instance name: test
HFRR routing instance handle: 0x9145740
HFRR sync BG scheduled: NO
HFRR RTS filter on: YES
HFRR delete BG scheduled: NO
HFRR ARP prefix limit: 0
HFRR ARP supplementary blackout timeout (in mins): 1
HFRR number of ARP routes learned: 4
HFRR number of FRR routes created: 2

```

show ingress-replication mvpn

| | |
|---------------------------------|---|
| Syntax | show ingress-replication mvpn |
| Release Information | Command introduced in Junos OS Release 10.4. |
| Description | Display the state and configuration of the ingress replication tunnels created for the MVPN application when using the mpls-internet-multicast routing instance type. |
| Required Privilege Level | View |
| List of Sample Output | show ingress-replication mvpn on page 956 |
| Output Fields | Table 225 on page 956 lists the output fields for the show ingress-replication mvpn command. Output fields are listed in the approximate order in which they appear. |

Table 225: show ingress-replication mvpn

| Field Name | Field Description |
|------------------------|--|
| Ingress tunnel | Identifies the MVPN ingress replication tunnel. |
| Application | Identifies the application (MVPN). |
| Unicast tunnels | List of unicast tunnels in use. |
| Leaf address | Address of the tunnel. |
| Tunnel type | Identifies the unicast tunnel type. |
| Mode | Indicates whether the tunnel was created as a new tunnel for the ingress replication, or if an existing tunnel was used. |
| State | Indicates whether the tunnel is Up or Down. |

Sample Output

show ingress-replication mvpn

```

user@host> show ingress-replication mvpn
Ingress Tunnel: mvpn:1
  Application: MVPN
  Unicast tunnels
    Leaf Address      Tunnel-type    Mode    State
    10.255.245.2      P2P LSP       New     Up
    10.255.245.4      P2P LSP       New     Up
Ingress Tunnel: mvpn:2
  Application: MVPN
  Unicast tunnels
    Leaf Address      Tunnel-type    Mode    State
    10.255.245.2      P2P LSP       Existing Up

```

show l2circuit connections

| | |
|---------------------------------|--|
| Syntax | <pre>show l2circuit connections <brief extensive summary> <down up up-down> <history> <interface <i>interface-name</i>> <logical-system (all <i>logical-system-name</i>)> <neighbor <i>neighbor</i>> <status></pre> |
| Release Information | Command introduced before Junos OS Release 7.4. Display enhancements in Junos OS Release 9.6. Display enhancements in Junos OS Release 10.2. Display enhancements in Junos OS Release 12.1. |
| Description | Display status information about Layer 2 virtual circuits from the local provider edge (PE) router to its neighbors. |
| Options | <p>none—Display standard information about Layer 2 virtual circuits on all interfaces for all neighbors.</p> <p>brief extensive summary—(Optional) Display the specified level of output. Use history to display information about connection history. Use status to display information about the connection and interface status.</p> <p>down up up-down—(Optional) Display nonoperational, operational, or both kinds of connections.</p> <p>history—(Optional) Display information about connection history.</p> <p>interface <i>interface-name</i>—(Optional) Show all Layer 2 virtual circuits on an interface.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>neighbor <i>neighbor</i>—(Optional) IP address of a specific neighbor.</p> <p>status—(Optional) Display information about the connection and interface status.</p> |
| Required Privilege Level | view |
| List of Sample Output | <p>show l2circuit connections on page 961</p> <p>show l2circuit connections interface on page 961</p> <p>show l2circuit connections extensive on page 962</p> |
| Output Fields | <p>Table 226 on page 958 lists the output fields for the show l2circuit connections command. Output fields are listed in the approximate order in which they appear.</p> |

Table 226: show l2circuit connections Output Fields

| Field Name | Field Description |
|--|---|
| Layer-2 Circuit Connections | Displays the legends for connection and interface status. |
| Neighbor | Remote PE neighbor. |
| Interface | Logical PE-to-CE interface on which the virtual circuit is configured. |
| Type | VC type: rmt (remote) or loc (local). |
| Legend for connection status (St) | <p>Status of the virtual circuit connection:</p> <ul style="list-style-type: none"> • EI—The local virtual circuit interface is configured with an encapsulation that is not supported. • MM—The two routers do not agree on an MTU value, which causes an MTU mismatch. • EM—The encapsulation type received on this virtual circuit from the neighbor does not match the local virtual circuit interface encapsulation type. • CM—The two routers do not agree on a control word, which causes a control word mismatch. • VM—The remote and local VLAN IDs do not match across the Layer 2 circuit. • OL—No advertisement has been received for this virtual circuit from the neighbor. There is no outgoing label available for use by this virtual circuit. • NC—The interface is not configured as a CCC or TCC interface. • BK—The virtual circuit has switched to a backup connection. • CB—The remote PE router is advertising a different cell bundle from that configured on the local PE router. • LD—The connection to the local site is signaled down, because the CE-facing interface to the local site is down. • RD—The remote neighbor is down. It has signaled a problem using the pseudowire status code. • NP—The router detects that interface hardware is not present. The hardware may be offline, a PIC may not be of the desired type, or the interface may be configured in a different routing instance. • Dn—The virtual circuit is down. • VC-Dn—The virtual circuit is down because there is no tunnel LSP from the local PE router to the neighbor. • UP—The virtual circuit is operational. • CF—The router cannot find enough bandwidth to the remote router to satisfy the Layer 2 circuit bandwidth requirement. • IB—The bit rate is incompatible for Time Division Multiplexing (TDM). • TDM—TDM is not configured correctly. • ST—The virtual circuit has been switched to a standby connection. • SP—The virtual circuit connection is using a static pseudowire. • RS—The remote site is in a standby state. • XX—The virtual circuit is down for an unknown reason. This is a programming error. |
| Time last up | Date and time the virtual circuit was last operational. |

Table 226: show l2circuit connections Output Fields (*continued*)

| Field Name | Field Description |
|-----------------------------|---|
| # Up trans | Number of times the virtual circuit came up. |
| <i>local-interface-name</i> | Name of the local PE-to-CE interface. |
| Status | Status of the local interface. |
| Up | Interface is operational. |
| Dn | Interface is not operational. |
| NP | Not present. Interface does not exist. |
| DS | Disabled. Interface has been administratively disabled. |
| WE | Wrong encapsulation. The interface is not configured as CCC. |
| UN | Interface status is initialized. |
| Encapsulation | Encapsulation of the local interface. |
| Remote PE | Prefix of the remote PE router. |
| Negotiated control-word | Whether the use of the control word has been negotiated for this virtual circuit: Yes (Null) or No . |
| Incoming label | Label used by the remote side of the virtual circuit to send packets destined to the local side. This label is routed to the local virtual circuit interface. |
| Outgoing label | Label used by the local side of the virtual circuit to send packets to the remote side of the virtual circuit. Packets originated on the local virtual circuit interface are encapsulated with this label before being placed on the tunnel LSP to the neighbor for this virtual circuit. This label is allocated by the neighbor and is used in demultiplexing incoming packets destined for this virtual circuit. |
| Negotiated PW status TLV | Displays the pseudowire status type, length, and value (TLV). TLVs are a method of encoding variable-length or optional information. If the pseudowire status TLV is used, the corresponding local or neighbor PE router status code is also displayed. |
| local PW status code | If the pseudowire status TLV is used, displays the local PE router status code. |
| Neighbor PW status code | If the pseudowire status TLV is used, displays the neighbor PE router status code. |
| Local interface | Name of the local interface used for the Layer 2 circuit connection. |
| Status | Status of the local interface (Up or Down). |

Table 226: show l2circuit connections Output Fields (*continued*)

| Field Name | Field Description |
|------------------------------|---|
| Encapsulation | Encapsulation configured for the local interface. |
| APS-active | Indicates that the interface belongs to the working circuit. |
| APS-inactive | Indicates that the interface belongs to the protect circuit. |
| Connection protection | Whether or not connection protection is configured for the Layer 2 circuit to the neighbor: Yes or No . |
| VC bandwidth | Bandwidth requirement of the Layer 2 circuit. |
| Time | Time at which the event occurred. |
| Connection History | <p>Event types logged in history.</p> <ul style="list-style-type: none"> • loc intf up—Local virtual circuit interface went up. • loc intf down—Local virtual circuit interface went down. • In lbl Update—Incoming label has been updated. • Out lbl Update—Outgoing label has been updated. • PE route changed—Route to PE router has been updated. • PE route down—Route to PE router is down. • rmt side marked—Remote side is marked. • VC Dn—Remote side indicated that its end of the virtual circuit is down (if the tunnel LSP from the remote side to the local side is down). • status update timer—Status update timer processing. It computes the state of the virtual circuit, and determines whether it should be advertised to or withdrawn from the remote side. |

Sample Output

**show l2circuit
connections**

user@host> show l2circuit connections

Layer-2 Circuit Connections:

Legend for connection status (St)

| | |
|---------------------------------|--------------------------------------|
| EI -- encapsulation invalid | NP -- interface h/w not present |
| MM -- mtu mismatch | Dn -- down |
| EM -- encapsulation mismatch | VC-Dn -- Virtual circuit Down |
| CM -- control-word mismatch | Up -- operational |
| VM -- vlan id mismatch | CF -- Call admission control failure |
| OL -- no outgoing label | IB -- TDM incompatible bitrate |
| NC -- intf encaps not CCC/TCC | TM -- TDM misconfiguration |
| BK -- Backup Connection | ST -- Standby Connection |
| CB -- rcvd cell-bundle size bad | SP -- Static Pseudowire |
| LD -- local site signaled down | RS -- remote site standby |
| RD -- remote site signaled down | XX -- unknown |

Legend for interface status

Up -- operational

Dn -- down

Neighbor: 10.255.245.51

| Interface | Type | St | Time last up | # Up trans |
|--|------|----|---------------------|------------|
| ge-2/0/2.600(vc 5) | rmt | Up | Dec 7 18:11:18 2009 | 1 |
| Remote PE: 10.255.245.51, Negotiated control-word: No | | | | |
| Incoming label: 299856, Outgoing label: 299808 | | | | |
| Negotiated PW status TLV: No | | | | |
| Local interface: ge-2/0/2.600, Status: Up, Encapsulation: VLAN | | | | |

Sample Output

**show l2circuit
connections interface**

user@host> show l2circuit connections interface t1-2/0/0:1:1.0

Layer-2 Circuit Connections:

Legend for connection status (St)

| | |
|---------------------------------|--------------------------------------|
| EI -- encapsulation invalid | NP -- interface h/w not present |
| MM -- mtu mismatch | Dn -- down |
| EM -- encapsulation mismatch | VC-Dn -- Virtual circuit Down |
| CM -- control-word mismatch | Up -- operational |
| VM -- vlan id mismatch | CF -- Call admission control failure |
| OL -- no outgoing label | IB -- TDM incompatible bitrate |
| NC -- intf encaps not CCC/TCC | TM -- TDM misconfiguration |
| BK -- Backup Connection | ST -- Standby Connection |
| CB -- rcvd cell-bundle size bad | SP -- Static Pseudowire |
| LD -- local site signaled down | RS -- remote site standby |
| RD -- remote site signaled down | XX -- unknown |

Legend for interface status

Up -- operational

Dn -- down

Neighbor: 10.1.1.1

| Interface | Type | St | Time last up | # Up trans |
|---|------|----|----------------------|------------|
| t1-2/0/0:1:1.0(vc 1)(SP) | rmt | Up | Apr 27 04:21:02 2011 | 1 |
| Remote PE: 10.1.1.1, Negotiated control-word: Yes (Non-null) | | | | |
| Incoming label: 1010001, Outgoing label: 1000001 | | | | |
| Negotiated PW status TLV: No | | | | |
| Local interface: t1-1/0/0:1:1.0, Status: Up, Encapsulation: SATOP-T1, | | | | |
| APS-active | | | | |
| Local interface: t1-2/0/0:1:1.0, Status: Up, Encapsulation: SATOP-T1, | | | | |

APS-inactive

Sample Output

**show l2circuit
connections extensive**

user@host>show l2circuit connections extensive
Layer-2 Circuit Connections:

Legend for connection status (St)

| | |
|---------------------------------|--------------------------------------|
| EI -- encapsulation invalid | NP -- interface h/w not present |
| MM -- mtu mismatch | Dn -- down |
| EM -- encapsulation mismatch | VC-Dn -- Virtual circuit Down |
| CM -- control-word mismatch | Up -- operational |
| VM -- vlan id mismatch | CF -- Call admission control failure |
| OL -- no outgoing label | IB -- TDM incompatible bitrate |
| NC -- intf encaps not CCC/TCC | TM -- TDM misconfiguration |
| BK -- Backup Connection | ST -- Standby Connection |
| CB -- rcvd cell-bundle size bad | SP -- Static Pseudowire |
| LD -- local site signaled down | RS -- remote site standby |
| RD -- remote site signaled down | XX -- unknown |

Legend for interface status

Up -- operational
Dn -- down

Neighbor: 10.255.49.149

| Interface | Type | St | Time last up | # Up trans |
|---------------|------|----|----------------------|------------|
| ae0.0(vc 100) | rmt | Up | Aug 31 09:36:12 2009 | 1 |

Remote PE: 10.255.49.149, Negotiated control-word: Yes (Null)

Incoming label: 299824, Outgoing label: 299776

Negotiated PW status TLV: Yes

Local PW status code: 0x00000000, Neighbor PW status code: 0x00000000

Local interface: ae0.0, Status: Up, Encapsulation: ETHERNET

Connection protection: Yes

Connection History:

| | | |
|----------------------|---------------------|--------|
| Aug 31 09:36:12 2009 | status update timer | |
| Aug 31 09:36:12 2009 | PE route changed | |
| Aug 31 09:36:12 2009 | Out lbl Update | 299776 |
| Aug 31 09:36:12 2009 | In lbl Update | 299824 |
| Aug 31 09:36:12 2009 | loc intf up | ae0.0 |

show l2vpn connections

| | |
|---------------------------------|---|
| Syntax | <pre>show l2vpn connections <brief extensive> <down up up-down> <history> <instance <i>instance</i>> <local-site <i>local-site</i>> <logical-system (all <i>logical-system-name</i>)> <remote-site <i>remote-site</i>> <status> <summary></pre> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | Display Layer 2 virtual private network (VPN) connections. |
| Options | <p>none—Display all Layer 2 VPN connections for all routing instances.</p> <p>brief extensive—(Optional) Display the specified level of output.</p> <p>down up up-down—(Optional) Display nonoperational, operational, or both kinds of connections.</p> <p>history—(Optional) Display information about connection history.</p> <p>instance <i>instance</i>—(Optional) Display connections for the specified routing instance only.</p> <p>local-site <i>local-site</i>—(Optional) Display connections for the specified Layer 2 VPN local site name or ID only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>remote-site <i>remote-site</i>—(Optional) Display connection for the specified Layer 2 VPN remote site ID only.</p> <p>status—(Optional) Display information about the connection and interface status.</p> <p>summary—(Optional) Display summary of all Layer 2 VPN connections information.</p> |
| Required Privilege Level | view |
| List of Sample Output | <p>show l2vpn connections on page 967</p> <p>show l2vpn connections extensive on page 967</p> |
| Output Fields | <p>Table 227 on page 964 lists the output fields for the show l2vpn connections command. Output fields are listed in the approximate order in which they appear.</p> |

Table 227: show l2vpn connections Output Fields

| Field Name | Field Description |
|------------------------|---|
| Instance | Name of Layer 2 VPN instance. |
| Local site | Name of local site. |
| Interface name | Name of interface. |
| Remote Site ID | Remote site ID. |
| Label Offset | Numbers within the label block that are skipped to find the next label base. |
| Label-base | Advertises the first label in a block of labels. A remote PE router uses this first label when sending traffic toward the advertising PE router. |
| Range | Advertises the label block size. |
| status-vector | Bit vector advertising the state of local PE-CE circuits to remote PE routers. A bit value of 0 indicates that the local circuit and LSP tunnel to the remote PE router are up, whereas a value of 1 indicates either one or both are down. |
| connection-site | Name of the connection site. |
| Type | Type of connection: loc (local) or rmt (remote). |
| St | Status of the connection. (For a list of possible values, see the Legend for connection status (St) field.) |
| Time last up | Time that the connection was last in the Up condition. |
| # Up trans | Number of transitions from Down to Up condition. |
| Local circuit | Address and status of local circuit. |
| Remote circuit | Address and status of remote circuit. |
| St | <p>Status of the Layer 2 VPN connection (corresponds with Legend for Connection Status):</p> <ul style="list-style-type: none"> EI—The local Layer 2 VPN interface is configured with an encapsulation that is not supported. EM—The encapsulation type received on this Layer 2 VPN connection from the neighbor does not match the local Layer 2 VPN connection interface encapsulation type. VC-Dn—The virtual circuit is currently down. CM—The two routers do not agree on a control word, which causes a control word mismatch. CN—The virtual circuit is not provisioned properly. OR—The label associated with the virtual circuit is out of range. |

Table 227: show l2vpn connections Output Fields (*continued*)

| Field Name | Field Description |
|-----------------------|---|
| | <ul style="list-style-type: none"> • OL—No advertisement has been received for this virtual circuit from the neighbor. There is no outgoing label available for use by this virtual circuit. • LD—All of the CE-facing interfaces to the local site are down. Therefore, the connection to the local site is signaled as down to the other PE routers. No pseudowires can be established. • RD—All the interfaces to the remote neighbor are down. Therefore, the remote site has been signaled as down to the other PE routers. No pseudowires can be established. • LN—The local site has lost path selection to the remote site and therefore no pseudowires can be established from this local site. • RN—The remote site has lost path selection to a local site or other remote site and therefore no pseudowires are established to this remote site. • XX—The Layer 2 VPN connection is down for an unknown reason. This is a programming error. • NC—The interface encapsulation is not configured as an appropriate CCC, TCC, or Layer 2 VPN encapsulation. • WE—The encapsulation configured for the interface does not match the encapsulation configured for the associated connection within the Layer 2 VPN routing instance. • NP—The router detects that interface hardware is not present. The hardware might be offline, a PIC might not be of the desired type, or the interface might be configured in a different routing instance. • ->—Only the outbound connection is up. • <-—Only the inbound connection is up. • Up—The Layer 2 VPN connection is operational. • Dn—The Layer 2 VPN connection is down. • CF—The router cannot find enough bandwidth to the remote router to satisfy the Layer 2 VPN connection bandwidth requirement. • SC—The local site identifier matches the remote site identifier. No pseudowire can be established between these two sites. You should configure different values for the local and remote site identifiers. • LM—The local site identifier is not the minimum designated, meaning it is not the lowest. There is another local site with a lower site identifier. Pseudowires are not being established to this local site, and the associated local site identifier is not being used to distribute Layer 2 VPN label blocks. However, this is not an error state. Traffic continues to be forwarded to the PE router interfaces connected to the local sites when the local sites are in this state. • RM—The remote site identifier is not the minimum designated, meaning it is not the lowest. There is another remote site connected to the same PE router which has lower site identifier. The PE router cannot establish a pseudowire to this remote site and the associated remote site identifier cannot be used to distribute VPLS label blocks. However, this is not an error state. Traffic can continue to be forwarded to the PE router interface connected to this remote site when the remote site is in this state. • IL—The incoming packets for the Layer 2 VPN connection have no MPLS label. |
| Remote PE | Address of the remote provider edge router. |
| Incoming label | Name of the incoming label. |

Table 227: show l2vpn connections Output Fields (*continued*)

| Field Name | Field Description |
|-------------------------|--|
| Outgoing label | Name of the outgoing label. |
| Time | Date and time of Layer 2 VPN connection event. |
| Event | Type of event. |
| Interface/Lbl/PE | Interface, label, or PE router. |

Sample Output

show l2vpn connections

```

user@host> show l2vpn connections
L2VPN Connections :
Instance : vpn-a
Local site: 2 (ce-2)
offset: 1, range: 3, label-base: 32768
  connection-site      Type  St  Time last up      # Up trans
  3 (3)                loc   Up   Jul 18 20:45:46 2001      1
    Local circuit: fe-0/0/0.1, Status: Up
    Remote circuit: fe-0/0/3.0, Status: Up
  1                    rmt   Up   Jul 18 21:47:25 2001      1
    Local circuit: fe-0/0/0.0, Status: Up
    Remote PE: 192.168.16.1
    Incoming label: 32768, Outgoing label: 32769
Local site: 3 (ce-3)
offset: 1, range: 2, label-base: 33792
  connection-site      Type  St  Time last up      # Up trans
  2 (ce-b)             loc   Up   Jul 18 20:45:46 2001      1
    Local circuit: fe-0/0/0.1, Status: Up
    Remote circuit: fe-0/0/3.0, Status: Up
  1                    rmt   Up   Jul 18 21:47:25 2001      1
    Local circuit: fe-0/0/3.1, Status: Up
    Remote PE: 192.168.16.1
    Incoming label: 33792, Outgoing label: 32770

```

show l2vpn connections extensive

```

user@host> show l2vpn connections extensive
L2VPN Connections:

Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch     WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down   NP -- interface hardware not present
CM -- control-word mismatch     -> -- only outbound connection is up
CN -- circuit not provisioned   <- -- only inbound connection is up
OR -- out of range              Up -- operational
OL -- no outgoing label         Dn -- down
LD -- local site signaled down  CF -- call admission control failure
RD -- remote site signaled down SC -- local and remote site ID collision
LN -- local site not designated LM -- local site ID not minimum designated
RN -- remote site not designated RM -- remote site ID not minimum designated
XX -- unknown connection status IL -- no incoming label

```

```

Instance: vpn-a
Local site: ce-a (1)
  Interface name      Remote Site ID
  fe-0/0/0.0          2
  Label Offset      Offset      Range
  32768              1          2
  connection-site      Type  St  Time last up      # Up trans
  2                    rmt   Up   Aug 3 00:08:14 2001      1
    Local circuit: fe-0/0/0.0, Status: Up
    Remote PE: 192.168.24.1
    Incoming label: 32769, Outgoing label: 32768
      Time          Event          Interface/Lbl/PE
      Aug 3 00:08:14 2001 PE route up
      Aug 3 00:08:14 2001 Out lbl Update      32768
      Aug 3 00:08:14 2001 In lbl Update      32769
      Aug 3 00:08:14 2001 ckt0 up             fe-0/0/0.0

```


show mvpn c-multicast

| | |
|---------------------------------|---|
| Syntax | show mvpn c-multicast <extensive summary> <instance-name <i>instance-name</i> > |
| Release Information | Command introduced in Junos OS Release 8.4. |
| Description | Display the multicast VPN customer multicast route information. |
| Options | extensive summary —(Optional) Display the specified level of output. instance-name <i>instance-name</i> —(Optional) Display output for the specified routing instance. |
| Required Privilege Level | view |
| List of Sample Output | show mvpn c-multicast on page 970 show mvpn c-multicast summary on page 970 show mvpn c-multicast extensive on page 970 |
| Output Fields | Table 228 on page 969 lists the output fields for the show mvpn c-multicast command. Output fields are listed in the approximate order in which they appear. |

Table 228: show mvpn c-multicast Output Fields

| Field Name | Field Description | Level of Output |
|------------------------------|---|------------------------|
| Instance | Name of the VPN routing instance. | summary extensive none |
| C-mcast IPv4 (S:G) | Customer router IPv4 multicast address. | extensive none |
| Ptnl | Provider tunnel attributes, <i>tunnel type:tunnel source, tunnel destination group</i> . | extensive none |
| St | State: <ul style="list-style-type: none"> • DS—Represents (S,G) and is created due to (*,G) • RM—Remote VPN route learned from the remote PE router • St display blank—SSM group join | extensive none |
| MVPN instance | Name of the multicast VPN routing instance | extensive none |
| C-multicast IPv4 route count | Number of customer multicast IPv4 routes associated with the multicast VPN routing instance. | summary |
| C-multicast IPv6 route count | Number of customer multicast IPv6 routes associated with the multicast VPN routing instance. | summary |

Sample Output

show mvpn c-multicast

```

user@host> show mvpn c-multicast
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
  C-mcast IPv4 (S:G)          Ptnl          St
  192.168.195.78/32:225.5.5.5/32 PIM-SM:10.255.14.144, 239.1.1.1      RM
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-B
  C-mcast IPv4 (S:G)          Ptnl          St
  192.168.195.94/32:226.6.6.6/32 PIM-SM:10.255.14.144, 239.2.0.0      RM

```

show mvpn c-multicast summary

```

user@host> show mvpn c-multicast summary
MVPN Summary:
Instance: VPN-A
  C-multicast IPv4 route count: 1
Instance: VPN-B
  C-multicast IPv4 route count: 2

```

show mvpn c-multicast extensive

```

user@host> show mvpn c-multicast extensive
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
  C-mcast IPv4 (S:G)          Ptnl          St
  192.168.195.78/32:225.5.5.5/32 PIM-SM:10.255.14.144, 239.1.1.1      RM
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-B
  C-mcast IPv4 (S:G)          Ptnl          St
  192.168.195.94/32:226.6.6.6/32 PIM-SM:10.255.14.144, 239.2.0.0      RM

```

show mvpn instance

| | |
|---------------------------------|---|
| Syntax | show mvpn instance <extensive summary> <instance <i>instance-name</i> > |
| Release Information | Command introduced in Junos OS Release 8.4. |
| Description | Display the multicast VPN routing instance information. |
| Options | extensive summary —(Optional) Display the specified level of output. instance <i>instance-name</i> —(Optional) Display statistics for the specified routing instance. |
| Required Privilege Level | view |
| List of Sample Output | show mvpn instance on page 973 show mvpn instance on page 973 show mvpn instance summary on page 973 show mvpn instance extensive on page 974 show mvpn instance summary (IPv6) on page 974 |
| Output Fields | Table 229 on page 971 lists the output fields for the show mvpn instance command. Output fields are listed in the approximate order in which they appear. |

Table 229: show mvpn instance Output Fields

| Field Name | Field Description | Level of Output |
|---------------------------|---|--------------------------------------|
| MVPN instance | Name of the multicast VPN routing instance | extensive none |
| Instance | Name of the VPN routing instance. | summary extensive none |
| Provider tunnel | Provider tunnel attributes, <i>tunnel type:tunnel source, tunnel destination group</i> . | extensive none |
| Neighbor | Address, type of provider tunnel (I-P-tnl , inclusive provider tunnel and S-P-tnl , selective provider tunnel) and provider tunnel for each neighbor. | extensive none |
| C-mcast IPv4 (S:G) | Customer IPv4 router multicast address. | extensive none |
| C-mcast IPv6 (S:G) | Customer IPv6 router multicast address. | extensive none |
| Ptnl | Provider tunnel attributes, <i>tunnel type:tunnel source, tunnel destination group</i> . | extensive none |
| St | State: <ul style="list-style-type: none"> DS—Represents (S,G) and is created due to (*,G) RM—Remote VPN route learned from the remote PE router St display blank—SSM group join | extensive none |

Table 229: show mvpn instance Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-------------------------------------|--|-----------------|
| Neighbor count | Number of neighbors associated with the multicast VPN routing instance. | summary |
| C-multicast IPv4 route count | Number of customer multicast IPv4 routes associated with the multicast VPN routing instance. | summary |
| C-multicast IPv6 route count | Number of customer multicast IPv6 routes associated with the multicast VPN routing instance. | summary |

Sample Output

```

show mvpn instance      user@host> show mvpn instance
                        MVPN instance:

                        Legend for provider tunnel
                        I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

                        Legend for c-multicast routes properties (Pr)
                        DS -- derived from (*, c-g)          RM -- remote VPN route
                        Instance: VPN-A
                        Provider tunnel: I-P-tnl:PIM-SM:10.255.14.144, 239.1.1.1
                        Neighbor          I-P-tnl
                        10.255.14.160      PIM-SM:10.255.14.160, 239.1.1.1
                        10.255.70.17       PIM-SM:10.255.70.17, 239.1.1.1
                        C-mcast IPv4 (S:G) Ptnl              St
                        192.168.195.78/32:225.5.5.5/32 PIM-SM:10.255.14.144, 239.1.1.1      RM
                        MVPN instance:

                        Legend for provider tunnel
                        I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

                        Legend for c-multicast routes properties (Pr)
                        DS -- derived from (*, c-g)          RM -- remote VPN route
                        Instance: VPN-B
                        Provider tunnel: I-P-tnl:PIM-SM:10.255.14.144, 239.2.0.0
                        Neighbor          I-P-tnl
                        10.255.14.160      PIM-SM:10.255.14.160, 239.2.0.0
                        10.255.70.17       PIM-SM:10.255.70.17, 239.2.0.0
                        C-mcast IPv4 (S:G) Ptnl              St
                        192.168.195.94/32:226.6.6.6/32 PIM-SM:10.255.14.144, 239.2.0.0      RM

```

```

show mvpn instance      user@host> show mvpn instance
                        MVPN instance:

                        Legend for provider tunnel
                        I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

                        Legend for c-multicast routes properties (Pr)
                        DS -- derived from (*, c-g)          RM -- remote VPN route
                        Instance : vpn-1
                        MVPN Mode : SPT-ONLY
                        Provider tunnel: I-P-tnl:LDP-P2MP:10.255.72.162, lsp-id 16777217
                        Neighbor          I-P-tnl
                        10.255.72.160      LDP-P2MP:10.255.72.160, lsp-id 16777217
                        10.255.72.166      RSVP-TE P2MP:10.255.72.166,
                        13054,10.255.72.166
                        10.255.72.168

```

Sample Output

```

show mvpn instance      user@host> show mvpn instance summary
summary
                        MVPN Summary:
                        Instance: VPN-A
                        Neighbor count: 2
                        C-multicast IPv4 route count: 1
                        Instance: VPN-B
                        Neighbor count: 4

```

C-multicast IPv4 route count: 2

Sample Output

**show mvpn instance
extensive**

```
user@host> show mvpn instance extensive
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
  Provider tunnel: I-P-tnl:PIM-SM:10.255.14.144, 239.1.1.1
  Neighbor                    I-P-tnl
  10.255.14.160                PIM-SM:10.255.14.160, 239.1.1.1
  10.255.70.17                 PIM-SM:10.255.70.17, 239.1.1.1
  C-mcast IPv4 (S:G)          Ptnl                    St
  192.168.195.78/32:225.5.5.5/32 PIM-SM:10.255.14.144, 239.1.1.1      RM
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-B
  Provider tunnel: I-P-tnl:PIM-SM:10.255.14.144, 239.2.0.0
  Neighbor                    I-P-tnl
  10.255.14.160                PIM-SM:10.255.14.160, 239.2.0.0
  10.255.70.17                 PIM-SM:10.255.70.17, 239.2.0.0
  C-mcast IPv4 (S:G)          Ptnl                    St
  192.168.195.94/32:226.6.6.6/32 PIM-SM:10.255.14.144, 239.2.0.0      RM
```

**show mvpn instance
summary (IPv6)**

```
user@host> show mvpn instance summary
MVPN Summary:
Instance: VPN-A
  C-multicast IPv6 route count: 2
Instance: VPN-B
  C-multicast IPv6 route count: 2
```


show mvpn neighbor

| | |
|---------------------------------|---|
| Syntax | show mvpn neighbor <extensive summary> <inet inet6> <instance <i>instance-name</i> neighbor-address <i>address</i> > <logical-system <i>logical-system-name</i> > |
| Release Information | Command introduced in Junos OS Release 8.4. |
| Description | Display multicast VPN neighbor information. |
| Options | <p>extensive summary—(Optional) Display the specified level of output for all multicast VPN neighbors.</p> <p>inet inet6—(Optional) Display IPv4 or IPv6 information for all multicast VPN neighbors.</p> <p>instance <i>instance-name</i> neighbor-address <i>address</i>—(Optional) Display multicast VPN neighbor information for the specified instance or the specified neighbor.</p> <p>logical-system <i>logical-system-name</i>—(Optional) Display multicast VPN neighbor information for the specified logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show mvpn neighbor on page 976 show mvpn neighbor extensive on page 976 show mvpn neighbor extensive on page 976 show mvpn neighbor instance-name on page 977 show mvpn neighbor neighbor-address on page 977 show mvpn neighbor neighbor-address summary on page 977 show mvpn neighbor neighbor-address extensive on page 977 show mvpn neighbor neighbor-address instance-name on page 978 |
| Output Fields | Table 230 on page 975 lists the output fields for the show mvpn neighbor command. Output fields are listed in the approximate order in which they appear. |

Table 230: show mvpn neighbor Output Fields

| Field Name | Field Description | Level of Output |
|------------------------|---|------------------------|
| MVPN instance | Name of the multicast VPN routing instance | extensive none |
| Instance | Name of the VPN routing instance. | summary extensive none |
| Neighbor | Address, type of provider tunnel (I-P-tnl, inclusive provider tunnel and S-P-tnl, selective provider tunnel) and provider tunnel for each neighbor. | extensive none |
| Provider tunnel | Provider tunnel attributes, <i>tunnel type:tunnel source, tunnel destination group</i> . | extensive none |

Sample Output

```

user@host> show mvpn neighbor
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)      RM -- remote VPN route
Instance: VPN-A
Neighbor                          I-P-tnl
10.255.14.160                     PIM-SM:10.255.14.160, 239.1.1.1
10.255.70.17                     PIM-SM:10.255.70.17, 239.1.1.1
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)      RM -- remote VPN route
Instance: VPN-B
Neighbor                          I-P-tnl
10.255.14.160                     PIM-SM:10.255.14.160, 239.2.0.0
10.255.70.17                     PIM-SM:10.255.70.17, 239.2.0.0

```

Sample Output

```

user@host> show mvpn neighbor extensive
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)      RM -- remote VPN route
Instance: VPN-A
Neighbor                          I-P-tnl
10.255.14.160                     PIM-SM:10.255.14.160, 239.1.1.1
10.255.70.17                     PIM-SM:10.255.70.17, 239.1.1.1
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)      RM -- remote VPN route
Instance: VPN-B
Neighbor                          I-P-tnl
10.255.14.160                     PIM-SM:10.255.14.160, 239.2.0.0
10.255.70.17                     PIM-SM:10.255.70.17, 239.2.0.0

user@host> show mvpn neighbor extensive
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

```

```

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: mvpn-a
Neighbor                             I-P-tnl
10.255.72.45
10.255.72.50                         LDP P2MP:10.255.72.50, lsp-id 1

```

Sample Output

**show mvpn neighbor
instance-name**

```

user@host> show mvpn neighbor instance-name VPN-A
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
Neighbor                             I-P-tnl
10.255.14.160                        PIM-SM:10.255.14.160, 239.1.1.1
10.255.70.17                         PIM-SM:10.255.70.17, 239.1.1.1

```

Sample Output

**show mvpn neighbor
neighbor-address**

```

user@host> show mvpn neighbor neighbor-address 10.255.14.160
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
Neighbor                             I-P-tnl
10.255.14.160                        PIM-SM:10.255.14.160, 239.1.1.1
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-B
Neighbor                             I-P-tnl
10.255.14.160                        PIM-SM:10.255.14.160, 239.2.0.0

```

Sample Output

**show mvpn neighbor
neighbor-address
summary**

```

user@host> show mvpn neighbor neighbor-address 10.255.70.17 summary
MVPN Summary:
Instance: VPN-A
Instance: VPN-B

```

Sample Output

```

user@host> show mvpn neighbor neighbor-address 10.255.70.17 extensive

```

show mvpn neighbor neighbor-address extensive

```

MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
  Neighbor                          I-P-tnl
  10.255.70.17                      PIM-SM:10.255.70.17, 239.1.1.1
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-B
  Neighbor                          I-P-tnl
  10.255.70.17                      PIM-SM:10.255.70.17, 239.2.0.0

```

Sample Output

show mvpn neighbor neighbor-address instance-name

```

user@host> show mvpn neighbor neighbor-address 10.255.70.17 instance-name VPN-A
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
  Neighbor                          I-P-tnl
  10.255.70.17                      PIM-SM:10.255.70.17, 239.1.1.1

```

show vpls connections

| | |
|---------------------------------|--|
| Syntax | <pre>show vpls connections <brief extensive> <down up up-down> <history> <instance <i>instance-name</i> local-site <i>local-site-name</i> remote-site <i>remote-site-name</i>> <logical-system (all <i>logical-system-name</i>)> <status> <summary></pre> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | (T Series and M Series routers, except for the M160 router) Display virtual private LAN service (VPLS) connection information. |
| Options | <p>none—Display information about all VPLS connections for all routing instances.</p> <p>brief extensive—(Optional) Display the specified level of output.</p> <p>down up up-down—(Optional) Display nonoperational, operational, or both types of connections.</p> <p>history—(Optional) Display information about connection history.</p> <p>instance <i>instance-name</i>—(Optional) Display the VPLS connections for the specified routing instance only.</p> <p>local-site <i>local-site-name</i>—(Optional) Display the VPLS connections for the specified local site name or ID only.</p> <p>remote-site <i>remote-site-name</i>—(Optional) Display the VPLS connections for the specified remote site name or ID only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>status—(Optional) Display information about the connection and interface status.</p> <p>summary—(Optional) Display summary of all VPLS connections information.</p> |
| Required Privilege Level | view |
| List of Sample Output | <p>show vpls connections on page 986</p> <p>show vpls connections extensive (Static VPLS Neighbors) on page 988</p> |
| Output Fields | <p>Table 231 on page 980 lists the output fields for the show vpls connections command. Output fields are listed in the approximate order in which they appear.</p> |

Table 231: show vpls connections Output Fields

| Field Name | Field Description |
|--------------------------------------|--|
| Instance | Name of the VPLS instance. |
| Local site | Name of the local site. |
| VPLS-id | Identifier for the VPLS site. |
| Number of local interfaces | Number of interfaces configured for the local site. |
| Number of local interfaces up | Number of interfaces configured for the local site that are currently up. |
| IRB interface present | Indicates whether or not an integrated routing and bridging (IRB) interface is present (yes or no). |
| Intf | <p>List of all of the interfaces configured for the local site. The types of interfaces can include VPLS virtual loopback tunnel interfaces and label-switched interfaces. Any interface that supports VPLS could be listed here.</p> <p>Virtual loopback tunnel interfaces are displayed using the vt-fpc/pic/port.nnnnn format. Label-switched interfaces are displayed using the lsi.nnnnn format. In both cases, nnnnn is a dynamically generated virtual port used to transport and receive packets from other provider edge (PE) routers in the VPLS domain.</p> <p>Each interface might include the following information:</p> <ul style="list-style-type: none"> • Identification as a VPLS interface • Name of the associated VPLS routing instance • Local site number • Remote site number • VPLS neighbor address • VPLS identifier |
| Interface flags | <p>Flag associated with the interface. Can include the following:</p> <ul style="list-style-type: none"> • VC-Down—The virtual circuit associated with this interface is down. |
| Label-base | First label in a block of labels. A remote PE router uses this first label when sending traffic toward the advertising PE router. |
| Offset | Displays the VPLS Edge (VE) block offset in the Layer 2 VPN NLRI. The VE block offset is used to identify a label block from which a particular label value is selected to setup a pseduowire for a remote site. The block offset value itself indicates the starting VE ID that maps to the label base contained in the VPLS NLRI advertisement. |
| Size | Label block size. |
| Range | Label block range. |

Table 231: show vpls connections Output Fields (*continued*)

| Field Name | Field Description |
|------------------------|---|
| Preference | Preference value advertised for a VPLS site. When multiple PE routers are assigned the same VE ID for multihoming, you might need to specify that a particular PE router acts as the designated forwarder by configuring the site preference value. The site preference indicates the degree of preference for a particular customer site. The site preference is one of the tie-breaking criteria used in a designated forwarder election. |
| status-vector | Bit vector advertising the state of local PE-CE circuits to remote PE routers. A bit value of 0 indicates that the local circuit and LSP tunnel to the remote PE router are up, whereas a value of 1 indicates either one or both are down. |
| connection-site | Name of the connection site. |
| Neighbor | IP address and VPLS identifier for the VPLS neighbor. |
| Type | Type of connection: loc (local) or rmt (remote). |

Table 231: show vpls connections Output Fields (*continued*)

| Field Name | Field Description |
|------------|-------------------|
| St | |

Table 231: show vpls connections Output Fields (*continued*)

| Field Name | Field Description |
|------------|---|
| | <p>Status of the VPLS connection (corresponds with Legend for Connection Status):</p> <ul style="list-style-type: none"> • EI—The local VPLS interface is configured with an encapsulation that is not supported. • EM—The encapsulation type received on this VPLS connection from the neighbor does not match the local VPLS connection interface encapsulation type. • VC-Dn—The virtual circuit is currently down. • CM—The two routers do not agree on a control word, which causes a control word mismatch. • CN—The virtual circuit is not provisioned properly. • OR—The label associated with the virtual circuit is out of range. • OL—No advertisement has been received for this virtual circuit from the neighbor. There is no outgoing label available for use by this virtual circuit. • LD—All of the CE-facing interfaces to the local site are down. Therefore, the connection to the local site is signaled as down to the other PE routers. No pseudowires can be established. • RD—All the interfaces to the remote neighbor are down. Therefore, the remote site has been signaled as down to the other PE routers. No pseudowires can be established. • LN—The local site has lost path selection to the remote site and therefore no pseudowires can be established from this local site. • RN—The remote site has lost path selection to a local site or other remote site and therefore no pseudowires are established to this remote site. • XX—The VPLS connection is down for an unknown reason. This is a programming error. • MM—The MTU for the local site and the remote site do not match. • BK—The router is using a backup connection. • PF—Profile parse failure. • RS—The remote site is in a standby state. • NC—The interface encapsulation is not configured as an appropriate CCC, TCC, or VPLS encapsulation. • WE—The encapsulation configured for the interface does not match the encapsulation configured for the associated connection within the VPLS routing instance. • NP—The router detects that interface hardware is not present. The hardware might be offline, a PIC might not be of the desired type, or the interface might be configured in a different routing instance. • ->—Only the outbound connection is up. • <-—Only the inbound connection is up. • Up—The VPLS connection is operational. • Dn—The VPLS connection is down. • CF—The router cannot find enough bandwidth to the remote router to satisfy the VPLS connection bandwidth requirement. • SC—The local site identifier matches the remote site identifier. No pseudowire can be established between these two sites. You should configure different values for the local and remote site identifiers. • LM—The local site identifier is not the minimum designated, meaning it is not |

Table 231: show vpls connections Output Fields (*continued*)

| Field Name | Field Description |
|---------------------------------|---|
| | <p>the lowest. There is another local site with a lower site identifier. Pseudowires are not being established to this local site, and the associated local site identifier is not being used to distribute VPLS label blocks. However, this is not an error state. Traffic continues to be forwarded to the PE router interfaces connected to the local sites when the local sites are in this state.</p> <ul style="list-style-type: none"> • RM—The remote site identifier is not the minimum designated, meaning it is not the lowest. There is another remote site connected to the same PE router which has lower site identifier. The PE router cannot establish a pseudowire to this remote site and the associated remote site identifier cannot be used to distribute VPLS label blocks. However, this is not an error state. Traffic can continue to be forwarded to the PE router interface connected to this remote site when the remote site is in this state. • IL—The incoming packets for the VPLS connection have no MPLS label. • MI—The configured mesh group identifier is in use by another system in the network. • ST—The router has switched to a standby connection. • PB—Profile busy. • SN—The VPLS neighbor is static. |
| Time last up | Time connection was last in the Up condition. |
| # Up trans | Number of transitions from Down to Up condition. |
| Status | <p>Status of the (local or remote circuit) local interface:</p> <ul style="list-style-type: none"> • Up—Operational • Dn—Down • NP—Not present • DS—Disabled • WE—Wrong encapsulation • UN—Uninitialized |
| Encapsulation | Type of encapsulation: VPLS . |
| Remote PE | Address of the remote provider edge router. |
| Negotiated control-word | Whether a control word has been negotiated: Yes or No . |
| Incoming label | Name of the incoming label. |
| Outgoing label | Name of the outgoing label. |
| Negotiated PW status TLV | Indicates whether or not the pseudowire status TLV has been negotiated for the VPLS connection. |

Table 231: show vpls connections Output Fields (*continued*)

| Field Name | Field Description |
|---------------------------|--|
| Local interface | Provides the following information about the local interface configured for the VPLS neighbor: <ul style="list-style-type: none"> • Name of the local interface • Status—Interface status (Up or Down) • Encapsulation—Interface encapsulation (for example, ETHERNET) • Description—Includes the VPLS instance name, the VPLS neighbor address, and the VPLS identifier |
| Time | Date and time of VPLS connection event. |
| Event | Type of event. |
| Interface/Lbl/PE | Interface, label, or PE router. |
| Connection History | Each entry can include the date, time, year, and the connection event. Connection events include any of a variety of events related to VPLS connections, such as route changes, label updates, and interfaces going down or coming up. |

Sample Output

```

user@host> show vpls connections
Layer-2 VPN connections:

Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch     WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down   NP -- interface hardware not present
CM -- control-word mismatch     -< -- only outbound connection is up
CN -- circuit not provisioned    >- -- only inbound connection is up
OR -- out of range              Up -- operational
OL -- no outgoing label         Dn -- down
LD -- local site signaled down  CF -- call admission control failure
RD -- remote site signaled down SC -- local and remote site ID collision
LN -- local site not designated LM -- local site ID not minimum designated
RN -- remote site not designated RM -- remote site ID not minimum designated
XX -- unn connection status     IL -- no incoming label
MM -- MTU mismatch              MI -- Mesh-Group ID not availble
BK -- Backup connection         ST -- Standby connection
PF -- Profile parse failure     PB -- Profile busy

Legend for interface status
Up -- operational
Dn -- down

Instance: vpls-1
Local site: 1 (11)
  Number of local interfaces: 1
  Number of local interfaces up: 1
  IRB interface present: no
  lt-1/3/0.10496
  vt-1/3/0.1048588    1      Intf - vpls vpls-1 local site 11 remote site 1
  vt-1/2/0.1048591    2      Intf - vpls vpls-1 local site 11 remote site 2
  vt-1/2/0.1048585    3      Intf - vpls vpls-1 local site 11 remote site 3
  vt-1/2/0.1048587    4      Intf - vpls vpls-1 local site 11 remote site 4
  vt-1/2/0.1048589    5      Intf - vpls vpls-1 local site 11 remote site 5
  vt-1/3/0.1048586    6      Intf - vpls vpls-1 local site 11 remote site 6
  vt-1/3/0.1048590    7      Intf - vpls vpls-1 local site 11 remote site 7
  vt-1/3/0.1048584    8      Intf - vpls vpls-1 local site 11 remote site 8

  Label-base      Offset      Size      Range      Preference
+ 800256          1          16        16          100

Timer Values:
  Startup wait time: 120 seconds
  New site wait-time: 20 seconds
  Collision detect time: 30 seconds
  Reclaim wait time: 748 milliseconds

connection-site      Type      St      Time last up      # Up trans
1                    rmt      Up      Apr 28 13:28:24 2009      2
  Remote PE: 124.1.2.1, Negotiated control-word: No
  Incoming label: 800256, Outgoing label: 800026
  Local interface: vt-1/3/0.1048588, Status: Up, Encapsulation: VPLS

```

```

Description: Intf - vpls vpls-1 local site 11 remote site 1
Connection History:
Apr 28 13:28:24 2009 status update timer
Apr 28 13:28:24 2009 PE route down
Apr 28 13:24:27 2009 status update timer
Apr 28 13:24:27 2009 loc intf up vt-1/3/0.1048588
Apr 28 13:24:27 2009 PE route changed
Apr 28 13:24:27 2009 Out lbl Update 800026
Apr 28 13:24:27 2009 In lbl Update 800256
Apr 28 13:24:27 2009 loc intf down
2 rmt Up Apr 28 13:28:24 2009 2
Remote PE: 124.1.7.1, Negotiated control-word: No
Incoming label: 800257, Outgoing label: 800034
Local interface: vt-1/2/0.1048591, Status: Up, Encapsulation: VPLS
Description: Intf - vpls vpls-1 local site 11 remote site 2
Connection History:
Apr 28 13:28:24 2009 status update timer
Apr 28 13:28:24 2009 PE route down
Apr 28 13:24:28 2009 status update timer
Apr 28 13:24:28 2009 loc intf up vt-1/2/0.1048591
Apr 28 13:24:28 2009 PE route changed
Apr 28 13:24:28 2009 Out lbl Update 800034
Apr 28 13:24:28 2009 In lbl Update 800257
Apr 28 13:24:28 2009 loc intf down
3 rmt Up Apr 28 13:28:24 2009 2
Remote PE: 124.1.4.1, Negotiated control-word: No
Incoming label: 800258, Outgoing label: 800026
Local interface: vt-1/2/0.1048585, Status: Up, Encapsulation: VPLS
Description: Intf - vpls vpls-1 local site 11 remote site 3
Connection History:
Apr 28 13:28:24 2009 status update timer
Apr 28 13:28:24 2009 PE route down
Apr 28 13:24:26 2009 status update timer
Apr 28 13:24:26 2009 loc intf up vt-1/2/0.1048585
Apr 28 13:24:26 2009 PE route changed
Apr 28 13:24:26 2009 Out lbl Update 800026
Apr 28 13:24:26 2009 In lbl Update 800258
Apr 28 13:24:26 2009 loc intf down
4 rmt Up Apr 28 13:28:24 2009 2
Remote PE: 124.1.6.1, Negotiated control-word: No
Incoming label: 800259, Outgoing label: 800026
Local interface: vt-1/2/0.1048587, Status: Up, Encapsulation: VPLS
Description: Intf - vpls vpls-1 local site 11 remote site 4
Connection History:
Apr 28 13:28:24 2009 status update timer
Apr 28 13:28:24 2009 PE route down
Apr 28 13:24:27 2009 status update timer
Apr 28 13:24:27 2009 loc intf up vt-1/2/0.1048587
Apr 28 13:24:27 2009 PE route changed
Apr 28 13:24:27 2009 Out lbl Update 800026
Apr 28 13:24:27 2009 In lbl Update 800259
Apr 28 13:24:27 2009 loc intf down
5 rmt Up Apr 28 13:28:24 2009 2
Remote PE: 124.1.3.1, Negotiated control-word: No
Incoming label: 800260, Outgoing label: 800034
Local interface: vt-1/2/0.1048589, Status: Up, Encapsulation: VPLS
Description: Intf - vpls vpls-1 local site 11 remote site 5
Connection History:
Apr 28 13:28:24 2009 status update timer
Apr 28 13:28:24 2009 PE route down
Apr 28 13:24:28 2009 status update timer

```

```

Apr 28 13:24:28 2009  loc intf up                vt-1/2/0.1048589
Apr 28 13:24:28 2009  PE route changed
Apr 28 13:24:28 2009  Out lbl Update                800034
Apr 28 13:24:27 2009  In lbl Update                800260
Apr 28 13:24:27 2009  loc intf down

```

show vpls connections extensive (Static VPLS Neighbors)

```

user@host> show vpls connections extensive instance red
Layer-2 VPN connections:

```

Legend for connection status (St)

```

EI -- encapsulation invalid      NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch     WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down   NP -- interface hardware not present
CM -- control-word mismatch     -> -- only outbound connection is up
CN -- circuit not provisioned   <- -- only inbound connection is up
OR -- out of range              Up -- operational
OL -- no outgoing label         Dn -- down
LD -- local site signaled down  CF -- call admission control failure
RD -- remote site signaled down SC -- local and remote site ID collision
LN -- local site not designated LM -- local site ID not minimum designated
RN -- remote site not designated RM -- remote site ID not minimum designated
XX -- unn connection status    IL -- no incoming label
MM -- MTU mismatch             MI -- Mesh-Group ID not availble
BK -- Backup connection        ST -- Standby connection
PF -- Profile parse failure     PB -- Profile busy
RS -- remote site standby       SN -- Static Neighbor

```

Legend for interface status

```

Up -- operational
Dn -- down

```

Instance: static

```

VPLS-id: 1
Number of local interfaces: 1
Number of local interfaces up: 1
ge-0/0/5.0
lsi.1049344                Intf - vpls static neighbor 10.255.114.3 vpls-id
1
Neighbor                    Type  St    Time last up      # Up trans
10.255.114.3(vpls-id 1)(SN) rmt Up    Mar  4 08:48:41 2010      1
Remote PE: 10.255.114.3, Negotiated control-word: No
Incoming label: 29696, Outgoing label: 29697
Negotiated PW status TLV: No
Local interface: lsi.1049344, Status: Up, Encapsulation: ETHERNET
Description: Intf - vpls static neighbor 10.255.114.3 vpls-id 1
Connection History:
Mar  4 08:48:41 2010  status update timer
Mar  4 08:48:41 2010  PE route changed
Mar  4 08:48:41 2010  Out lbl Update                29697
Mar  4 08:48:41 2010  In lbl Update                29696
Mar  4 08:48:41 2010  loc intf up                lsi.1049344

```

```

user@PE1> show vpls connections extensive (Multihoming with FEC 129)
Layer-2 VPN connections:

```

Legend for connection status (St)

```

EI -- encapsulation invalid      NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch     WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down   NP -- interface hardware not present
CM -- control-word mismatch     -> -- only outbound connection is up

```

CN -- circuit not provisioned <- -- only inbound connection is up
 OR -- out of range Up -- operational
 OL -- no outgoing label Dn -- down
 LD -- local site signaled down CF -- call admission control failure
 RD -- remote site signaled down SC -- local and remote site ID collision
 LN -- local site not designated LM -- local site ID not minimum designated
 RN -- remote site not designated RM -- remote site ID not minimum designated
 XX -- unknown connection status IL -- no incoming label
 MM -- MTU mismatch MI -- Mesh-Group ID not available
 BK -- Backup connection ST -- Standby connection
 PF -- Profile parse failure PB -- Profile busy
 RS -- remote site standby SN -- Static Neighbor
 LB -- Local site not best-site RB -- Remote site not best-site
 VM -- VLAN ID mismatch

Legend for interface status

Up -- operational
 Dn -- down

Instance: green

L2vpn-id: 100:100

Local-id: 1.1.1.2

Number of local interfaces: 2

Number of local interfaces up: 2

ge-0/3/1.0

ge-0/3/3.0

lsi.101711873

Intf - vpls green local-id 1.1.1.2 remote-id

1.1.1.4 neighbor 1.1.1.4

| Remote-id | Type | St | Time last up | # Up trans |
|-----------|------|----|----------------------|------------|
| 1.1.1.4 | rmt | Up | Jan 31 13:49:52 2012 | 1 |

Remote PE: 1.1.1.4, Negotiated control-word: No

Incoming label: 262146, Outgoing label: 262146

Local interface: lsi.101711873, Status: Up, Encapsulation: ETHERNET

Description: Intf - vpls green local-id 1.1.1.2 remote-id 1.1.1.4 neighbor

1.1.1.4

Connection History:

| | |
|----------------------|---------------------------|
| Jan 31 13:49:52 2012 | status update timer |
| Jan 31 13:49:52 2012 | PE route changed |
| Jan 31 13:49:52 2012 | Out lbl Update 262146 |
| Jan 31 13:49:52 2012 | In lbl Update 262146 |
| Jan 31 13:49:52 2012 | loc intf up lsi.101711873 |

Multi-home:

| Local-site | Id | Pref | State |
|------------|----|------|-------|
| test | 1 | 100 | Up |

Number of interfaces: 1

Number of interfaces up: 1

ge-0/3/1.0

Received multi-homing advertisements:

| Remote-PE | Pref | flag | Description |
|-----------|------|------|-------------|
| 1.1.1.4 | 100 | 0x0 | |

show vpls flood event-queue

| | |
|---------------------------------|---|
| Syntax | show vpls flood event-queue |
| Release Information | Command introduced in Junos OS Release 8.0. |
| Description | Display the pending events in the VPLS flood queue. |
| Options | This command has no options. |
| Required Privilege Level | view |
| List of Sample Output | show vpls flood event-queue on page 991 |
| Output Fields | Table 232 on page 990 lists the output fields for the show vpls flood event-queue command. Output fields are listed in the approximate order in which they appear. |

Table 232: show vpls flood event-queue Output Fields

| Field Name | Field Description |
|------------------------------|--|
| Current Pending Event | Provides information on the current event in the VPLS flood event queue. |
| Name | Name of the event. |
| Owner Name | Name of the interface associated with the flood event. |
| Pending Op | Pending operation for the event. |
| Last Error | Name of the last error encountered. |
| Number of Retries | Number of attempts made to update the event queue. |
| Pending Event List | List of the events awaiting processing. |
| Event Name | Name of the event. |
| Pending Op | Pending operation for the event. |
| Event Identifier | Name of the interface associated with the flood event. |

Sample Output

`show vpls flood
event-queue`

```
user@host> show vpls flood event-queue
Current Pending Event
  Name:          Flood Nexthop
  Owner Name:ge-4/3/0.0
  Pending Op: ADD
  Last Error:ENOMEM
  Number of Retries:3
  Pending Event List:
  Event Name      Pending Op      Event Identifier
  Flood Nexthop   ADD          ge-4/3/0.0
  Flood Route     ADD          ge-4/3/0.0
```

show vpls flood instance

| | |
|---------------------------------|--|
| Syntax | show vpls flood instance <brief detail extensive> <instance-name> <logical-system <i>logical-system-name</i> > |
| Release Information | Command introduced in Junos OS Release 8.0. |
| Description | Display VPLS information related to the flood process. |
| Options | <p>none—Display VPLS information related to the flood process for all routing instances.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>instance-name—(Optional) Display VPLS information related to the flood process for the specified routing instance.</p> <p>logical-system <i>logical-system-name</i>—(Optional) Display VPLS information related to the flood process for the specified logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show vpls flood instance on page 993 show vpls flood instance logical-system-name on page 993 show vpls flood instance detail on page 993 |
| Output Fields | Table 233 on page 992 lists the output fields for the show vpls flood instance command. Output fields are listed in the approximate order in which they appear. |

Table 233: show vpls flood instance Output Fields

| Field Name | Field Description |
|-----------------------|---|
| Logical system | Name of the logical system. |
| Name | Name of the VPLS routing instance. |
| CEs | Number of CE routers connected to the VPLS instance. |
| VEs | Number of VE routers connected to the VPLS instance. |
| Flood routes | List of all flood routes associated with the VPLS instance. |
| Prefix | Prefix for the route. |
| Type | Type of route. |
| Owner | VPLS routing instance or interface associated with the route. |
| Nh type | Next-hop type. For example, flood for a flood route. |

Table 233: show vpls flood instance Output Fields (*continued*)

| Field Name | Field Description |
|------------|--------------------------------------|
| Nhindex | Next-hop index number for the route. |

Sample Output

show vpls flood instance

```
user@host> show vpls flood instance

Logical system: __juniper_ls1__
Name: green
CEs: 1
VEs: 1
Flood Routes:
  Prefix    Type           Owner           NhType    NhIndex
  default   ALL_CE_FLOOD   green           flood      383
  0x47/16   CE_FLOOD       fe-1/2/1.0     flood      388
```

show vpls flood instance logical-system-name

```
user@host: __juniper_ls1__> show vpls flood instance juniper_ls1

Logical system: __juniper_ls1__
Name: green
CEs: 1
VEs: 1
Flood Routes:
  Prefix    Type           Owner           NhType    NhIndex
  default   ALL_CE_FLOOD   green           flood      383
  0x47/16   CE_FLOOD       fe-1/2/1.0     flood      388
```

show vpls flood instance detail

```
user@host: __juniper_ls1__> show vpls flood instance detail

Logical system: __juniper_ls1__
Name: green
CEs: 1
VEs: 1
Flood Routes:
  Prefix    Type           Owner           NhType    NhIndex
  default   ALL_CE_FLOOD   green           flood      383
  0x47/16   CE_FLOOD       fe-1/2/1.0     flood      388
```

show vpls flood route

| | |
|---------------------------------|---|
| Syntax | show vpls flood route (all-ce-flood instance-name <i>instance-name</i> <logical-system-name <i>logical-system-name</i> > ce-flood interface <i>interface-name</i>) |
| Release Information | Command introduced in Junos OS Release 8.0. |
| Description | Display VPLS route information related to the flood process for either the specified routing instance or the specified interface. |
| Options | <p>all-ce-flood—Display the flood next-hop route for all customer edge routers for traffic coming from the core of the network.</p> <p>ce-flood interface <i>interface-name</i>—Display the flood next-hop route for traffic coming from the specified customer edge interface.</p> <p>instance-name <i>instance-name</i>—Display the flood routes for the specified instance.</p> <p>logical-system-name <i>logical-system-name</i>—(Optional) Specify the logical system whose flood routes you want to display. You can only specify the default logical system name for VPLS. The default logical system name is __juniper_ls1__ (the name must be entered in the command with the underscore characters).</p> |
| Required Privilege Level | view |
| List of Sample Output | show vpls flood route all-ce-flood on page 995 show vpls flood route ce-flood on page 995 |
| Output Fields | Table 234 on page 994 lists the output for the show vpls flood route command. Output fields are listed in the approximate order in which they appear. |

Table 234: show vpls flood route Output Fields

| Field Name | Field Description |
|------------------------|---|
| Flood route prefix | Prefix for the flood route. |
| Flood route type | Type of flood route (either CE_FLOOD or ALL_CE_FLOOD). |
| Flood route owner | VPLS routing instance or interface associated with the flood route. |
| Nexthop type | Next-hop type. For example, flood for a flood route. |
| Nexthop index | Next-hop index number for the route. |
| Interfaces flooding to | Interfaces to which VPLS routes are being flooded. |
| Name | Name of the interface. |

Table 234: show vpls flood route Output Fields (*continued*)

| Field Name | Field Description |
|------------|-----------------------------------|
| Type | Type of VPLS router (CE or VE). |
| Nh type | Next-hop type. |
| Index | Index number for the flood route. |

Sample Output

**show vpls flood route
all-ce-flood**

```
user@host: __juniper_lsl__> show vpls flood route all-ce-flood logical-system-name
__juniper_lsl__instance-name green
```

```
Flood route prefix: default
Flood route type: ALL_CE_FLOOD
Flood route owner: green
Nexthop type: flood
Nexthop index: 383
  Interfaces Flooding to:
    Name          Type          NhType          Index
    fe-1/2/1.0    CE
```

**show vpls flood route
ce-flood**

```
user@host: __juniper_lsl__> show vpls flood route ce-flood interface fe-1/2/1.0
```

```
Flood route prefix: 0x47/16
Flood route type: CE_FLOOD
Flood route owner: fe-1/2/1.0
Nexthop type: flood
Nexthop index: 388
  Interfaces Flooding to:
    Name          Type          NhType          Index
    lsi.49152      VE            indr            262142
```

show vpls mac-table

| | |
|---------------------------------|---|
| Syntax | <pre>show vpls mac-table <brief detail extensive summary> <bridge-domain <i>bridge-domain-name</i>> <instance <i>instance-name</i>> <interface <i>interface-name</i>> <logical-system (all <i>logical-system-name</i>)> <mac-address> <vlan-id <i>vlan-id-number</i>></pre> |
| Release Information | Command introduced in Junos OS Release 8.5. |
| Description | (MX960 routers only) Display learned VPLS MAC address information. |
| Options | <p>none—Display all learned VPLS MAC address information.</p> <p>brief detail extensive summary—(Optional) Display the specified level of output.</p> <p>bridge-domain <i>bridge-domain-name</i>—(Optional) Display learned VPLS MAC addresses for the specified bridge domain.</p> <p>instance <i>instance-name</i>—(Optional) Display learned VPLS MAC addresses for the specified instance.</p> <p>interface <i>interface-name</i>—(Optional) Display learned VPLS MAC addresses for the specified instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Display learned VPLS MAC addresses for all logical systems or for the specified logical system.</p> <p>mac-address—(Optional) Display the specified learned VPLS MAC address information..</p> <p>vlan-id <i>vlan-id-number</i>—(Optional) Display learned VPLS MAC addresses for the specified VLAN.</p> |
| Required Privilege Level | view |
| List of Sample Output | <p>show vpls mac-table on page 998</p> <p>show vpls mac-table count on page 998</p> <p>show vpls mac-table detail on page 999</p> <p>show vpls mac-table extensive on page 999</p> |
| Output Fields | <p>Table 235 on page 996 describes the output fields for the show bridge mac-table command. Output fields are listed in the approximate order in which they appear.</p> |

Table 235: show vpls mac-table Output fields

| Field Name | Field Description |
|------------------|-------------------------------|
| Routing instance | Name of the routing instance. |

Table 235: show vpls mac-table Output fields (*continued*)

| Field Name | Field Description |
|---------------------------|--|
| Bridging domain | Name of the bridging domain. |
| MAC address | MAC address or addresses learned on a logical interface. |
| MAC flags | Status of MAC address learning properties for each interface: <ul style="list-style-type: none"> • S—Static MAC address configured. • D—Dynamic MAC address learned. • SE—MAC accounting is enabled. • NM—Nonconfigured MAC. |
| Logical interface | Name of the logical interface. |
| MAC count | Number of MAC addresses learned on a specific routing instance or interface. |
| Learning interface | Logical interface or logical Label Switched Interface (LSI) the address is learned on. |
| Learn VLAN ID/VLAN | VLAN ID of the routing instance or bridge domain in which the MAC address was learned. |
| Layer 2 flags | Debugging flags signifying that the MAC address is present in various lists. |
| Epoch | Spanning Tree Protocol epoch number identifying when the MAC address was learned. Used for debugging. |
| Sequence number | Sequence number assigned to this MAC address. Used for debugging. |
| Learning mask | Mask of Packet Forwarding Engines where this MAC address was learned. Used for debugging. |
| IPC generation | Creation time of the logical interface when this MAC address was learned. Used for debugging. |

Sample Output

show vpls mac-table

```

user@host> show vpls mac-table
MAC flags (S -static MAC, D -dynamic MAC,
           SE -Statistics enabled, NM -Non configured MAC)

Routing instance : vpls_ldp1
VLAN : 223
  MAC          MAC      Logical
  address      flags    interface
  00:90:69:9c:1c:5d  D      ge-0/2/5.400

MAC flags (S -static MAC, D -dynamic MAC,
           SE -Statistics enabled, NM -Non configured MAC)

Routing instance : vpls_red
VLAN : 401
  MAC          MAC      Logical
  address      flags    interface
  00:00:aa:12:12:12  D      lsi.1051138
  00:05:85:74:9f:f0  D      lsi.1051138

```

show vpls mac-table count

```

user@host> show vpls mac-table count
0 MAC address learned in routing instance __juniper_private1__

MAC address count per interface within routing instance:
  Logical interface      MAC count
  1c-0/0/0.32769         0
  1c-0/1/0.32769         0
  1c-0/2/0.32769         0
  1c-2/0/0.32769         0
  1c-0/3/0.32769         0
  1c-2/1/0.32769         0
  1c-9/0/0.32769         0
  1c-11/0/0.32769        0
  1c-2/2/0.32769         0
  1c-9/1/0.32769         0
  1c-11/1/0.32769        0
  1c-2/3/0.32769         0
  1c-9/2/0.32769         0
  1c-11/2/0.32769        0
  1c-11/3/0.32769        0
  1c-9/3/0.32769         0

MAC address count per learn VLAN within routing instance:
  Learn VLAN ID      MAC count
  0                   0

1 MAC address learned in routing instance vpls_ldp1

MAC address count per interface within routing instance:
  Logical interface      MAC count
  lsi.1051137            0
  ge-0/2/5.400           1

MAC address count per learn VLAN within routing instance:
  Learn VLAN ID      MAC count
  0                   1

```


1 MAC address learned in routing instance vpls_red

MAC address count per interface within routing instance:

| Logical interface | MAC count |
|-------------------|-----------|
| ge-0/2/5.300 | 1 |

MAC address count per learn VLAN within routing instance:

| Learn VLAN ID | MAC count |
|---------------|-----------|
| 0 | 1 |

show vpls mac-table detail

user@host> show vpls mac-table detail

MAC address: 00:90:69:9c:1c:5d
 Routing instance: vpls_ldp1
 Learning interface: ge-0/2/5.400
 Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
 Epoch: 0 Sequence number: 1
 Learning mask: 0x1 IPC generation: 0

MAC address: 00:90:69:9c:1c:5d
 Routing instance: vpls_red
 Learning interface: ge-0/2/5.300
 Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
 Epoch: 0 Sequence number: 1
 Learning mask: 0x1 IPC generation: 0

show vpls mac-table extensive

user@host> show vpls mac-table extensive

MAC address: 00:00:aa:12:12:12
 Routing instance: vpls_ldp1
 Learning interface: lsi.1051137
 Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
 Epoch: 0 Sequence number: 1
 Learning mask: 0x1 IPC generation: 0

MAC address: 00:05:85:74:9f:f0
 Routing instance: vpls_ldp1
 Learning interface: lsi.1051137
 Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
 Epoch: 0 Sequence number: 1
 Learning mask: 0x1 IPC generation: 0

MAC address: 00:90:69:9c:1c:5d
 Routing instance: vpls_ldp1
 Learning interface: ge-0/2/5.400
 Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
 Epoch: 0 Sequence number: 1
 Learning mask: 0x1 IPC generation: 0

MAC address: 00:00:aa:12:12:12
 Routing instance: vpls_red
 Learning interface: lsi.1051138
 Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
 Epoch: 0 Sequence number: 0
 Learning mask: 0x1 IPC generation: 0

MAC address: 00:05:85:74:9f:f0
 Routing instance: vpls_red
 Learning interface: lsi.1051138
 Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
 Epoch: 0 Sequence number: 0
 Learning mask: 0x1 IPC generation: 0

show vpls statistics

| | |
|---------------------------------|---|
| Syntax | show vpls statistics <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before Junos OS Release 7.4. |
| Description | (T Series and M Series routers, except for the M160 router) Display virtual private LAN service (VPLS) statistics. |
| Options | <p>none—Display VPLS statistics for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Display VPLS statistics for a specific VPLS routing instance only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | show vpls statistics on page 1003 show vpls statistics instance on page 1003 |
| Output Fields | Table 236 on page 1001 lists the output fields for the show vpls statistics command. Output fields are listed in the approximate order in which they appear. |

Table 236: show vpls statistics Output Fields

| Field Name | Field Description |
|------------------------------------|---|
| Instance | Name of the VPLS instance. |
| Local interface | Name of the local VPLS virtual loopback tunnel interface, <i>vt-fpc/plc/port.nnnnn</i> , where <i>nnnnn</i> is a dynamically generated virtual port used to transport and receive packets from other provider edge (PE) routers in the VPLS domain. |
| Index | Number associated with the next hop. |
| Remote provider edge router | Address of the remote provider edge router. |
| Multicast packets | Number of multicast packets received. |
| Multicast bytes | Number of multicast bytes received. |
| Flood packets | Number of VPLS flood packets received. |
| Flood bytes | Number of VPLS flood bytes received. |

Table 236: show vpls statistics Output Fields (*continued*)

| Field Name | Field Description |
|--------------------------|---|
| Current MAC count | Number of MAC addresses learned by the interface and the configured maximum limit on the number of MAC addresses that can be learned. |

Sample Output

show vpls statistics

```
user@host> show vpls statistics
```

VPLS statistics:

Instance: green

```
Local interface: fe-2/2/1.0, Index: 69
Multicast packets:      1
Multicast bytes   :      60
Flooded packets   :      18
Flooded bytes    :     2556
Current MAC count:      1
```

```
Local interface: lt-0/3/0.2, Index: 72
Multicast packets:      3
Multicast bytes   :     153
Flooded packets   :      1
Flooded bytes    :      51
Current MAC count:      1
```

```
Local interface: lsi.32769, Index: 75
Current MAC count:      0
```

```
Local interface: lsi.32771, Index: 77
Remote PE: 10.255.14.222
Current MAC count:      2
```

Instance: red

```
Local interface: vt-0/3/0.32768, Index: 74
Multicast packets:      0
Multicast bytes   :      0
Flooded packets   :      0
Flooded bytes    :      0
Current MAC count:      0
```

```
Local interface: vt-0/3/0.32770, Index: 76
Multicast packets:      0
Multicast bytes   :      0
Flooded packets   :      0
Flooded bytes    :      0
Current MAC count:      0
```

show vpls statistics
instance

```
user@host> show vpls statistics instance red
```

Layer-2 VPN Statistics:

Instance: red

```
Local interface: vt-3/2/0.32768, Index: 73
Remote provider edge router: 10.255.17.35
Multicast packets:      0
Multicast bytes   :      0
Flood packets    :      0
Flood bytes     :      0
Current MAC count:      1 (Limit 20)
```


PART 6

Index

- [Index on page 1007](#)
- [Index of Statements and Commands on page 1019](#)

Index

Symbols

| | |
|--|------|
| #, comments in configuration statements..... | xvii |
| (), in syntax descriptions..... | xvii |
| < >, in syntax descriptions..... | xvi |
| [], in configuration statements..... | xvii |
| { }, in configuration statements..... | xvii |
| (pipe), in syntax descriptions..... | xvii |

A

| | |
|---|-----|
| administrative groups <i>See</i> groups | |
| advertisements, displaying | 497 |
| ANCP | |
| CoS state | |
| displaying ANCP..... | 10 |
| local access loop | |
| displaying..... | 19 |
| neighbor | |
| clearing..... | 4 |
| neighbors | |
| displaying..... | 13 |
| request OAM per interface..... | 8 |
| request OAM per neighbor..... | 9 |
| subscriber | |
| clearing..... | 6 |
| subscribers | |
| displaying..... | 19 |
| AS paths | |
| distribution of, displaying..... | 478 |
| domain information, displaying..... | 482 |
| matching regular expressions, displaying..... | 504 |
| summary of, displaying..... | 484 |
| ASN | |
| BGP community routes, displaying..... | 512 |
| automatic bandwidth allocation | |
| LSPs..... | 811 |
| Automatic Multicast Tunneling <i>See</i> AMT | |
| autonomous system number <i>See</i> ASN | |
| autonomous system paths <i>See</i> AS paths | |

B

| | |
|---|------|
| backup PE groups | |
| multicast, displaying..... | 178 |
| bandwidth, allocating for LSPs..... | 811 |
| best routes, displaying..... | 506 |
| BFD | |
| sessions | |
| clearing..... | 27 |
| displaying..... | 28 |
| BGP | |
| community ASN, displaying routes..... | 512 |
| community name, displaying routes..... | 514 |
| damping parameters | |
| clearing..... | 38 |
| displaying..... | 74 |
| damping routes, displaying..... | 516 |
| groups | |
| general information, displaying..... | 44 |
| traffic statistics, displaying..... | 51 |
| neighbors | |
| clearing connections..... | 39 |
| displaying..... | 53 |
| summary information, displaying..... | 69 |
| table | |
| clearing..... | 41 |
| BGP Monitoring Protocol | |
| displaying | |
| statistics..... | 43 |
| Bidirectional Forwarding Detection <i>See</i> BFD | |
| binding state of DHCP client | |
| clearing..... | 732 |
| displaying..... | 745 |
| binding state of DHCPv6 client | |
| clearing..... | 737 |
| displaying..... | 753 |
| bootstrap routers, displaying..... | 230 |
| braces, in configuration statements..... | xvii |
| brackets | |
| angle, in syntax descriptions..... | xvi |
| square, in configuration statements..... | xvii |
| bridge domain | |
| Layer 2, displaying..... | 903 |
| C | |
| CAC | |
| displaying for LSPs..... | 830 |
| call admission control <i>See</i> CAC | |
| CCC | |
| connections, displaying..... | 813 |

| | |
|--|------|
| circuit cross-connect See CCC | |
| clear (ospf ospf3) database command..... | 393 |
| clear (ospf ospf3) database-protection command..... | 396 |
| clear (ospf ospf3) io-statistics command..... | 397 |
| clear (ospf ospf3) neighbor command..... | 398 |
| clear (ospf ospf3) overload command..... | 400 |
| clear (ospf ospf3) statistics command..... | 401 |
| clear ancp neighbor command..... | 4 |
| clear ancp subscriber command..... | 6 |
| clear bfd adaptation command..... | 26 |
| clear bfd session command..... | 27 |
| clear bgp damping command..... | 38 |
| clear bgp neighbor command..... | 39 |
| clear bgp table command..... | 41 |
| clear bridge mac-table command..... | 900 |
| clear dhcp relay binding command..... | 732 |
| clear dhcp relay statistics command..... | 734 |
| clear dhcpv6 relay binding command..... | 737 |
| clear dhcpv6 relay statistics command..... | 741 |
| clear error bpdu command..... | 901 |
| clear error mac-rewrite command..... | 902 |
| clear esis adjacency command..... | 78 |
| clear esis statistics command..... | 79 |
| clear firewall command..... | 712 |
| clear helper statistics command..... | 743 |
| clear igmp membership command..... | 92 |
| clear igmp snooping membership command..... | 95 |
| clear igmp snooping statistics command..... | 96 |
| clear igmp statistics command..... | 97 |
| clear ipv6 neighbors command..... | 302 |
| clear ipv6 router-advertisement command..... | 303 |
| clear isis adjacency command..... | 311 |
| clear isis database command..... | 313 |
| clear isis overload command..... | 315 |
| clear isis statistics command..... | 317 |
| clear ldp neighbor command..... | 777 |
| clear ldp session command..... | 778 |
| clear ldp statistics command..... | 779 |
| clear lldp neighbor command..... | 364 |
| clear lldp statistics command..... | 365 |
| clear mld membership command..... | 99 |
| clear mld statistics command..... | 100 |
| clear msdp cache command..... | 101 |
| clear msdp statistics command..... | 102 |
| clear multicast bandwidth-admission command..... | 103 |
| clear multicast forwarding-cache command..... | 105 |
| clear multicast scope command..... | 106 |
| clear multicast sessions command..... | 107 |
| clear multicast snooping statistics command..... | 108 |
| clear multicast statistics command..... | 109 |
| clear pgm negative-acknowledgments command..... | 110 |
| clear pgm source-path-messages command..... | 111 |
| clear pgm statistics command..... | 112 |
| clear pim join command..... | 113 |
| clear pim join-distribution command..... | 114 |
| clear pim register command..... | 116 |
| clear pim statistics command..... | 122 |
| clear rip general-statistics command..... | 692 |
| clear rip statistics command..... | 693 |
| clear ripng general-statistics command..... | 702 |
| clear ripng statistics command..... | 703 |
| clear rsvp session command..... | 866 |
| clear rsvp statistics command..... | 868 |
| clear spanning-tree protocol-migration command..... | 928 |
| clear spanning-tree statistics command..... | 929 |
| clear vpls mac-address command..... | 949 |
| clear vpls mac-table command..... | 950 |
| comments, in configuration statements..... | xvii |
| community ASN, displaying routes..... | 512 |
| community name, displaying routes..... | 514 |
| conditions routing policy..... | 768 |
| Constrained Shortest Path First See CSPF | |
| conventions text and syntax..... | xvi |
| CSPF statistics, displaying..... | 832 |
| curly braces, in configuration statements..... | xvii |
| customer support..... | xvii |
| contacting JTAC..... | xvii |
| D | |
| damping parameters, BGP clearing..... | 38 |
| displaying..... | 74 |
| damping routes, BGP displaying..... | 516 |

-
- DHCP
 - relay binding
 - clearing.....732
 - relay binding state
 - displaying.....745
 - relay statistics
 - clearing.....734
 - displaying.....750
 - DHCP client
 - binding state
 - clearing.....732
 - displaying.....745
 - statistics
 - clearing.....734
 - DHCPv6
 - relay binding
 - clearing.....737
 - relay binding state
 - displaying.....753
 - relay statistics
 - clearing.....741
 - displaying.....759
 - DHCPv6 client
 - binding state
 - displaying.....753
 - statistics
 - clearing.....741
 - DHCPv6client
 - binding state
 - clearing.....737
 - DiffServ
 - classes, displaying for MPLS.....834
 - Distance Vector Multicast Routing Protocol *See* DVMRP
 - DVMRP
 - documentation
 - comments on.....xvii
 - Draft-rosen MVPNs
 - data MDT cache, displaying.....248
 - MDT tunnels
 - displaying.....250
 - DVMRP
 - groups, displaying.....221
 - interfaces, displaying.....126
 - neighbors, displaying.....128
 - prefixes, displaying.....130
 - prunes, displaying active.....132
 - dynamic overload bit, resetting for IS-IS.....315
 - dynamic tunnel database, displaying.....952
 - E
 - End System-to-Intermediate System *See* ES-IS
 - ES-IS
 - adjacencies
 - clearing.....78
 - displaying.....80
 - interfaces, displaying.....82
 - statistics
 - clearing.....79
 - displaying.....84
 - export route information, displaying.....539
 - F
 - FEC filters
 - displaying for LDP.....784
 - firewall
 - filter version
 - displaying.....721
 - statistics
 - displaying.....714
 - firewall filters
 - log information, displaying.....722
 - policed packets, displaying.....728
 - statistics
 - clearing.....712
 - displaying.....725
 - flooding
 - Layer 2 bridging, displaying.....905
 - font conventions.....xvi
 - forwarding equivalence class *See* FEC
 - forwarding options
 - DHCP relay agent.....732, 734, 745, 750
 - DHCPv6 relay agent.....737, 741, 753, 759
 - UDP statistics
 - clearing.....743
 - displaying.....762
 - forwarding table
 - interfaces, displaying.....574
 - multicast information, displaying.....195
 - multicast snooping information,
 - displaying.....212
 - route entries, displaying.....560
 - G
 - generalized MPLS *See* GMPLS
 - GMPLS
 - link-management information, displaying
 - all.....816
 - peers.....820

| | |
|--|----------|
| routing process..... | 822 |
| statistics..... | 825 |
| traffic-engineered links..... | 827 |
| groups | |
| BGP | |
| general information, displaying..... | 44 |
| traffic statistics, displaying..... | 51 |
| DVMRP, displaying..... | 221 |
| IGMP membership, displaying..... | 134 |
| MLD | |
| clearing..... | 99 |
| displaying..... | 155 |
| MPLS, displaying administrative..... | 829 |
| PIM | |
| general information, displaying..... | 235 |
| usage information, displaying..... | 221 |
| H | |
| hidden routes, displaying..... | 578 |
| hostnames | |
| IS-IS, displaying..... | 342 |
| I | |
| icons defined, notice..... | xv |
| IGMP | |
| group membership, displaying..... | 134 |
| interfaces, displaying..... | 138 |
| PIM-to-IGMP message translation information, | |
| displaying..... | 191 |
| snooping (interface)..... | 142 |
| snooping (membership)..... | 145 |
| snooping (statistics)..... | 149 |
| statistics, displaying..... | 152 |
| interface learning information | |
| Layer 2, displaying..... | 924 |
| interfaces, displaying | |
| in the forwarding table..... | 574 |
| Intermediate System-to-Intermediate System See | |
| IS-IS | |
| Internet Group Management Protocol See IGMP | |
| invalid routes, displaying..... | 600 |
| IP IGMP snooping | |
| membership | |
| clearing..... | 95 |
| statistics | |
| clearing..... | 96 |
| IP multicast..... | 87 |
| announced sessions, displaying..... | 207 |
| backup PE groups, displaying..... | 178 |
| bandwidth admission | |
| clearing..... | 103 |
| flow map information, displaying..... | 180 |
| forwarding cache, clearing..... | 105 |
| forwarding table, displaying..... | 195 |
| forwarding-cache statistics | |
| displaying..... | 182 |
| interface information, displaying..... | 184 |
| network information, displaying..... | 186 |
| next-hop table, displaying..... | 188 |
| PIM-to-IGMP message translation information, | |
| displaying..... | 191 |
| PIM-to-MLD message translation information, | |
| displaying..... | 193 |
| RPF calculations, displaying..... | 201 |
| SAP announcements, displaying..... | 298 |
| scope, clearing..... | 106 |
| scoped information, displaying..... | 205 |
| sessions, clearing..... | 107 |
| statistics | |
| clearing..... | 109 |
| displaying..... | 218 |
| <i>See also</i> DVMRP, MDT, MLD, MSDP, PGM, PIM, | |
| AMT | |
| IP multicast snooping | |
| forwarding table, displaying..... | 212 |
| statistics | |
| clearing..... | 108 |
| displaying..... | 210, 215 |
| IPv6 | |
| neighbor cache information | |
| clearing..... | 302 |
| displaying..... | 304 |
| router advertisements | |
| clearing..... | 303 |
| displaying..... | 306 |
| IS-IS | |
| adjacency database entries, clearing..... | 311 |
| authentication, displaying..... | 323 |
| backup coverage | |
| displaying..... | 325 |
| backup MPLS LSPs..... | 327 |
| backup paths | |
| SPF calculations..... | 329 |
| dynamic overload bit, resetting..... | 315 |
| hostname database, displaying..... | 342 |
| interfaces, displaying..... | 343 |

| | |
|---|----------|
| link-state database entries | |
| clearing..... | 313 |
| displaying..... | 335 |
| neighbors, displaying..... | 319 |
| routes, displaying..... | 351 |
| SPF calculations, displaying..... | 355 |
| traffic statistics | |
| clearing..... | 317 |
| displaying..... | 360 |
| J | |
| join states, clearing PIM..... | 113 |
| join states, redistributing..... | 114 |
| L | |
| Label Distribution Protocol See LDP | |
| label switched path, LSP..... | 807 |
| Layer 2 BPDU loop prevention | |
| clearing errors..... | 901 |
| Layer 2 bridging..... | 899 |
| bridge domain, displaying..... | 903 |
| flooding, displaying..... | 905 |
| interface learning information, displaying..... | 924 |
| learning information, displaying..... | 920, 921 |
| learning properties, displaying..... | 922 |
| MAC address table, clearing..... | 900 |
| MAC address table, displaying..... | 912 |
| statistics, displaying..... | 916 |
| Layer 2 circuits | |
| connections, displaying..... | 957 |
| redundant pseudowires..... | 951 |
| Layer 2 protocol tunneling | |
| clearing errors..... | 902 |
| interface information, displaying..... | 926 |
| Layer 2 switching..... | 899 |
| Layer 2 VPNs | |
| connections, displaying..... | 963 |
| VPLS connections, displaying..... | 979 |
| VPLS statistics, displaying..... | 1001 |
| Layer 3 VPNs | |
| dynamic tunnel database, displaying..... | 952 |
| PIM MDTs, displaying..... | 244 |
| PIM, displaying..... | 252 |
| LDP | |
| database entries, displaying..... | 780 |
| FEC filters, displaying..... | 784 |
| interfaces, displaying..... | 785 |
| LSPs, displaying..... | 789 |
| neighbors | |
| clearing connections..... | 777 |
| displaying..... | 787 |
| routes, displaying..... | 791 |
| sessions | |
| clearing..... | 778 |
| displaying..... | 795 |
| statistics | |
| clearing..... | 779 |
| displaying..... | 801 |
| traffic statistics, displaying..... | 805 |
| learning information | |
| Layer 2, displaying..... | 920, 921 |
| learning properties | |
| Layer 2, displaying..... | 922 |
| link-state paths See LSPs | |
| LLDP | |
| displaying..... | 366 |
| local information, displaying..... | 369 |
| neighbor | |
| clearing..... | 364 |
| neighbors, displaying..... | 371 |
| remote global statistics, displaying..... | 375 |
| statistics | |
| clearing..... | 365 |
| statistics, displaying..... | 377 |
| LSPs | |
| bandwidth allocation, adjusting..... | 811 |
| CAC information, displaying..... | 830 |
| clearing..... | 809 |
| LDP, displaying..... | 789 |
| MPLS, displaying..... | 838 |
| M | |
| MAC address table | |
| Layer 2, clearing address information..... | 900 |
| Layer 2, displaying..... | 912 |
| manuals | |
| comments on..... | xvii |
| martians, displaying..... | 600 |
| MDT | |
| displaying information..... | 244 |
| MDT join TLV | |
| displaying advertisements received..... | 248 |
| MLD | |
| group membership | |
| clearing..... | 99 |
| displaying..... | 155 |
| interfaces, displaying..... | 160 |

| | |
|--|---|
| <ul style="list-style-type: none"> PIM-to-MLD message translation information, displaying.....193 statistics <ul style="list-style-type: none"> clearing.....100 displaying.....164 | <ul style="list-style-type: none"> MVLP <ul style="list-style-type: none"> configuration <ul style="list-style-type: none"> show.....380, 382, 385, 386 dynamic-vlan-memberships <ul style="list-style-type: none"> show.....384 statistics <ul style="list-style-type: none"> show.....388 |
| <ul style="list-style-type: none"> MPLS.....775 <ul style="list-style-type: none"> administrative groups, displaying.....829 CCC connections, displaying.....813 CSPF statistics, displaying.....832 DiffServ classes, displaying.....834 interfaces, displaying.....836 labels, displaying routes.....594 link-management information, displaying <ul style="list-style-type: none"> all.....816 peers.....820 routing process.....822 statistics.....825 traffic-engineered links.....827 LSPs, displaying.....851 static LSPs, displaying.....853 See also LDP, LSPs, RSVP, traffic engineering database MSDP <ul style="list-style-type: none"> cache entries, clearing.....101 general information, displaying.....167 message source information, displaying.....169 peer statistics <ul style="list-style-type: none"> clearing.....102 displaying.....174 source-active cache, displaying.....171 MSTP <ul style="list-style-type: none"> configuration <ul style="list-style-type: none"> displaying.....941 multicast See IP multicast multicast distribution trees See MDT Multicast Listener Discovery See MLD Multicast Source Discovery Protocol See MSDP multicast virtual private network See MVPN multicast VPNs <ul style="list-style-type: none"> customer multicast routes, displaying.....969 neighbors, displaying.....975 routing instances, displaying.....971 Multiprotocol Label Switching See MPLS MVPN.....248, 250 <ul style="list-style-type: none"> displaying information.....252 See also Draft-Rosen MVPNs | <ul style="list-style-type: none"> N <ul style="list-style-type: none"> neighbor <ul style="list-style-type: none"> ANCP.....4 LLDP.....364 next hops <ul style="list-style-type: none"> multicast entries, displaying.....188 resolution database, displaying.....641 routes sent to, displaying.....602 notice icons defined.....xv O <ul style="list-style-type: none"> Open Shortest Path First See OSPF OSPF <ul style="list-style-type: none"> backup coverage <ul style="list-style-type: none"> displaying.....403 backup MPLS LSPs <ul style="list-style-type: none"> displaying.....406 backup neighbor paths.....408 backup paths <ul style="list-style-type: none"> SPF calculations.....410 context identifier, displaying.....418 interfaces, displaying.....440 link-state database entries, displaying <ul style="list-style-type: none"> version 2.....420 version 3.....429 neighbors <ul style="list-style-type: none"> clearing connections.....398 displaying.....452 overview <ul style="list-style-type: none"> displaying.....459 routing table entries, displaying.....464 SPF calculations, displaying.....448 statistics, general <ul style="list-style-type: none"> clearing.....401 displaying.....470 statistics, I/O <ul style="list-style-type: none"> clearing.....397 displaying.....446 overload bit, resetting for IS-IS.....315 |

P

parentheses, in syntax descriptions.....xvii

PGM

negative acknowledgments

clearing.....110

displaying.....224

source path messages

clearing.....111

displaying.....226

statistics

clearing.....112

displaying.....227

PIM

bootstrap routers, displaying.....230

groups

general information, displaying.....235

usage information, displaying.....221

interfaces

displaying.....232

join states, clearing.....113

MDTs, displaying.....244

MVPN, displaying.....252

neighbors, displaying.....253

PIM-to-IGMP message translation information,
displaying.....191

PIM-to-MLD message translation information,
displaying.....193

prune states, clearing.....113

redistributing join states.....114

register

clearing.....116

RPF, displaying source state.....283

RPs

displaying.....258

statistics

clearing.....122

displaying.....285

policers, displaying.....728

Pragmatic General Multicast *See* PGM

Protocol Independent Multicast *See* PIM

prune states, clearing PIM.....113

prunes, DVMRP, displaying.....132

pseudowires

redundant switchover.....951

R

redundant pseudowires, switchover.....951

regular expressions

AS paths, displaying matching routes.....504

IP multicast scope

clearing.....106

IP multicast sessions

clearing.....107

displaying.....207

LSPs, clearing.....809

target communities, displaying.....541

rendezvous points *See* RPs

replication

of BGP configuration.....67

request ancp oam interface command.....8

request ancp oam neighbor command.....9

request l2circuit-switchover command.....951

request mpls lsp adjust-autobandwidth
command.....811

request pim multicast-tunnel rebalance
command.....125

Resource Reservation Protocol *See* RSVP

reverse path forwarding *See* RPF

RIP

general statistics

clearing.....692

displaying.....694

neighbors

displaying.....696

statistics

clearing.....693

displaying.....698

RIPng

general statistics

clearing.....702

displaying.....704

neighbors

displaying.....705

statistics

clearing.....703

displaying.....707

route advertisements, displaying.....497

route, displaying

next-hop.....602

router advertisements

IPv6

clearing.....303

displaying.....306

routes, displaying

active.....486

active path.....492

| | | | |
|--|-----|--|--------------------|
| advertising protocol..... | 497 | RSVP | |
| all..... | 502 | interfaces, displaying..... | 869 |
| AS paths | | neighbors, displaying..... | 875 |
| distribution of..... | 478 | sessions | |
| domain information..... | 482 | clearing..... | 866 |
| regular expressions, matching..... | 504 | displaying..... | 880 |
| summary of..... | 484 | statistics | |
| best..... | 506 | clearing..... | 868 |
| brief information..... | 509 | displaying..... | 890 |
| CCC..... | 511 | version, displaying..... | 894 |
| community ASN..... | 512 | S | |
| community name..... | 514 | SAP session announcements, displaying..... | 298 |
| damping, BGP..... | 516 | Session Announcement Protocol See SAP | |
| detailed information..... | 522 | shortest path first..... | 355 |
| extensive information..... | 543 | show (ospf ospf3) backup coverage | |
| flow validation..... | 558 | command..... | 403 |
| hidden..... | 578 | show (ospf ospf3) backup lsp..... | 406 |
| in a prefix range..... | 627 | show (ospf ospf3) backup neighbor..... | 408 |
| in a specific routing table..... | 662 | show (ospf ospf3) backup spf..... | 410 |
| in the forwarding table..... | 560 | show (ospf ospf3) interface command..... | 440 |
| in the LDP internal topology table..... | 791 | show (ospf ospf3) io-statistics command..... | 446 |
| inactive path..... | 581 | show (ospf ospf3) log command..... | 448 |
| inactive prefix..... | 585 | show (ospf ospf3) neighbor command..... | 452 |
| instances..... | 587 | show (ospf ospf3) overview command..... | 459 |
| learned from a specific address..... | 652 | show (ospf ospf3) route command..... | 464 |
| learned from a specific protocol..... | 615 | show (ospf ospf3) statistics command..... | 470 |
| learned from snooping..... | 644 | show ancp cos command..... | 10 |
| LSP..... | 596 | show ancp neighbor command..... | 13 |
| martian..... | 600 | show ancp subscriber command..... | 19 |
| matching the specified address..... | 537 | show as-path command..... | 478 |
| MPLS labels..... | 594 | show as-path domain command..... | 482 |
| next-hop resolution..... | 641 | show as-path summary command..... | 484 |
| not associated with a community..... | 608 | show bfd session command..... | 28 |
| policy-based route export..... | 539 | show bgp bmp command..... | 43 |
| received through a neighbor..... | 632 | show bgp group command..... | 44 |
| sent to a specific interface..... | 611 | show bgp group traffic-statistics command..... | 51 |
| summary statistics..... | 658 | show bgp neighbor command..... | 53 |
| terse information..... | 673 | show bgp replication command..... | 67 |
| Routing Information Protocol See RIP | | show bgp summary command..... | 69 |
| Routing Information Protocol next generation See | | show bridge domain command..... | 903 |
| RIPng | | show bridge flood command..... | 905 |
| routing policies | | show bridge mac-table command..... | 912 |
| displaying..... | 766 | show bridge statistics command..... | 916 |
| testing the configuration for..... | 770 | show command..... | 380, 382, 385, 386 |
| RPF | | show connections command..... | 813 |
| calculations, displaying..... | 201 | show dhcp relay binding command..... | 745 |
| PIM source state, displaying..... | 283 | show dhcp relay statistics command..... | 750 |
| RPs | | show dhcpv6 relay binding command..... | 753 |
| displaying..... | 258 | | |

| | |
|--|-----|
| show dhcpv6 relay statistics command..... | 759 |
| show dvmrp interfaces command..... | 126 |
| show dvmrp neighbors command..... | 128 |
| show dvmrp prefix command..... | 130 |
| show dvmrp prunes command..... | 132 |
| show dynamic-tunnels database command..... | 952 |
| show esis adjacency command..... | 80 |
| show esis interface command..... | 82 |
| show esis statistics command..... | 84 |
| show firewall command..... | 714 |
| show firewall filter version command..... | 721 |
| show firewall log command..... | 722 |
| show firewall prefix-action-stats command..... | 725 |
| show firewall templates-in-use command..... | 726 |
| show helper statistics command..... | 762 |
| show igmp group command..... | 134 |
| show igmp interface command..... | 138 |
| show igmp snooping interface command..... | 142 |
| show igmp snooping membership command..... | 145 |
| show igmp snooping statistics command..... | 149 |
| show igmp statistics command..... | 152 |
| show ipv6 neighbors command..... | 304 |
| show ipv6 router-advertisement command..... | 306 |
| show isis adjacency command..... | 319 |
| show isis authentication command..... | 323 |
| show isis backup coverage..... | 325 |
| show isis backup spf results..... | 329 |
| show isis context-identifier command..... | 333 |
| show isis database command..... | 335 |
| show isis hostname command..... | 342 |
| show isis interface command..... | 343 |
| show isis overview command..... | 348 |
| show isis route command..... | 351 |
| show isis spf command..... | 355 |
| show isis statistics command..... | 360 |
| show isis-backup label-switched-path command..... | 327 |
| show l2-learning global-information command..... | 920 |
| show l2-learning global-mac-count command..... | 921 |
| show l2-learning instance command..... | 922 |
| show l2-learning interface command..... | 924 |
| show l2circuit connections command..... | 957 |
| show l2vpn connections command..... | 963 |
| show ldp database command..... | 780 |
| show ldp fec-filters command..... | 784 |
| show ldp interface command..... | 785 |
| show ldp neighbor command..... | 787 |
| show ldp path command..... | 789 |
| show ldp route command..... | 791 |
| show ldp session command..... | 795 |
| show ldp statistics command..... | 801 |
| show ldp traffic-statistics command..... | 805 |
| show link-management command..... | 816 |
| show link-management peer command..... | 820 |
| show link-management routing command..... | 822 |
| show link-management statistics command..... | 825 |
| show link-management te-link command..... | 827 |
| show lldp command..... | 366 |
| show lldp local-information command..... | 369 |
| show lldp neighbors command..... | 371 |
| show lldp remote-global-statistics command..... | 375 |
| show lldp statistics command..... | 377 |
| show mac-rewrite interface command..... | 926 |
| show mld group command..... | 155 |
| show mld interface command..... | 160 |
| show mld statistics command..... | 164 |
| show mpls admin-groups command..... | 829 |
| show mpls call-admission-control command..... | 830 |
| show mpls cspf command..... | 832 |
| show mpls diffserv-te command..... | 834 |
| show mpls interface command..... | 836 |
| show mpls lsp command..... | 838 |
| show mpls path command..... | 851 |
| show mpls static-lsp command..... | 853 |
| show msdp command..... | 167 |
| show msdp source command..... | 169 |
| show msdp source-active command..... | 171 |
| show msdp statistics command..... | 174 |
| show multicast backup-pe-groups command..... | 178 |
| show multicast flow-map command..... | 180 |
| show multicast forwarding-cache statistics command..... | 182 |
| show multicast interface command..... | 184 |
| show multicast mrinfo command..... | 186 |
| show multicast next-hops command..... | 188 |
| show multicast pim-to-igmp-proxy command..... | 191 |
| show multicast pim-to-mld-proxy command..... | 193 |
| show multicast route command..... | 195 |
| show multicast rpf command..... | 201 |
| show multicast scope command..... | 205 |
| show multicast sessions command..... | 207 |
| show multicast snooping next-hops command..... | 210 |
| show multicast snooping route command..... | 212 |
| show multicast snooping statistics command..... | 215 |
| show multicast statistics command..... | 218 |

| | |
|--|----------|
| show multicast usage command..... | 221 |
| show mvpn c-multicast command..... | 969 |
| show mvpn instance command..... | 971 |
| show mvpn neighbor command..... | 975 |
| show mvrp statistics command..... | 388 |
| show ospf context-identifier command..... | 418 |
| show ospf database command..... | 420 |
| show ospf3 database command..... | 429 |
| show pgm negative-acknowledgments command..... | 224 |
| show pgm source-path-messages command..... | 226 |
| show pgm statistics command..... | 227 |
| show pim bootstrap command..... | 230 |
| show pim interfaces command..... | 232 |
| show pim join command..... | 235 |
| show pim mdt command..... | 244, 252 |
| show pim mdt data-mdt-joins command..... | 248 |
| show pim mdt data-mdt-limit command..... | 250 |
| show pim neighbors command..... | 253 |
| show pim rps command..... | 258 |
| show pim source command..... | 283 |
| show pim statistics command..... | 285 |
| show policer command..... | 728 |
| show policy command..... | 766 |
| show policy conditions command..... | 768 |
| show policy damping command..... | 74 |
| show rip general-statistics command..... | 694 |
| show rip neighbor command..... | 696 |
| show rip statistics command..... | 698 |
| show ripng general-statistics command..... | 704 |
| show ripng neighbor command..... | 705 |
| show ripng statistics command..... | 707 |
| show route active-path command..... | 492 |
| show route advertising-protocol command..... | 497 |
| show route all command..... | 502 |
| show route aspath-regex command..... | 504 |
| show route best command..... | 506 |
| show route brief command..... | 509 |
| show route ccc command..... | 511 |
| show route command..... | 486 |
| show route community command..... | 512 |
| show route community-name command..... | 514 |
| show route damping command..... | 516 |
| show route detail command..... | 522 |
| show route exact command..... | 537 |
| show route export command..... | 539 |
| show route export vrf-target command..... | 541 |
| show route extensive command..... | 543 |
| show route flow validation command..... | 558 |
| show route forwarding-table command..... | 560 |
| show route forwarding-table interface-name command..... | 574 |
| show route hidden command..... | 578 |
| show route inactive-path command..... | 581 |
| show route inactive-prefix command..... | 585 |
| show route instance command..... | 587 |
| show route label command..... | 594 |
| show route label-switched-path command..... | 596 |
| show route martians command..... | 600 |
| show route next-hop command..... | 602 |
| show route no-community command..... | 608 |
| show route output command..... | 611 |
| show route protocol command..... | 615 |
| show route range command..... | 627 |
| show route receive-protocol command..... | 632 |
| show route resolution command..... | 641 |
| show route snooping command..... | 644 |
| show route source-gateway command..... | 652 |
| show route summary command..... | 658 |
| show route table command..... | 662 |
| show route terse command..... | 673 |
| show rsvp interface command..... | 869 |
| show rsvp neighbor command..... | 875 |
| show rsvp session command..... | 880 |
| show rsvp statistics command..... | 890 |
| show rsvp version command..... | 894 |
| show sap listen command..... | 298 |
| show spanning-tree bridge command..... | 930 |
| show spanning-tree interface command..... | 935 |
| show spanning-tree mstp configuration command..... | 941 |
| show spanning-tree statistics command..... | 943 |
| show ted database command..... | 856 |
| show ted link command..... | 861 |
| show ted protocol command..... | 863 |
| show vpls connections command..... | 979 |
| show vpls flood event-queue command..... | 990 |
| show vpls flood instance command..... | 992 |
| show vpls flood route command..... | 994 |
| show vpls mac-table command..... | 996 |
| show vpls statistics command..... | 1001 |
| snooping (interface) IGMP..... | 142 |
| snooping (membership) IGMP..... | 145 |
| snooping (statistics) IGMP..... | 149 |
| snooping routes, displaying..... | 644 |

- source gateway addresses, displaying.....652
 - Spanning Tree Protocol *See* STP
 - SPF calculations, displaying.....448
 - static LSPs
 - MPLS, displaying.....853
 - statistics
 - Layer 2 bridging, displaying.....916
 - LLDP.....365
 - STP
 - bridge
 - displaying.....930
 - interface
 - displaying.....935
 - protocol
 - clearing.....928
 - statistics
 - clearing.....929
 - displaying.....943
 - subscriber
 - ANCP.....6
 - support, technical *See* technical support
 - syntax conventions.....xvi
- T**
- technical support
 - contacting JTAC.....xvii
 - TED *See* traffic engineering database
 - test msdp command.....299
 - test policy command.....770
 - traffic engineering database
 - database entries, displaying.....856
 - link information, displaying.....861
 - protocols learned from, displaying.....863
- U**
- UDP, statistics
 - clearing.....743
 - displaying.....762
 - User Datagram Protocol *See* UDP
- V**
- version
 - RSVP, displaying.....894
 - virtual private LAN service *See* VPLS
 - VPLS
 - connections, displaying.....979
 - statistics, displaying.....1001
 - VPN routing and forwarding *See* VRF
 - VPNs *See* Layer 2 VPNs, Layer 3 VPNs
- VRF
 - route export, displaying.....541

Index of Statements and Commands

C

| | |
|--|-----|
| clear (ospf ospf3) database command..... | 393 |
| clear (ospf ospf3) database-protection command..... | 396 |
| clear (ospf ospf3) io-statistics command..... | 397 |
| clear (ospf ospf3) neighbor command..... | 398 |
| clear (ospf ospf3) overload command..... | 400 |
| clear (ospf ospf3) statistics command..... | 401 |
| clear ancp neighbor command..... | 4 |
| clear ancp subscriber command..... | 6 |
| clear bfd adaptation command..... | 26 |
| clear bfd session command..... | 27 |
| clear bgp damping command..... | 38 |
| clear bgp neighbor command..... | 39 |
| clear bgp table command..... | 41 |
| clear bridge mac-table command..... | 900 |
| clear dhcp relay binding command..... | 732 |
| clear dhcp relay statistics command..... | 734 |
| clear dhcpv6 relay binding command..... | 737 |
| clear dhcpv6 relay statistics command..... | 741 |
| clear error bpdu command..... | 901 |
| clear error mac-rewrite command..... | 902 |
| clear esis adjacency command..... | 78 |
| clear esis statistics command..... | 79 |
| clear firewall command..... | 712 |
| clear helper statistics command..... | 743 |
| clear igmp membership command..... | 92 |
| clear igmp snooping membership command..... | 95 |
| clear igmp snooping statistics command..... | 96 |
| clear igmp statistics command..... | 97 |
| clear ipv6 neighbors command..... | 302 |
| clear ipv6 router-advertisement command..... | 303 |
| clear isis adjacency command..... | 311 |
| clear isis database command..... | 313 |
| clear isis overload command..... | 315 |
| clear isis statistics command..... | 317 |

| | |
|--|-----|
| clear ldp neighbor command..... | 777 |
| clear ldp session command..... | 778 |
| clear ldp statistics command..... | 779 |
| clear lldp neighbor command..... | 364 |
| clear lldp statistics command..... | 365 |
| clear mld membership command..... | 99 |
| clear mld statistics command..... | 100 |
| clear msdp cache command..... | 101 |
| clear msdp statistics command..... | 102 |
| clear multicast bandwidth-admission command..... | 103 |
| clear multicast forwarding-cache command..... | 105 |
| clear multicast scope command..... | 106 |
| clear multicast sessions command..... | 107 |
| clear multicast snooping statistics command..... | 108 |
| clear multicast statistics command..... | 109 |
| clear pgm negative-acknowledgments command..... | 110 |
| clear pgm source-path-messages command..... | 111 |
| clear pgm statistics command..... | 112 |
| clear pim join command..... | 113 |
| clear pim join-distribution command..... | 114 |
| clear pim register command..... | 116 |
| clear pim statistics command..... | 122 |
| clear rip general-statistics command..... | 692 |
| clear rip statistics command..... | 693 |
| clear ripng general-statistics command..... | 702 |
| clear ripng statistics command..... | 703 |
| clear rsvp session command..... | 866 |
| clear rsvp statistics command..... | 868 |
| clear spanning-tree protocol-migration command..... | 928 |
| clear spanning-tree statistics command..... | 929 |
| clear vpls mac-address command..... | 949 |
| clear vpls mac-table command..... | 950 |

R

| | |
|--|-----|
| request ancp oam interface command..... | 8 |
| request ancp oam neighbor command..... | 9 |
| request l2circuit-switchover command..... | 951 |
| request mpls lsp adjust-autobandwidth command..... | 811 |
| request pim multicast-tunnel rebalance command..... | 125 |

S

| | |
|---|-----|
| show (ospf ospf3) backup coverage command..... | 403 |
| show (ospf ospf3) backup lsp..... | 406 |

| | |
|--|--------------------|
| show (ospf ospf3) backup neighbor..... | 408 |
| show (ospf ospf3) backup spf..... | 410 |
| show (ospf ospf3) interface command..... | 440 |
| show (ospf ospf3) io-statistics command..... | 446 |
| show (ospf ospf3) log command..... | 448 |
| show (ospf ospf3) neighbor command..... | 452 |
| show (ospf ospf3) overview command..... | 459 |
| show (ospf ospf3) route command..... | 464 |
| show (ospf ospf3) statistics command..... | 470 |
| show ancp cos command..... | 10 |
| show ancp neighbor command..... | 13 |
| show ancp subscriber command..... | 19 |
| show as-path command..... | 478 |
| show as-path domain command..... | 482 |
| show as-path summary command..... | 484 |
| show bfd session command..... | 28 |
| show bgp bmp command..... | 43 |
| show bgp group command..... | 44 |
| show bgp group traffic-statistics command..... | 51 |
| show bgp neighbor command..... | 53 |
| show bgp replication command..... | 67 |
| show bgp summary command..... | 69 |
| show bridge domain command..... | 903 |
| show bridge flood command..... | 905 |
| show bridge mac-table command..... | 912 |
| show bridge statistics command..... | 916 |
| show command..... | 380, 382, 385, 386 |
| show connections command..... | 813 |
| show dhcp relay binding command..... | 745 |
| show dhcp relay statistics command..... | 750 |
| show dhcpv6 relay binding command..... | 753 |
| show dhcpv6 relay statistics command..... | 759 |
| show dvmp interfaces command..... | 126 |
| show dvmp neighbors command..... | 128 |
| show dvmp prefix command..... | 130 |
| show dvmp prunes command..... | 132 |
| show dynamic-tunnels database command..... | 952 |
| show esis adjacency command..... | 80 |
| show esis interface command..... | 82 |
| show esis statistics command..... | 84 |
| show firewall command..... | 714 |
| show firewall filter version command..... | 721 |
| show firewall log command..... | 722 |
| show firewall prefix-action-stats command..... | 725 |
| show firewall templates-in-use command..... | 726 |
| show helper statistics command..... | 762 |
| show igmp group command..... | 134 |
| show igmp interface command..... | 138 |
| show igmp snooping interface command..... | 142 |
| show igmp snooping membership command..... | 145 |
| show igmp snooping statistics command..... | 149 |
| show igmp statistics command..... | 152 |
| show ipv6 neighbors command..... | 304 |
| show ipv6 router-advertisement command..... | 306 |
| show isis adjacency command..... | 319 |
| show isis authentication command..... | 323 |
| show isis backup coverage..... | 325 |
| show isis backup spf results..... | 329 |
| show isis database command..... | 335 |
| show isis hostname command..... | 342 |
| show isis interface command..... | 343 |
| show isis overview command..... | 348 |
| show isis route command..... | 351 |
| show isis spf command..... | 355 |
| show isis statistics command..... | 360 |
| show isis-backup label-switched-path command..... | 327 |
| show l2-learning global-information command..... | 920 |
| show l2-learning global-mac-count command..... | 921 |
| show l2-learning instance command..... | 922 |
| show l2-learning interface command..... | 924 |
| show l2circuit connections command..... | 957 |
| show l2vpn connections command..... | 963 |
| show ldp database command..... | 780 |
| show ldp fec-filters command..... | 784 |
| show ldp interface command..... | 785 |
| show ldp neighbor command..... | 787 |
| show ldp path command..... | 789 |
| show ldp route command..... | 791 |
| show ldp session command..... | 795 |
| show ldp statistics command..... | 801 |
| show ldp traffic-statistics command..... | 805 |
| show link-management command..... | 816 |
| show link-management peer command..... | 820 |
| show link-management routing command..... | 822 |
| show link-management statistics command..... | 825 |
| show link-management te-link command..... | 827 |
| show lldp command..... | 366 |
| show lldp local-information command..... | 369 |
| show lldp neighbors command..... | 371 |
| show lldp remote-global-statistics command..... | 375 |
| show lldp statistics command..... | 377 |
| show mac-rewrite interface command..... | 926 |
| show mld group command..... | 155 |
| show mld interface command..... | 160 |
| show mld statistics command..... | 164 |

| | | | |
|--|----------|--|-----|
| show mpls admin-groups command..... | 829 | show pim rps command..... | 258 |
| show mpls call-admission-control command..... | 830 | show pim source command..... | 283 |
| show mpls cspf command..... | 832 | show pim statistics command..... | 285 |
| show mpls diffserv-te command..... | 834 | show policer command..... | 728 |
| show mpls interface command..... | 836 | show policy command..... | 766 |
| show mpls lsp command..... | 838 | show policy conditions command..... | 768 |
| show mpls path command..... | 851 | show policy damping command..... | 74 |
| show mpls static-lsp command..... | 853 | show rip general-statistics command..... | 694 |
| show msdp command..... | 167 | show rip neighbor command..... | 696 |
| show msdp source command..... | 169 | show rip statistics command..... | 698 |
| show msdp source-active command..... | 171 | show ripng general-statistics command..... | 704 |
| show msdp statistics command..... | 174 | show ripng neighbor command..... | 705 |
| show multicast backup-pe-groups command..... | 178 | show ripng statistics command..... | 707 |
| show multicast flow-map command..... | 180 | show route active-path command..... | 492 |
| show multicast forwarding-cache statistics command..... | 182 | show route advertising-protocol command..... | 497 |
| show multicast interface command..... | 184 | show route all command..... | 502 |
| show multicast minfo command..... | 186 | show route aspath-regex command..... | 504 |
| show multicast next-hops command..... | 188 | show route best command..... | 506 |
| show multicast pim-to-igmp-proxy command..... | 191 | show route brief command..... | 509 |
| show multicast pim-to-mld-proxy command..... | 193 | show route ccc command..... | 511 |
| show multicast route command..... | 195 | show route command..... | 486 |
| show multicast rpf command..... | 201 | show route community command..... | 512 |
| show multicast scope command..... | 205 | show route community-name command..... | 514 |
| show multicast sessions command..... | 207 | show route damping command..... | 516 |
| show multicast snooping next-hops command..... | 210 | show route detail command..... | 522 |
| show multicast snooping route command..... | 212 | show route exact command..... | 537 |
| show multicast snooping statistics command..... | 215 | show route export command..... | 539 |
| show multicast statistics command..... | 218 | show route export vrf-target command..... | 541 |
| show multicast usage command..... | 221 | show route extensive command..... | 543 |
| show mvpn c-multicast command..... | 969 | show route flow validation command..... | 558 |
| show mvpn instance command..... | 971 | show route forwarding-table command..... | 560 |
| show mvpn neighbor command..... | 975 | show route forwarding-table interface-name command..... | 574 |
| show mvrp statistics command..... | 388 | show route hidden command..... | 578 |
| show ospf context-identifier command..... | 418 | show route inactive-path command..... | 581 |
| show ospf database command..... | 420 | show route inactive-prefix command..... | 585 |
| show ospf3 database command..... | 429 | show route instance command..... | 587 |
| show pgm negative-acknowledgments command..... | 224 | show route label command..... | 594 |
| show pgm source-path-messages command..... | 226 | show route label-switched-path command..... | 596 |
| show pgm statistics command..... | 227 | show route martians command..... | 600 |
| show pim bootstrap command..... | 230 | show route next-hop command..... | 602 |
| show pim interfaces command..... | 232 | show route no-community command..... | 608 |
| show pim join command..... | 235 | show route output command..... | 611 |
| show pim mdt command..... | 244, 252 | show route protocol command..... | 615 |
| show pim mdt data-mdt-joins command..... | 248 | show route range command..... | 627 |
| show pim mdt data-mdt-limit command..... | 250 | show route receive-protocol command..... | 632 |
| show pim neighbors command..... | 253 | show route resolution command..... | 641 |
| | | show route snooping command..... | 644 |
| | | show route source-gateway command..... | 652 |

| | |
|---|------|
| show route summary command..... | 658 |
| show route table command..... | 662 |
| show route terse command..... | 673 |
| show rsvp interface command..... | 869 |
| show rsvp neighbor command..... | 875 |
| show rsvp session command..... | 880 |
| show rsvp statistics command..... | 890 |
| show rsvp version command..... | 894 |
| show sap listen command..... | 298 |
| show spanning-tree bridge command..... | 930 |
| show spanning-tree interface command..... | 935 |
| show spanning-tree mstp configuration command..... | 941 |
| show spanning-tree statistics command..... | 943 |
| show ted database command..... | 856 |
| show ted link command..... | 861 |
| show ted protocol command..... | 863 |
| show vpls connections command..... | 979 |
| show vpls flood event-queue command..... | 990 |
| show vpls flood instance command..... | 992 |
| show vpls flood route command..... | 994 |
| show vpls mac-table command..... | 996 |
| show vpls statistics command..... | 1001 |

T

| | |
|--------------------------|-----|
| test msdp command..... | 299 |
| test policy command..... | 770 |