



JUNOS® Software

Routing Protocols and Policies Command Reference

Release 9.6

Juniper Networks, Inc.

1194 North Mathilda Avenue
Sunnyvale, California 94089
USA

408-745-2000

www.juniper.net

Published: 2009-07-16

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, the Juniper Networks logo, JUNOS, NetScreen, ScreenOS, and Steel-Belted Radius are registered trademarks of Juniper Networks, Inc. in the United States and other countries. JUNOSe is a trademark 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® Software Routing Protocols and Policies Command Reference

Release 9.6

Copyright © 2009, Juniper Networks, Inc.

All rights reserved. Printed in USA.

Writing: Philomena Dolla, Mark Barnard, Andrea Couvrey, Walter Goralski, Fran Singer, Albert Statti, Ines Salazar, Roy Spencer, and Rekha J

Editing: Stella Hackell, Nancy Kurahashi, Sonia Saruba, Joanne McClintock

Cover Design: Edmonds Design

Revision History

July 2009—R1 JUNOS 9.6

The information in this document is current as of the date listed in the revision history.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. The JUNOS Software 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

READ THIS END USER LICENSE AGREEMENT ("AGREEMENT") BEFORE DOWNLOADING, INSTALLING, OR USING THE SOFTWARE. BY DOWNLOADING, INSTALLING, OR USING THE SOFTWARE OR OTHERWISE EXPRESSING YOUR AGREEMENT TO THE TERMS CONTAINED HEREIN, YOU (AS CUSTOMER OR IF YOU ARE NOT THE CUSTOMER, AS A REPRESENTATIVE/AGENT AUTHORIZED TO BIND THE CUSTOMER) CONSENT TO BE BOUND BY THIS AGREEMENT. IF YOU DO NOT OR CANNOT AGREE TO THE TERMS CONTAINED HEREIN, THEN (A) DO NOT DOWNLOAD, INSTALL, OR USE THE SOFTWARE, AND (B) YOU MAY CONTACT JUNIPER NETWORKS REGARDING LICENSE TERMS.

1. **The Parties.** The parties to this Agreement are (i) Juniper Networks, Inc. (if the Customer's principal office is located in the Americas) or Juniper Networks (Cayman) Limited (if the Customer's principal office is located outside the Americas) (such applicable entity being referred to herein as "Juniper"), and (ii) the person or organization that originally purchased from Juniper or an authorized Juniper reseller the applicable license(s) for use of the Software ("Customer") (collectively, the "Parties").

2. **The Software.** In this Agreement, "Software" means the program modules and features of the Juniper or Juniper-supplied software, for which Customer has paid the applicable license or support fees to Juniper or an authorized Juniper reseller, or which was embedded by Juniper in equipment which Customer purchased from Juniper or an authorized Juniper reseller. "Software" also includes updates, upgrades and new releases of such software. "Embedded Software" means Software which Juniper has embedded in or loaded onto the Juniper equipment and any updates, upgrades, additions or replacements which are subsequently embedded in or loaded onto the equipment.

3. **License Grant.** Subject to payment of the applicable fees and the limitations and restrictions set forth herein, Juniper grants to Customer a non-exclusive and non-transferable license, without right to sublicense, to use the Software, in executable form only, subject to the following use restrictions:

- a. Customer shall use Embedded Software solely as embedded in, and for execution on, Juniper equipment originally purchased by Customer from Juniper or an authorized Juniper reseller.
- b. Customer shall use the Software on a single hardware chassis having a single processing unit, or as many chassis or processing units for which Customer has paid the applicable license fees; provided, however, with respect to the Steel-Belted Radius or Odyssey Access Client software only, Customer shall use such Software on a single computer containing a single physical random access memory space and containing any number of processors. Use of the Steel-Belted Radius or IMS AAA software on multiple computers or virtual machines (e.g., Solaris zones) requires multiple licenses, regardless of whether such computers or virtualizations are physically contained on a single chassis.
- c. Product purchase documents, paper or electronic user documentation, and/or the particular licenses purchased by Customer may specify limits to Customer's use of the Software. Such limits may restrict use to a maximum number of seats, registered endpoints, concurrent users, sessions, calls, connections, subscribers, clusters, nodes, realms, devices, links, ports or transactions, or require the purchase of separate licenses to use particular features, functionalities, services, applications, operations, or capabilities, or provide throughput, performance, configuration, bandwidth, interface, processing, temporal, or geographical limits. In addition, such limits may restrict the use of the Software to managing certain kinds of networks or require the Software to be used only in conjunction with other specific Software. Customer's use of the Software shall be subject to all such limitations and purchase of all applicable licenses.
- d. For any trial copy of the Software, Customer's right to use the Software expires 30 days after download, installation or use of the Software. Customer may operate the Software after the 30-day trial period only if Customer pays for a license to do so. Customer may not extend or create an additional trial period by re-installing the Software after the 30-day trial period.
- e. The Global Enterprise Edition of the Steel-Belted Radius software may be used by Customer only to manage access to Customer's enterprise network. Specifically, service provider customers are expressly prohibited from using the Global Enterprise Edition of the Steel-Belted Radius software to support any commercial network access services.

The foregoing license is not transferable or assignable by Customer. No license is granted herein to any user who did not originally purchase the applicable license(s) for the Software from Juniper or an authorized Juniper reseller.

4. **Use Prohibitions.** Notwithstanding the foregoing, the license provided herein does not permit the Customer to, and Customer agrees not to and shall not: (a) modify, unbundle, reverse engineer, or create derivative works based on the Software; (b) make unauthorized copies of the Software (except as necessary for backup purposes); (c) rent, sell, transfer, or grant any rights in and to any copy of the Software, in any form, to any third party; (d) remove any proprietary notices, labels, or marks on or in any copy of the Software or any product in which the Software is embedded; (e) distribute any copy of the Software to any third party, including as may be embedded in Juniper equipment sold in the secondhand market; (f) use any 'locked' or key-restricted feature, function, service, application, operation, or capability without first purchasing the applicable license(s) and obtaining a valid key from Juniper, even if such feature, function, service, application, operation, or capability is enabled without a key; (g) distribute any key for the Software provided by Juniper to any third party; (h) use the Software in any manner that extends or is broader than the uses purchased by Customer from Juniper or an authorized Juniper reseller; (i) use Embedded Software on non-Juniper equipment; (j) use Embedded Software (or make it available for use) on Juniper equipment that the Customer did not originally purchase from Juniper or an authorized Juniper reseller; (k) disclose the results of testing or benchmarking of the Software to any third party without the prior written consent of Juniper; or (l) use the Software in any manner other than as expressly provided herein.

5. **Audit.** Customer shall maintain accurate records as necessary to verify compliance with this Agreement. Upon request by Juniper, Customer shall furnish such records to Juniper and certify its compliance with this Agreement.

6. **Confidentiality.** The Parties agree that aspects of the Software and associated documentation are the confidential property of Juniper. As such, Customer shall exercise all reasonable commercial efforts to maintain the Software and associated documentation in confidence, which at a minimum includes restricting access to the Software to Customer employees and contractors having a need to use the Software for Customer's internal business purposes.

7. **Ownership.** Juniper and Juniper's licensors, respectively, retain ownership of all right, title, and interest (including copyright) in and to the Software, associated documentation, and all copies of the Software. Nothing in this Agreement constitutes a transfer or conveyance of any right, title, or interest in the Software or associated documentation, or a sale of the Software, associated documentation, or copies of the Software.

8. **Warranty, Limitation of Liability, Disclaimer of Warranty.** The warranty applicable to the Software shall be as set forth in the warranty statement that accompanies the Software (the "Warranty Statement"). Nothing in this Agreement shall give rise to any obligation to support the Software. Support services may be purchased separately. Any such support shall be governed by a separate, written support services agreement. TO THE MAXIMUM EXTENT PERMITTED BY LAW, JUNIPER SHALL NOT BE LIABLE FOR ANY LOST PROFITS, LOSS OF DATA, OR COSTS OR PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES, OR FOR ANY SPECIAL, INDIRECT, OR CONSEQUENTIAL DAMAGES ARISING OUT OF THIS AGREEMENT, THE SOFTWARE, OR ANY JUNIPER OR JUNIPER-SUPPLIED SOFTWARE. IN NO EVENT SHALL JUNIPER BE LIABLE FOR DAMAGES ARISING FROM UNAUTHORIZED OR IMPROPER USE OF ANY JUNIPER OR JUNIPER-SUPPLIED SOFTWARE, EXCEPT AS EXPRESSLY PROVIDED IN THE WARRANTY STATEMENT TO THE EXTENT PERMITTED BY LAW, JUNIPER DISCLAIMS ANY AND ALL WARRANTIES IN AND TO THE SOFTWARE (WHETHER EXPRESS, IMPLIED, STATUTORY, OR OTHERWISE), INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NONINFRINGEMENT. IN NO EVENT DOES JUNIPER WARRANT THAT THE SOFTWARE, OR ANY EQUIPMENT OR NETWORK RUNNING THE SOFTWARE, WILL OPERATE WITHOUT ERROR OR INTERRUPTION, OR WILL BE FREE OF VULNERABILITY TO INTRUSION OR ATTACK. In no event shall Juniper's or its suppliers' or licensors' liability to Customer, whether in contract, tort (including negligence), breach of warranty, or otherwise, exceed the price paid by Customer for the Software that gave rise to the claim, or if the Software is embedded in another Juniper product, the price paid by Customer for such other product. Customer acknowledges and agrees that Juniper has set its prices and entered into this Agreement in reliance upon the disclaimers of warranty and the limitations of liability set forth herein, that the same reflect an allocation of risk between the Parties (including the risk that a contract remedy may fail of its essential purpose and cause consequential loss), and that the same form an essential basis of the bargain between the Parties.

9. **Termination.** Any breach of this Agreement or failure by Customer to pay any applicable fees due shall result in automatic termination of the license granted herein. Upon such termination, Customer shall destroy or return to Juniper all copies of the Software and related documentation in Customer's possession or control.

10. **Taxes.** All license fees payable under this agreement are exclusive of tax. Customer shall be responsible for paying Taxes arising from the purchase of the license, or importation or use of the Software. If applicable, valid exemption documentation for each taxing jurisdiction shall be provided to Juniper prior to invoicing, and Customer shall promptly notify Juniper if their exemption is revoked or modified. All payments made by Customer shall be net of any applicable withholding tax. Customer will provide reasonable assistance to Juniper in connection with such withholding taxes by promptly: providing Juniper with valid tax receipts and other required documentation showing Customer's payment of any withholding taxes; completing appropriate applications that would reduce the amount of withholding tax to be paid; and notifying and assisting Juniper in any audit or tax proceeding related to transactions hereunder. Customer shall comply with all applicable tax laws and regulations, and Customer will promptly pay or reimburse Juniper for all costs and damages related to any liability incurred by Juniper as a result of Customer's non-compliance or delay with its responsibilities herein. Customer's obligations under this Section shall survive termination or expiration of this Agreement.

11. **Export.** Customer agrees to comply with all applicable export laws and restrictions and regulations of any United States and any applicable foreign agency or authority, and not to export or re-export the Software or any direct product thereof in violation of any such restrictions, laws or regulations, or without all necessary approvals. Customer shall be liable for any such violations. The version of the Software supplied to Customer may contain encryption or other capabilities restricting Customer's ability to export the Software without an export license.

12. **Commercial Computer Software.** The Software is "commercial computer software" and is provided with restricted rights. Use, duplication, or disclosure by the United States government is subject to restrictions set forth in this Agreement and as provided in DFARS 227.7201 through 227.7202-4, FAR 12.212, FAR 27.405(b)(2), FAR 52.227-19, or FAR 52.227-14(ALT III) as applicable.

13. **Interface Information.** To the extent required by applicable law, and at Customer's written request, Juniper shall provide Customer with the interface information needed to achieve interoperability between the Software and another independently created program, on payment of applicable fee, if any. Customer shall observe strict obligations of confidentiality with respect to such information and shall use such information in compliance with any applicable terms and conditions upon which Juniper makes such information available.

14. **Third Party Software.** Any licensor of Juniper whose software is embedded in the Software and any supplier of Juniper whose products or technology are embedded in (or services are accessed by) the Software shall be a third party beneficiary with respect to this Agreement, and such licensor or vendor shall have the right to enforce this Agreement in its own name as if it were Juniper. In addition, certain third party software may be provided with the Software and is subject to the accompanying license(s), if any, of its respective owner(s). To the extent portions of the Software are distributed under and subject to open source licenses obligating Juniper to make the source code for such portions publicly available (such as the GNU General Public License ("GPL") or the GNU Library General Public License ("LGPL")), Juniper will make such source code portions (including Juniper modifications, as appropriate) available upon request for a period of up to three years from the date of distribution. Such request can be made in writing to Juniper Networks, Inc., 1194 N. Mathilda Ave., Sunnyvale, CA 94089, ATTN: General Counsel. You may obtain a copy of the GPL at <http://www.gnu.org/licenses/gpl.html>, and a copy of the LGPL at <http://www.gnu.org/licenses/lgpl.html>.

15. **Miscellaneous.** This Agreement shall be governed by the laws of the State of California without reference to its conflicts of laws principles. The provisions of the U.N. Convention for the International Sale of Goods shall not apply to this Agreement. For any disputes arising under this Agreement, the Parties hereby consent to the personal and exclusive jurisdiction of, and venue in, the state and federal courts within Santa Clara County, California. This Agreement constitutes the entire and sole agreement between Juniper and the Customer with respect to the Software, and supersedes all prior and contemporaneous

agreements relating to the Software, whether oral or written (including any inconsistent terms contained in a purchase order), except that the terms of a separate written agreement executed by an authorized Juniper representative and Customer shall govern to the extent such terms are inconsistent or conflict with terms contained herein. No modification to this Agreement nor any waiver of any rights hereunder shall be effective unless expressly assented to in writing by the party to be charged. If any portion of this Agreement is held invalid, the Parties agree that such invalidity shall not affect the validity of the remainder of this Agreement. This Agreement and associated documentation has been written in the English language, and the Parties agree that the English version will govern. (For Canada: Les parties aux présentes confirment leur volonté que cette convention de même que tous les documents y compris tout avis qui s'y rattache, soient rédigés en langue anglaise. (Translation: The parties confirm that this Agreement and all related documentation is and will be in the English language)).

Abbreviated Table of Contents

About This Guide

xvii

Part 1

Protocols

| | | |
|------------|--------------------------------------------------------|-----|
| Chapter 1 | ANCP Operational Mode Commands | 3 |
| Chapter 2 | BFD Operational Mode Commands | 15 |
| Chapter 3 | BGP Operational Mode Commands | 25 |
| Chapter 4 | ES-IS Operational Mode Commands | 57 |
| Chapter 5 | IP Multicast Operational Mode Commands | 65 |
| Chapter 6 | IPv6 Operational Mode Commands | 203 |
| Chapter 7 | IS-IS Operational Mode Commands | 211 |
| Chapter 8 | LLDP Operational Mode Commands | 255 |
| Chapter 9 | OSPF Operational Mode Commands | 269 |
| Chapter 10 | Protocol-Independent Routing Operational Mode Commands | 321 |
| Chapter 11 | RIP Operational Mode Commands | 487 |
| Chapter 12 | RIPng Operational Mode Commands | 495 |

Part 2

Policy Framework

| | | |
|------------|-------------------------------------------|-----|
| Chapter 13 | Firewall Filter Operational Mode Commands | 505 |
| Chapter 14 | Forwarding Operational Mode Commands | 515 |
| Chapter 15 | Routing Policy Operational Mode Commands | 531 |

Part 3

MPLS

| | | |
|------------|--------------------------------|-----|
| Chapter 16 | LDP Operational Mode Commands | 539 |
| Chapter 17 | MPLS Operational Mode Commands | 565 |
| Chapter 18 | RSVP Operational Mode Commands | 609 |

Part 4

Layer 2 Bridging and Switching Operational Mode Commands

| | | |
|------------|----------------------------------------------------------|-----|
| Chapter 19 | Layer 2 Bridging and Switching Operational Mode Commands | 637 |
| Chapter 20 | Spanning Tree Operational Mode Commands | 663 |

Part 5

VPNs

| | | |
|------------|-------------------------------|-----|
| Chapter 21 | VPN Operational Mode Commands | 683 |
|------------|-------------------------------|-----|

Part 6

Index

| | |
|----------------------------------|-----|
| Index | 725 |
| Index of Statements and Commands | 735 |

Table of Contents

| | | |
|------------------|---------------------------------------------|-------------|
| | About This Guide | xvii |
| | JUNOS Documentation and Release Notes | xvii |
| | Objectives | xvii |
| | Audience | xviii |
| | Supported Platforms | xix |
| | Using the Indexes | xix |
| | Documentation Conventions | xix |
| | Documentation Feedback | xxi |
| | Requesting Technical Support | xxi |
| Part 1 | Protocols | |
| Chapter 1 | ANCP Operational Mode Commands | 3 |
| | clear ancp neighbor | 4 |
| | show ancp cos | 5 |
| | show ancp neighbor | 8 |
| | show ancp subscriber | 10 |
| Chapter 2 | BFD Operational Mode Commands | 15 |
| | clear bfd adaptation | 16 |
| | clear bfd session | 17 |
| | show bfd session | 18 |
| Chapter 3 | BGP Operational Mode Commands | 25 |
| | clear bgp damping | 26 |
| | clear bgp neighbor | 27 |
| | clear bgp table | 29 |
| | show bgp bmp | 30 |
| | show bgp group | 31 |
| | show bgp group traffic-statistics | 37 |
| | show bgp neighbor | 39 |
| | show bgp summary | 50 |
| | show policy damping | 54 |

| | | |
|------------------|-----------------------------------------------|-----------|
| Chapter 4 | ES-IS Operational Mode Commands | 57 |
| | clear esis adjacency | 58 |
| | clear esis statistics | 59 |
| | show esis adjacency | 60 |
| | show esis interface | 62 |
| | show esis statistics | 64 |
| Chapter 5 | IP Multicast Operational Mode Commands | 65 |
| | clear igmp membership | 69 |
| | clear igmp snooping membership | 72 |
| | clear igmp snooping statistics | 73 |
| | clear igmp statistics | 74 |
| | clear mld membership | 76 |
| | clear mld statistics | 77 |
| | clear msdp cache | 78 |
| | clear msdp statistics | 79 |
| | clear multicast scope | 80 |
| | clear multicast snooping statistics | 81 |
| | clear multicast bandwidth-admission | 82 |
| | clear multicast sessions | 84 |
| | clear multicast statistics | 85 |
| | clear pgm negative-acknowledgments | 86 |
| | clear pgm source-path-messages | 87 |
| | clear pgm statistics | 88 |
| | clear pim join | 89 |
| | clear pim register | 90 |
| | clear pim statistics | 91 |
| | show dvmrp interfaces | 93 |
| | show dvmrp neighbors | 95 |
| | show dvmrp prefix | 97 |
| | show dvmrp prunes | 99 |
| | show igmp group | 101 |
| | show igmp interface | 104 |
| | show igmp snooping interface | 107 |
| | show igmp snooping membership | 109 |
| | show igmp snooping statistics | 112 |
| | show igmp statistics | 115 |
| | show mld group | 118 |
| | show mld interface | 122 |
| | show mld statistics | 125 |
| | show msdp | 128 |
| | show msdp source | 130 |
| | show msdp source-active | 132 |
| | show msdp statistics | 134 |
| | show multicast flow-map | 136 |
| | show multicast interface | 138 |
| | show multicast minfo | 140 |

| | |
|------------------------------------------|-----|
| show multicast next-hops | 142 |
| show multicast pim-to-igmp-proxy | 144 |
| show multicast pim-to-mld-proxy | 145 |
| show multicast route | 146 |
| show multicast rpf | 150 |
| show multicast scope | 154 |
| show multicast sessions | 156 |
| show multicast snooping route | 158 |
| show multicast snooping statistics | 160 |
| show multicast statistics | 163 |
| show multicast usage | 166 |
| show pgm negative-acknowledgments | 168 |
| show pgm source-path-messages | 170 |
| show pgm statistics | 171 |
| show pim bootstrap | 174 |
| show pim interfaces | 176 |
| show pim join | 179 |
| show pim mdt | 184 |
| show pim mvpn | 187 |
| show pim neighbors | 188 |
| show pim rps | 192 |
| show pim source | 197 |
| show pim statistics | 199 |
| show sap listen | 201 |
| test msdp | 202 |

Chapter 6**IPv6 Operational Mode Commands****203**

| | |
|---------------------------------------|-----|
| clear ipv6 neighbors | 204 |
| clear ipv6 router-advertisement | 205 |
| show ipv6 neighbors | 206 |
| show ipv6 router-advertisement | 207 |

Chapter 7**IS-IS Operational Mode Commands****211**

| | |
|--------------------------------------------|-----|
| clear isis adjacency | 213 |
| clear isis database | 214 |
| clear isis overload | 216 |
| clear isis statistics | 217 |
| show isis adjacency | 219 |
| show isis authentication | 223 |
| show isis backup coverage | 225 |
| show isis backup label-switched-path | 227 |
| show isis backup spf results | 228 |
| show isis database | 231 |
| show isis hostname | 237 |
| show isis interface | 238 |
| show isis overview | 241 |
| show isis route | 244 |

| | |
|----------------------------|-----|
| show isis spf | 247 |
| show isis statistics | 252 |

Chapter 8 **LLDP Operational Mode Commands** **255**

| | |
|------------------------------------------|-----|
| clear lldp neighbor | 256 |
| clear lldp statistics | 257 |
| show lldp | 258 |
| show lldp local-information | 260 |
| show lldp neighbors | 262 |
| show lldp remote-global-statistics | 265 |
| show lldp statistics | 267 |

Chapter 9 **OSPF Operational Mode Commands** **269**

| | |
|------------------------------------------|-----|
| clear (ospf ospf3) database | 271 |
| clear (ospf ospf3) io-statistics | 274 |
| clear (ospf ospf3) neighbor | 275 |
| clear ospf overload | 276 |
| clear (ospf ospf3) statistics | 277 |
| show ospf database | 279 |
| show ospf3 database | 286 |
| show (ospf ospf3) interface | 296 |
| show (ospf ospf3) io-statistics | 301 |
| show (ospf ospf3) log | 302 |
| show (ospf ospf3) neighbor | 304 |
| show (ospf ospf3) overview | 309 |
| show (ospf ospf3) route | 313 |
| show (ospf ospf3) statistics | 318 |

Chapter 10 **Protocol-Independent Routing Operational Mode Commands** **321**

| | |
|---------------------------------------|-----|
| show as-path | 324 |
| show as-path domain | 328 |
| show as-path summary | 330 |
| show route | 331 |
| show route active-path | 335 |
| show route advertising-protocol | 339 |
| show route all | 343 |
| show route aspath-regex | 344 |
| show route best | 346 |
| show route brief | 349 |
| show route ccc | 351 |
| show route community | 352 |
| show route community-name | 353 |
| show route damping | 355 |
| show route detail | 360 |
| show route exact | 374 |
| show route export | 376 |

| | | |
|-------------------|--------------------------------------------------|------------|
| | show route export vrf-target | 378 |
| | show route extensive | 380 |
| | show route flow validation | 392 |
| | show route forwarding-table | 394 |
| | show route hidden | 405 |
| | show route inactive-path | 408 |
| | show route inactive-prefix | 411 |
| | show route instance | 413 |
| | show route label | 420 |
| | show route label-switched-path | 422 |
| | show route martians | 423 |
| | show route next-hop | 425 |
| | show route no-community | 431 |
| | show route output | 434 |
| | show route protocol | 438 |
| | show route range | 446 |
| | show route receive-protocol | 450 |
| | show route resolution | 457 |
| | show route snooping | 460 |
| | show route source-gateway | 468 |
| | show route summary | 474 |
| | show route table | 476 |
| | show route terse | 483 |
| Chapter 11 | RIP Operational Mode Commands | 487 |
| | clear rip general-statistics | 488 |
| | clear rip statistics | 489 |
| | show rip general-statistics | 490 |
| | show rip neighbor | 491 |
| | show rip statistics | 493 |
| Chapter 12 | RIPng Operational Mode Commands | 495 |
| | clear ripng general-statistics | 496 |
| | clear ripng statistics | 497 |
| | show ripng general-statistics | 498 |
| | show ripng neighbor | 499 |
| | show ripng statistics | 500 |
| Part 2 | Policy Framework | |
| Chapter 13 | Firewall Filter Operational Mode Commands | 505 |
| | clear firewall | 506 |
| | show firewall | 507 |
| | show firewall log | 510 |

| | |
|-----------------------------------------|-----|
| show firewall prefix-action-stats | 512 |
| show policer | 513 |

| | | |
|-------------------|---------------------------------------------|------------|
| Chapter 14 | Forwarding Operational Mode Commands | 515 |
|-------------------|---------------------------------------------|------------|

| | |
|-----------------------------------|-----|
| clear dhcp relay binding | 516 |
| clear dhcp relay statistics | 518 |
| clear helper statistics | 520 |
| show dhcp relay binding | 521 |
| show dhcp relay statistics | 525 |
| show helper statistics | 528 |

| | | |
|-------------------|-------------------------------------------------|------------|
| Chapter 15 | Routing Policy Operational Mode Commands | 531 |
|-------------------|-------------------------------------------------|------------|

| | |
|------------------------------|-----|
| show policy | 532 |
| show policy conditions | 534 |
| test policy | 536 |

| | |
|---------------|-------------|
| Part 3 | MPLS |
|---------------|-------------|

| | | |
|-------------------|--------------------------------------|------------|
| Chapter 16 | LDP Operational Mode Commands | 539 |
|-------------------|--------------------------------------|------------|

| | |
|-----------------------------------|-----|
| clear ldp neighbor | 540 |
| clear ldp session | 541 |
| clear ldp statistics | 542 |
| show ldp database | 543 |
| show ldp fec-filters | 546 |
| show ldp interface | 547 |
| show ldp neighbor | 549 |
| show ldp path | 551 |
| show ldp route | 553 |
| show ldp session | 555 |
| show ldp statistics | 559 |
| show ldp traffic-statistics | 562 |

| | | |
|-------------------|---------------------------------------|------------|
| Chapter 17 | MPLS Operational Mode Commands | 565 |
|-------------------|---------------------------------------|------------|

| | |
|---------------------------------------------|-----|
| clear mpls lsp | 567 |
| request mpls lsp adjust-autobandwidth | 569 |
| show connections | 570 |
| show link-management | 573 |
| show link-management peer | 576 |
| show link-management routing | 578 |
| show link-management statistics | 581 |
| show link-management te-link | 583 |
| show mpls admin-groups | 585 |

| | |
|----------------------------------------|-----|
| show mpls call-admission-control | 586 |
| show mpls cspf | 588 |
| show mpls diffserv-te | 590 |
| show mpls interface | 591 |
| show mpls lsp | 592 |
| show mpls path | 600 |
| show ted database | 601 |
| show ted link | 605 |
| show ted protocol | 607 |

Chapter 18 RSVP Operational Mode Commands 609

| | |
|-----------------------------|-----|
| clear rsvp session | 610 |
| clear rsvp statistics | 612 |
| show rsvp interface | 613 |
| show rsvp neighbor | 618 |
| show rsvp session | 622 |
| show rsvp statistics | 629 |
| show rsvp version | 632 |

Part 4 Layer 2 Bridging and Switching Operational Mode Commands

Chapter 19 Layer 2 Bridging and Switching Operational Mode Commands 637

| | |
|-------------------------------------------|-----|
| clear bridge mac-table | 638 |
| clear error bpdu | 639 |
| clear error mac-rewrite | 640 |
| show bridge domain | 641 |
| show bridge flood | 643 |
| show bridge mac-table | 650 |
| show bridge statistics | 653 |
| show l2-learning global-information | 656 |
| show l2-learning global-mac-count | 657 |
| show l2-learning instance | 658 |
| show l2-learning interface | 659 |
| show mac-rewrite interface | 661 |

Chapter 20 Spanning Tree Operational Mode Commands 663

| | |
|----------------------------------------------|-----|
| clear spanning-tree protocol-migration | 664 |
| clear spanning-tree statistics | 665 |
| show spanning-tree bridge | 666 |
| show spanning-tree interface | 671 |
| show spanning-tree mstp configuration | 676 |
| show spanning-tree statistics | 678 |

Part 5**VPNs**

Chapter 21**VPN Operational Mode Commands****683**

| | |
|-------------------------------------|-----|
| clear vpls mac-address | 685 |
| clear vpls mac-table | 686 |
| request l2circuit-switchover | 687 |
| show dynamic-tunnels database | 688 |
| show l2circuit connections | 690 |
| show l2vpn connections | 694 |
| show mvpn c-multicast | 698 |
| show mvpn instance | 700 |
| show mvpn neighbor | 703 |
| show vpls connections | 706 |
| show vpls flood event-queue | 711 |
| show vpls flood instance | 712 |
| show vpls flood route | 714 |
| show vpls mac-table | 716 |
| show vpls statistics | 720 |

Part 6**Index**

| | |
|----------------------------------------|-----|
| Index | 725 |
| Index of Statements and Commands | 735 |

About This Guide

This preface provides the following guidelines for using the *JUNOS® Software Routing Protocols and Policies Command Reference*:

- JUNOS Documentation and Release Notes on page xvii
- Objectives on page xvii
- Audience on page xviii
- Supported Platforms on page xix
- Using the Indexes on page xix
- Documentation Conventions on page xix
- Documentation Feedback on page xxi
- Requesting Technical Support on page xxi

JUNOS Documentation and Release Notes

For a list of related JUNOS 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 Software 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 JUNOS Software 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 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 Software commands that you use to monitor and troubleshoot routing protocols, protocol-independent

features, and policies, including firewall filters, forwarding options, and routing policies.

For additional commands, see these guides:

- *JUNOS System Basics and Services Command Reference*
- *JUNOS Interfaces Command Reference*



NOTE: For additional information about JUNOS Software—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 Routing Protocols Configuration Guide*—Includes configuration statements and guidelines for routing protocols and protocol-independent features.
- *JUNOS Policy Framework Configuration Guide*—Includes configuration statements and guidelines for policies, including firewall filters, forwarding options, and routing policies.
- *JUNOS MPLS Applications Configuration Guide*—Includes configuration statements and guidelines for Multiprotocol Label Switching (MPLS) traffic engineering.
- *JUNOS 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, JUNOS Software currently supports the following platforms:

- J Series
- M Series
- MX 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 xix 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 xx 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 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 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> |
| Plain text like this | Represents names of configuration statements, commands, files, and directories; IP addresses; 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> >; |
| (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 [<i>community-ids</i>] |

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

| Convention | Description | Examples |
|--------------------------------|----------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Indentation and braces ({ }) | Identify a level in the configuration hierarchy. | [edit] routing-options { static { route default { nexthop address; 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 name
- Document part number
- Page number
- Software release version (not required for Network Operations Guides [NOGs])

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/customers/support/downloads/710059.pdf>.
- Product warranties—For product warranty information, visit <http://www.juniper.net/support/warranty/>.
- JTAC Hours of Operation —The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

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

- Find CSC offerings: <http://www.juniper.net/customers/support/>
- Search for known bugs: <http://www2.juniper.net/kb/>
- Find product documentation: <http://www.juniper.net/techpubs/>
- Find solutions and answer questions using our Knowledge Base: <http://kb.juniper.net/>
- Download the latest versions of software and review release notes: <http://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <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 located at <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 15
- BGP Operational Mode Commands on page 25
- ES-IS Operational Mode Commands on page 57
- IP Multicast Operational Mode Commands on page 65
- IPv6 Operational Mode Commands on page 203
- IS-IS Operational Mode Commands on page 211
- LLDP Operational Mode Commands on page 255
- OSPF Operational Mode Commands on page 269
- Protocol-Independent Routing Operational Mode Commands on page 321
- RIP Operational Mode Commands on page 487
- RIPng Operational Mode Commands on page 495

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. | clear ancp neighbor |
| Display ANCP class-of-service information. | show ancp cos |
| Display ANCP neighbor information. | show ancp neighbor |
| Display ANCP subscriber information. | show ancp subscriber |



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

clear ancp neighbor

| Syntax | clear ancp neighbor <ip-address <i>ip-address</i> > <system-name <i>mac-address</i> > | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|------------------|--------------|------------------|--------------|------------|-------------------|-------------|---|------|------------|-------------------|-------------|---|------|------------|-------------------|-------------|---|------|------------|-------------------|-------------|---|------|------------|-------------|-------|------------------|--------------|------------|-------------------|-------------|---|------|------------|-------------------|-------------|---|------|------------|-------------------|-------------|---|------|
| Release Information | Command introduced in JUNOS Release 9.4. | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| Description | Clear the connection with all ANCP neighbors or with the specified ANCP neighbor. | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| Options | 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 Topics | show ancp neighbor | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| List of Sample Output | clear ancp neighbor on page 4 clear ancp neighbor on page 4 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 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. | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| clear ancp neighbor | <p>The following sample output displays the connections with ANCP neighbors before and after the clear ancp neighbor command was issued.</p> <pre>user@host> show ancp neighbor</pre> <table><tr><th>IP Address</th><th>MAC Address</th><th>State</th><th>Subscriber Count</th><th>Capabilities</th></tr><tr><td>10.10.10.2</td><td>ba:ad:be:ef:10:10</td><td>Established</td><td>5</td><td>Topo</td></tr><tr><td>12.12.12.2</td><td>ba:ad:be:ef:10:12</td><td>Established</td><td>5</td><td>Topo</td></tr><tr><td>13.13.13.2</td><td>ba:ad:be:ef:10:13</td><td>Established</td><td>5</td><td>Topo</td></tr><tr><td>14.14.14.2</td><td>ba:ad:be:ef:10:14</td><td>Established</td><td>5</td><td>Topo</td></tr></table> <pre>user@host> clear ancp ip-address 10.10.10.2 user@host> show ancp neighbor</pre> <table><tr><th>IP Address</th><th>MAC Address</th><th>State</th><th>Subscriber Count</th><th>Capabilities</th></tr><tr><td>12.12.12.2</td><td>ba:ad:be:ef:10:12</td><td>Established</td><td>5</td><td>Topo</td></tr><tr><td>13.13.13.2</td><td>ba:ad:be:ef:10:13</td><td>Established</td><td>5</td><td>Topo</td></tr><tr><td>14.14.14.2</td><td>ba:ad:be:ef:10:14</td><td>Established</td><td>5</td><td>Topo</td></tr></table> | 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 | 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 |
| 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 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 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 neighbor | user@host> clear ancp neighbor | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |

show ancp cos

| | |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show ancp cos <identifier <i>identifier</i> > <last-update> <pending-update> |
| Release Information | Command introduced in JUNOS Release 9.4. |
| Description | Display information about the CoS state for subscriber traffic. |
| Options | <p><i>identifier identifier</i>—(Optional) Display information about the local loops for the specified access identifier.</p> <p><i>last-update</i>—(Optional) Display the most recently updated CoS information.</p> <p><i>pending-update</i>—(Optional) Display the pending update of CoS information.</p> |
| Required Privilege Level | view |
| List of Sample Output | <p>show ancp cos on page 6</p> <p>show ancp cos last-update on page 6</p> <p>show ancp cos pending-update on page 6</p> |
| Output Fields | Table 4 on page 5 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. |
| 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. |

Table 4: show ancp cos Output Fields (continued)

| Field Name | Field Description |
|------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Rate Update Time | Time at which the shaping rate was last updated. |
| Type | Subscriber access type: <code>ifl</code> indicates that a single VLAN carries subscriber traffic and <code>iflset</code> 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. |

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 Release 9.4. |
| Description | Display information about all ANCP neighbors or the specified ANCP neighbor. |
| Options | <p>brief detail extensive terse—(Optional) Display the specified level of detail.</p> <p>ip-address <i>ip-address</i>—(Optional) IP address of the ANCP neighbor (access node).</p> <p>system-name <i>mac-address</i>—(Optional) MAC address of the ANCP neighbor (access node).</p> |
| Required Privilege Level | view |
| Related Topics | clear ancp neighbor |
| List of Sample Output | <p>show ancp neighbor on page 9</p> <p>show ancp neighbor ip-address on page 9</p> |
| Output Fields | Table 5 on page 8 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 | <p>State of the ANCP adjacency:</p> <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). |
| Capabilities | Negotiated ANCP capability. Currently, only topology discovery is available. |
| TCP Port | TCP port on which ANCP messages are exchanged. |
| System Instance | Number identifying the ANCP link instance from the edge device's perspective. |

Table 5: show ancp neighbor Output Fields (continued)

| Field Name | Field Description |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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. |

show ancp neighbor user@host> **show ancp neighbor**

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

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 subscriber

| | |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show ancp subscriber <brief detail extensive terse> <identifier <i>identifier</i> > <neighbor <i>ip-address</i> > |
| Release Information | Command introduced in JUNOS 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 extensive terse—(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 specified access node.</p> |
| Required Privilege Level | view |
| List of Sample Output | <p>show ancp subscriber brief on page 12</p> <p>show ancp subscriber neighbor ip-address detail on page 12</p> <p>show ancp subscriber identifier identifier detail on page 13</p> |
| Output Fields | Table 6 on page 10 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). |
| 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. |

Table 6: show ancp subscriber Output Fields (continued)

| Field Name | Field Description |
|--------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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. |
| 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 |

Table 6: show ancp subscriber Output Fields (continued)

| Field Name | Field Description |
|-------------------------------------|-------------------------------------------------------------------------------------------------------|
| 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. |

**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
neighbor ip-address
detail**user@host> **show ancp subscriber neighbor 10.11.11.2 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 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 15 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. | clear bfd adaptation |
| Clear BFD sessions. | clear bfd session |
| Display BFD session statistics. | show bfd session |



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 | clear bfd adaptation <address <i>session-address</i> > <discriminator <i>discr-number</i> > |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | 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. |
| 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 16 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| 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>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| 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 Topics | show bfd session |
| List of Sample Output | clear bfd session on page 17 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| clear bfd session | user@host> clear bfd session |

show bfd session

| | |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | <pre>show bfd session <brief detail extensive summary> <address address> <discriminator discriminator> <logical-system (all logical-system-name)> <prefix address></pre> |
| Release Information | <p>Command introduced before JUNOS Release 7.4.</p> <p>The discriminator and address options were introduced in JUNOS Release 8.2.</p> <p>The prefix option was introduced in JUNOS Release 9.0.</p> |
| Description | Display information about active Bidirectional Forwarding Detection (BFD) sessions. |
| Options | <p>none—(Same as brief) Display information about active BFD sessions.</p> <p>brief detail extensive summary—(Optional) Display the specified level of output.</p> <p>address address—(Optional) Display information about the BFD session for the specified neighbor address.</p> <p>discriminator discriminator—(Optional) Display information about the BFD session using the specified local discriminator.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>prefix address—(Optional) Display information about all of the BFD sessions for the specified LDP forwarding equivalence class (FEC).</p> |
| Required Privilege Level | view |
| Related Topics | clear bfd session |
| List of Sample Output | <p>show bfd session on page 21</p> <p>show bfd session brief on page 22</p> <p>show bfd session detail on page 22</p> <p>show bfd session detail (with Authentication) on page 22</p> <p>show bfd session address extensive on page 22</p> <p>show bfd session extensive on page 23</p> <p>show bfd session extensive (with Authentication) on page 23</p> <p>show bfd session summary on page 23</p> |
| Output Fields | Table 8 on page 19 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 may be configured on a BFD session. | extensive |
| algo | BFD authentication algorithm being used for a specific client: keyed-md5 , keyed-sha-1 , meticulous-keyed-md5 , meticulous-keyed-sha-1 , or simple-password . BFD authentication information for a client is provided in a single line and includes the keychain , algo , and mode parameters. Multiple clients may be configured on a BFD session. | extensive |
| mode | Level of BFD authentication enforcement being used by a specific client: strict or loose . Strict enforcement indicates authentication is configured at both ends of the session (the default). Loose enforcement indicates that one end of the session may not be authenticated. BFD authentication information for a client is provided in a single line and includes the keychain , algo , and mode parameters. Multiple clients may be configured on a BFD session. | extensive |
| Local diagnostic | Local diagnostic information about failing BFD sessions. | detail extensive |

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

| Field Name | Field Description | Level of Output |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| 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 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 |
| 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 |

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

| Field Name | Field Description | Level of Output |
|-------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| 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 may not be authenticated. <p>This information is only shown if BFD authentication is configured.</p> | extensive |
| sessions | Total number of active BFD sessions. | All levels |
| clients | Total number of clients that are hosting active BFD sessions. | All levels |
| Cumulative transmit rate | Total number of BFD control packets transmitted per second on all active sessions. | All levels |
| Cumulative receive rate | Total number of BFD control packets received per second on all active sessions. | All levels |
| Multi-hop, min-recv-TTL | Minimum time to live (TTL) accepted if the session is configured for multihop. | extensive |
| route table | Route table used if the session is configured for multihop. | extensive |
| local address | Local address of source used if the session is configured for multihop. | extensive |

```

show bfd session  user@host> show bfd session

                                Transmit
Address      State   Interface  Detect Time  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 21.

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
                                     Transmit
Address      State      Interface      Detect Time  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
  Min async interval 0.200, min slow interval 1.000
  Adaptive async tx interval 0.200, rx interval 0.200
  Local min tx interval 0.200, min rx interval 0.200, multiplier 3
  Remote min tx interval 0.100, min rx interval 0.100, multiplier 3
  Threshold transmission interval 0.000, Threshold for detection time 0.000
  Local discriminator 11, remote discriminator 80
  Echo mode disabled/inactive

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 00:30, previous down time 00:00:01
  Local diagnostic NbrSignal, remote diagnostic AdminDown
  Remote state Up, version 1
  Replicated
  Min async interval 0.200, min slow interval 1.000
  Adaptive async tx interval 0.200, rx interval 0.200
  Local min tx interval 0.200, min rx interval 0.200, multiplier 3
  Remote min tx interval 0.200, min rx interval 0.200, multiplier 3
  Threshold transmission interval 0.000, Threshold for detection time 0.000
  Local discriminator 12, remote discriminator 11
  Echo mode disabled/inactive
  Remote is control-plane independent

2 sessions, 2 clients

Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps

show bfd session extensive (with Authentication)
user@host>show bfd session extensive
                                     Detect   Transmit
Address      State      Interface      Time     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 25 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 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” on page 321.



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> > |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Clear Border Gateway Protocol (BGP) route flap damping information. |
| Options | <p>none—Clear all BGP route flap damping information.</p> <p><i>logical-system</i> (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><i>prefix</i>—(Optional) Clear route flap damping information for only the specified destination prefix.</p> |
| Required Privilege Level | clear |
| Related Topics | show policy damping show route damping |
| List of Sample Output | clear bgp damping on page 26 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| clear bgp damping | user@host> clear bgp damping |

clear bgp neighbor

Syntax clear bgp neighbor
 <as *as-number*>
 <instance *instance-name*>
 <logical-system (all | *logical-system-name*)>
 <neighbor>
 <soft | soft-inbound>
 <soft-minimum-igp>

Release Information Command introduced before JUNOS Release 7.4.

Description Perform one of the following tasks:

- Change the state of one or more Border Gateway Protocol (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 none—Change the state of all BGP neighbors to IDLE.

as *as-number*—(Optional) Apply this command only to neighbors in the specified autonomous system (AS).

instance *instance-name*—(Optional) Apply this command only to neighbors for the specified routing instance.

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

neighbor—(Optional) IP address of a BGP peer. Apply this command only to the specified neighbor.

soft—(Optional) Reapply any export policies and send refresh updates to neighbors without clearing the state.

soft-inbound—(Optional) Reapply any import policies and send refresh updates to neighbors without clearing the state.

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

Required Privilege Level clear

Related Topics show bgp neighbor

List of Sample Output clear bgp neighbor on page 28

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

clear bgp neighbor user@host> **clear bgp neighbor**

clear bgp table

| | |
|---------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | clear bgp table <i>table-name</i> <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced in JUNOS Release 9.0. |
| Description | Request BGP to refresh routes in a specified routing table. |
| Options | <i>table-name</i> —Request BGP to refresh routes in the specified table. (logical-system (all <i>logical-system-name</i>))—(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Additional Information | In some cases, a prefix limit is associated with a routing table for a VPN instance . When this limit is exceeded, as for example, because of a network misconfiguration, some routes might not get 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 BGP to refresh routes in a VPN instance table. |
| Required Privilege Level | clear |
| List of Sample Output | clear bgp table private.inet.0 on page 29 |
| Output Fields | This command produces no output. |
| clear bgp table private.inet.0 | user@host> clear bgp table private.inet.0 |

show bgp bmp

| | |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show bgp bmp |
| Release Information | Command introduced in JUNOS Release 9.5. |
| 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 30 |
| Output Fields | Table 10 on page 30 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 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 the transmission of BMP data to the monitoring station.. |
| Memory limit | Threshold, in bytes, at which the router stops collecting BMP data if it is exceeded. |
| Memory-connect retry timeout | Amount of time, in seconds, after which the router attempts to resume a BMP session that was ended after the configured memory threshold was exceeded. |

```

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 | show bgp group <brief detail summary> <group-name> <instance instance-name> <logical-system (all logical-system-name)> <rtf> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display information about the configured Border Gateway Protocol (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>instance instance-name—(Optional) Display information about a particular BGP peer in the specified instance. 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 34</p> <p>show bgp group on page 34</p> <p>show bgp group brief on page 35</p> <p>show bgp group detail on page 35</p> <p>show bgp group rtf detail on page 35</p> <p>show bgp group summary on page 36</p> <p>show bgp group summary on page 36</p> |
| Output Fields | Table 11 on page 31 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 |
| 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 router. | brief detail none |

Table 11: show bgp group Output Fields (continued)

| Field Name | Field Description | Level of Output |
|------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------|
| Name | Name of a specific BGP group. | brief detail none |
| Flags | Flags associated with the BGP group. This field is used by Juniper Networks Customer Support. | 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 interval, in seconds, that updates to multiple exit discriminator (MED) are delayed. Also displays the time remaining before the interval is set to expire | All |
| 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 |
| 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 an established session was established in the main router 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 router, 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 2/4/4/0 indicates 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 |
| <i>ip-addresses</i> | 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 |
| <i>inet.number</i> | <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 | <p>Information about the routing table.</p> <ul style="list-style-type: none"> ■ Received prefixes—Total number of prefixes from the peer, both active and inactive, that are in the routing table. ■ Active prefixes—Number of prefixes received from the peer that are active in the routing table. ■ Suppressed due to damping—Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols. ■ Advertised prefixes—Number of prefixes advertised to a peer. ■ Received external prefixes—Total number of prefixes from the external BGP (EBGP) peers, both active and inactive, that are in the routing table. ■ Active external prefixes—Number of prefixes received from the EBGP peers that are active in the routing table. ■ Externals suppressed—Number of routes received from EBGP peers currently inactive because of damping or other reasons. ■ Received internal prefixes—Total number of prefixes from the IBGP peers, both active and inactive, that are in the routing table. ■ Active internal prefixes—Number of prefixes received from the IBGP peers that are active in the routing table. ■ Internals suppressed—Number of routes received from IBGP peers currently inactive because of damping or other reasons. ■ RIB State—Status of the graceful restart process for this routing table: BGP restart is complete, BGP restart in progress, VPN restart in progress, or VPN restart is complete. | detail |
| Groups | Total number of groups. | All levels |
| Peers | Total number of peers. | All levels |
| External | Total number of external peers. | All levels |
| Internal | Total number of internal peers. | All levels |
| Down peers | Total number of unavailable peers. | All levels |
| Flaps | Total number of flaps that occurred. | All levels |
| Table | Name of a routing table. | brief, none |
| Tot Paths | Total number of paths. | 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 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) router 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 |

```

show bgp group user@host> show bgp group
Group Type: Internal      AS: 21                      Local AS: 21
Name: from_vpn04_to_other Index: 0          Flags: <>
Holdtime: 0
Total peers: 3           Established: 3
10.255.14.178+179
10.255.71.24+179
10.255.14.182+179
inet.0: 2/7/0

Group Type: External
Name: from_vpn04_to_vpn06 Index: 1          Local AS: 21
Export: [ internal-and-bgp ]                Flags: <Export Eval>
Holdtime: 0
Traffic Statistics Interval: 300
Total peers: 1           Established: 1
100.1.3.2+2910
inet.0: 5/10/0

Groups: 2  Peers: 4   External: 1   Internal: 3   Down peers: 0   Flaps: 2
Table      Tot Paths  Act Paths  Suppressed  History  Damp State  Pending
inet.0           17           7           0           0           0           0

show bgp group user@host> show bgp group
Group Type: External                      Local AS: 65500
Name: as65501peers      Index: 0          Flags: Export <Eval>
Export: [ export-policy ]
Holdtime: 0
Total peers: 1           Established: 1
192.168.4.222+179
Trace options: all

```



```
Trace file: /var/log/bgp size 10485760 files 10
inet.0: 7/10/9/0
inet.2: 0/0/0/0
```

| | | | | | |
|-----------|-----------|-------------|-------------|---------------|--------------------|
| Groups: 1 | Peers: 1 | External: 1 | Internal: 0 | Down peers: 0 | Flaps: 0 |
| Table | Tot Paths | Act Paths | Suppressed | History | Damp State Pending |
| inet.0 | 10 | 7 | 0 | 0 | 0 0 |
| inet.2 | 0 | 0 | 0 | 0 | 0 0 |

show bgp group brief The output for the `show bgp group brief` command is identical to that for the `show bgp group` command. For sample output, see `show bgp group` on page 34.

```
user@host> show bgp group detail
Group Type: Internal      AS: 21                      Local AS: 21
Name: from_vpn04_to_other Index: 0      Flags: <>
Holdtime: 0
Total peers: 3            Established: 3
10.255.14.178+179
10.255.71.24+179
10.255.14.182+179
Route Queue Timer: unset Route Queue: empty
Table inet.0
  Active prefixes:        2
  Received prefixes:      7
  Suppressed due to damping: 0
  Advertised prefixes:    5

Group Type: External      Local AS: 21
Name: from_vpn04_to_vpn06 Index: 1      Flags:<Export Eval>
Export: [ internal-and-bgp ]
Holdtime: 0
Traffic Statistics Interval: 300
Total peers: 1            Established: 1
100.1.3.2+2910
Route Queue Timer: unset Route Queue: empty
Table inet.0
  Active prefixes:        5
  Received prefixes:      10
  Suppressed due to damping: 0
  Advertised prefixes:    6

Groups: 2 Peers: 4 External: 1 Internal: 3 Down peers: 0 Flaps: 2
Table inet.0
  Received prefixes:      17
  Active prefixes:        7
  Suppressed due to damping: 0
  Received external prefixes: 10
  Active external prefixes: 5
  Externals suppressed:   0
  Received internal prefixes: 7
  Active internal prefixes: 2
  Internals suppressed:   0
RIB State: BGP restart is complete
```

```
user@host> show bgp group rtf detail
Group: asbr
Receive mask: 00000001
Table: bgp.rtarget.0      Flags: Filter Entries: 4
Target                    Mask
```

| | |
|------------|----------|
| 109:1/64 | 00000001 |
| 109:2/64 | 00000001 |
| 701:1/64 | 00000001 |
| 10458:2/64 | 00000001 |

Group: mesh_0

Receive mask: 0000000e

Table: bgp.rtarget.0

Flags: Filter Entries: 12

| Target | Mask |
|-------------|----------|
| 109:1/64 | 00000002 |
| 701:1/64 | 00000002 |
| 701:2/64 | 00000002 |
| 10458:1/64 | 0000000e |
| 10458:2/64 | 00000006 |
| 10458:3/64 | 00000006 |
| 10458:5/64 | 00000006 |
| 10458:6/64 | 00000004 |
| 10458:7/64 | 00000008 |
| 10458:8/64 | 00000008 |
| 10458:10/64 | 00000002 |

**show bgp group
summary**

user@host> show bgp group summary

| Group | Type | Peers | Established | Active/Received/Damped |
|---------------------|----------|----------|-------------|------------------------|
| from_vpn04_to_other | Internal | 3 | 3 | |
| inet.0 | | : 2/7/0 | | |
| from_vpn04_to_vpn06 | External | 1 | 1 | |
| inet.0 | | : 5/10/0 | | |

Groups: 2 Peers: 4 External: 1 Internal: 3 Down peers: 0 Flaps: 2
inet.0 : 7/17/0 External: 5/10/0 Internal: 2/7/0

**show bgp group
summary**

user@host> show bgp group summary

| Group | Type | Peers | Established | Active/Received/Accepted/Damped |
|-------------------------------------------------|----------|------------|-------------|---------------------------------|
| as65501peers | External | 1 | 1 | |
| Trace options: all | | | | |
| Trace file: /var/log/bgp size 10485760 files 10 | | | | |
| inet.0 | | : 7/10/9/0 | | |
| inet.2 | | : 0/0/0/0 | | |

Groups: 1 Peers: 1 External: 1 Internal: 0 Down peers: 0 Flaps: 0
inet.0 : 7/10/9/0 External: 7/10/9/0 Internal: 0/0/0/0
inet.2 : 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 Release 7.4. |
| Description | Display the traffic statistics for configured Border Gateway Protocol (BGP) groups. |
| Options | none—Display traffic statistics for all BGP groups. brief detail—(Optional) Display the specified level of output. <i>group-name</i> —(Optional) Display BGP traffic statistics for only the specified group. 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 bgp group traffic-statistics (Per-Group-Label Not Configured) on page 37 show bgp group traffic-statistics (Per-Group-Label Configured) on page 38 |
| Output Fields | Table 12 on page 37 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. |

```

show bgp group      user@host> show bgp group traffic-statistics
traffic-statistics Group Name: ext1      Group Index: 0      NLRI: inet-labeled-unicast
                    FEC                      Packets          Bytes      EgressAS    AdvLabel

```

```

(Per-Group-Label Not Configured) 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 user@host> show bgp group traffic-statistics

```

```

(Per-Group-Label Configured) 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 | show bgp neighbor <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> <neighbor-address> <orf (<i>neighbor-address</i> detail) |
| Release Information | Command introduced before JUNOS Release 7.4. orf option introduced in JUNOS Release 9.2. |
| Description | Display information about Border Gateway Protocol (BGP) peers. |
| Options | none—Display information about all BGP peers. instance <i>instance-name</i> —(Optional) Display information about BGP peers 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. neighbor-address—(Optional) Display information for only the BGP peer at the specified IP address. orf (<i>neighbor-address</i> detail)—(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. |
| Additional Information | For information about the local-address, nlri, hold-time, and preference statements, see the <i>JUNOS Routing Protocols Configuration Guide</i> . |
| Required Privilege Level | view |
| Related Topics | clear bgp neighbor |
| List of Sample Output | show bgp neighbor (CLNS) on page 44 show bgp neighbor (Layer 2 VPN) on page 45 show bgp neighbor (Layer 3 VPN) on page 47 show bgp neighbor neighbor-address on page 48 show bgp neighbor neighbor-address on page 48 show bgp neighbor orf neighbor-address detail on page 49 |
| Output Fields | Table 13 on page 39 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's port number. |
| AS | AS number of the peer. |

Table 13: show bgp neighbor Output Fields (continued)

| Field Name | Field Description |
|------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Local | Address of the local router. The address is followed by the peer's 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, 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. |
| 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. |

Table 13: show bgp neighbor Output Fields (continued)

| Field Name | Field Description |
|------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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 router sent a BGP keepalive message to the peer. ■ Open—The local router sent a BGP open message to the peer. ■ OpenFail—The local router did not receive an acknowledgment of a BGP open message from the peer. ■ RecvKeepAlive—The local router received a BGP keepalive message from the peer. ■ RecvNotify—The local router received a BGP notification message from the peer. ■ RecvOpen—The local router received a BGP open message from the peer. ■ RecvUpdate—The local router 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. |
| Options | <p>Configured BGP options:</p> <ul style="list-style-type: none"> ■ AddressFamily—Configured address family: <code>inet</code> or <code>inet-vpn</code>. ■ AuthKeyChain—Authentication key chain is enabled. ■ GracefulRestart—Graceful restart is configured. ■ HoldTime—Hold time configured with the <code>hold-time</code> statement. The hold time is three times the interval at which keepalive messages are sent. ■ Local Address—Address configured with the <code>local-address</code> statement. ■ Multihop—Allow BGP connections to external peers that are not on a directly connected network. ■ NLRI—Configured MBGP state for the BGP group: <code>multicast</code>, <code>unicast</code>, or both if you have configured <code>nlri</code> any. ■ Peer AS—Configured peer autonomous system (AS). ■ Preference—Preference value configured with the <code>preference</code> statement. ■ Refresh—Configured to refresh automatically when the policy changes. ■ Rib-group—Configured routing table group. |

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

| Field Name | Field Description |
|-------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Authentication key change | Name of the authentication key chain enabled. |
| Authentication algorithm | 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 router. |
| 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. |
| 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. |
| Number of flaps | Number of times the BGP session has gone down and then come back up. |
| Peer ID | Router identifier of the peer. |
| Peer Index | Index of this peer in its group. |
| Local ID | Router identifier of the local router. |
| Local Interface | Name of the interface on the local router. |
| Active holdtime | Hold time the local router 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 the 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. |

Table 13: show bgp neighbor Output Fields (continued)

| Field Name | Field Description |
|-------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Peer supports Refresh capability | Remote peer supports Route Refresh Capability, and is able to send and request full route table readvertisement. 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. |
| 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. |
| 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. |
| 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 router checked. |

Table 13: show bgp neighbor Output Fields (continued)

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

```

show bgp neighbor      user@host> show bgp neighbor
(CLNS)                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 aaa.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 l2vpn
NLRI peer can save forwarding state: inet-vpn-unicast l2vpn
NLRI that peer saved forwarding for: inet-vpn-unicast l2vpn
NLRI that restart is negotiated for: inet-vpn-unicast l2vpn
NLRI of received end-of-rib markers: inet-vpn-unicast l2vpn
Table bgp.l3vpn.0 Bit: 10000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          10
  Received prefixes:        10
  Suppressed due to damping: 0
Table bgp.l2vpn.0 Bit: 20000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        1
  Suppressed due to damping: 0
Table BGP-INET.inet.0 Bit: 30000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2
  Suppressed due to damping: 0
Table BGP-L.inet.0 Bit: 40000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2
  Suppressed due to damping: 0
Table LDP.inet.0 Bit: 50000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        1
  Suppressed due to damping: 0
Table OSPF.inet.0 Bit: 60000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2
  Suppressed due to damping: 0
Table RIP.inet.0 Bit: 70000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2
  Suppressed due to damping: 0
Table STATIC.inet.0 Bit: 80000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          1

```

```

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

```

**show bgp neighbor
neighbor-address**

```

user@host> show bgp neighbor 192.168.4.222
Peer: 192.168.4.222+4902 AS 65501 Local: 192.168.4.221+179 AS 65500
Type: External State: Established Flags: <Sync>

```

```

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

| | |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show bgp summary <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display Border Gateway Protocol (BGP) summary information. |
| Options | <p>none—Display BGP summary information for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Display information for the specified instance only. 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 52</p> <p>show bgp summary (When a Peer Is Established) on page 52</p> <p>show bgp summary (CLNS) on page 52</p> <p>show bgp summary (Layer 2 VPN) on page 52</p> <p>show bgp summary (Layer 3 VPN) on page 53</p> |
| Output Fields | Table 14 on page 50 describes the output fields for the show bgp summary command. Output fields are listed in the approximate order in which they appear. |

Table 14: show bgp summary Output Fields

| Field Name | Field Description |
|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Groups | Number of BGP groups. |
| Peers | Number of BGP peers. |
| Down peers | Number of down BGP peers. |
| Table | Name of routing table. |
| Tot Paths | Total number of paths. |
| Act Paths | Number of active routes. |
| Suppressed | Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols. |
| History | Number of withdrawn routes stored locally to keep track of damping history. |
| Damp State | Number of routes with a figure of merit greater than zero, but still active because the value has not reached the threshold at which suppression occurs. |

Table 14: show bgp summary Output Fields (continued)

| Field Name | Field Description |
|------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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 | Count of the 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. |
| 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 an established session was established in the main router 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 router, 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 2/4/4/0 indicates 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 router 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> |

**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        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        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        0
inet.0          0         0          0          0        0      0        0
Peer        AS      InPkt    OutPkt    OutQ    Flaps Last Up/Dwn
Up/Dwn State|#Active/Received/Damped...
10.255.245.35 65299        72        74        0        1      19:00 Establ
  bgp.l2vpn.0: 1/1/0
  frame-vpn.l2vpn.0: 1/1/0
10.255.245.36 65299     2164     2423        0        4      19:50 Establ
  bgp.l2vpn.0: 0/0/0
  frame-vpn.l2vpn.0: 0/0/0
10.255.245.37 65299        36        37        0        4      17:07 Establ
  inet.0: 0/0/0
10.255.245.39 65299       138       168        0        6      53:48 Establ
  bgp.l2vpn.0: 0/0/0
  frame-vpn.l2vpn.0: 0/0/0
10.255.245.69 65299       134       140        0        6      53:42 Establ
  inet.0: 0/0/0

```

show bgp summary
(Layer 3 VPN)

```
user@host> show bgp summary
Groups: 2 Peers: 2 Down peers: 0
Table      Tot Paths  Act Paths Suppressed  History  Damp State Pending
bgp.l3vpn.0      2      2      0      0      0      0      0
Peer          AS      InPkt      OutPkt      OutQ      Flaps Last Up/Dwn
State|#Active/Received/Damped...
10.39.1.5      2      21      22      0      0      6:26 Establ
  VPN-AB.inet.0: 1/1/0
10.255.71.15   1      19      21      0      0      6:17 Establ
  bgp.l3vpn.0: 2/2/0
  VPN-A.inet.0: 1/1/0
  VPN-AB.inet.0: 2/2/0
  VPN-B.inet.0: 1/1/0
```

show policy damping

| | |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show policy damping <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display information about Border Gateway Protocol (BGP) route flap damping parameters. |
| Options | <p>none—(Same as logical-system all) Display information about BGP route flap damping parameters for all logical systems.</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 to the probability of future instability of a router. 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 Topics | clear bgp damping show route damping “Configuring BGP Flap Damping Parameters” in the <i>JUNOS Policy Framework Configuration Guide</i> |
| List of Sample Output | show policy damping on page 55 |
| Output Fields | Table 15 on page 54 describes the output fields for the show policy damping command. Output fields are listed in the approximate order in which they appear. |

Table 15: show policy damping Output Fields

| Field Name | Field Description |
|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Half-life | 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. |
| 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. |

Table 15: show policy damping Output Fields (*continued*)

| Field Name | Field Description |
|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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. |

```

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

```


Chapter 4

ES-IS Operational Mode Commands

Table 16 on page 57 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 16: 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 Routing Protocols Configuration Guide*.

clear esis adjacency

| | |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | clear esis adjacency <instance <i>instance-name</i> > <interface <i>interface-name</i> > < <i>neighbor</i> > |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | (J Series routers only) Clear End System-to-Intermediate System (ES-IS) adjacencies. |
| Options | <p>none—Clear all ES-IS adjacencies.</p> <p>instance <i>instance-name</i>—(Optional) Clear adjacencies for the specified routing instance only.</p> <p>interface <i>interface-name</i>—(Optional) Clear adjacencies for the specified interface only.</p> <p><i>neighbor</i>—(Optional) Clear adjacencies for the specified neighbor only.</p> |
| Required Privilege Level | clear |
| Related Topics | show esis adjacency |
| List of Sample Output | clear esis adjacency on page 58 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| 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 Release 7.4. |
| Description | (J Series routers only) 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 Topics | show esis statistics |
| List of Sample Output | clear esis statistics on page 59 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| 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 Release 7.4. |
| Description | (J 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 Topics | clear esis adjacency |
| List of Sample Output | <p>show esis adjacency on page 61</p> <p>show esis adjacency brief on page 61</p> <p>show esis adjacency detail on page 61</p> <p>show esis adjacency extensive on page 61</p> |
| Output Fields | Table 17 on page 60 describes the output fields for the show esis adjacency command. Output fields are listed in the approximate order in which they appear. |

Table 17: 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 |
| Advertised holdtime | Holdtime interval advertised by this neighbor. | detail extensive |

Table 17: show esis adjacency Output Fields (continued)

| Field Name | Field Description | Level of Output |
|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| 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 |

show esis adjacency user@host> **show esis adjacency**

| | | | |
|------|-------------------------------------------------|--------|------------|
| Nbr | NSAP/NET | Hold | Interface |
| Type | | (secs) | |
| IS | 47.0005.80ff.f800.0000.0108.0001.0102.5501.6008 | 135 | fe-0/0/0.0 |

show esis adjacency brief The output for the show esis adjacency brief command is identical to that for the show esis adjacency command. For sample output, see show esis adjacency on page 61.

show esis adjacency detail user@host> **show esis adjacency detail**

NSAP/NET: 47.0005.80ff.f800.0000.0108.0001.0102.5501.6008, Type: IS
Interface: fe-0/0/0.0, Advertised hold time: 180 secs, Expires in: 173 secs
SNPA: 0:5:85:c1:73:71

show esis adjacency extensive user@host> **show esis adjacency extensive**

NSAP/NET: 47.0005.80ff.f800.0000.0108.0001.0102.5501.6008, Type: IS
Interface: fe-0/0/0.0, Advertised hold time: 180 secs, Expires in: 167 secs
SNPA: 0:5:85:c1:73:71
Transition log:

| | | | |
|---------------------|-------|--------------|-------------|
| When | State | Event | Down reason |
| Sun Nov 26 22:07:35 | Up | Received ISH | |

show esis interface

| | |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show esis interface <brief detail extensive> <instance <i>instance-name</i> > <interface <i>interface-name</i> > |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | (J 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 | <p>show esis interface on page 63</p> <p>show esis interface brief on page 63</p> <p>show esis interface detail on page 63</p> <p>show esis interface extensive on page 63</p> |
| Output Fields | Table 18 on page 62 describes the output fields for the show esis interface command. Output fields are listed in the approximate order in which they appear. |

Table 18: 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 |
| End system configuration timer | Time, in seconds, for the end system to configure itself for ES-IS. | detail extensive |

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

| Field Name | Field Description | Level of Output |
|-------------------|----------------------------------------------|------------------|
| Interface index | Index value. | detail extensive |
| NET used in hello | Network entity title used in hello messages. | detail extensive |

```

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 63.

```

show esis interface user@host> show esis interface detail
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 63.

show esis statistics

| | |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show esis statistics <instance <i>instance-name</i> > |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | (J Series routers only) Display End System-to-Intermediate System (ES-IS) packet statistics. |
| Options | none—Display ES-IS packet statistics for all routing instances. instance <i>instance-name</i> —(Optional) Display ES-IS statistics for the specified routing instance only. |
| Required Privilege Level | view |
| Related Topics | clear esis statistics |
| List of Sample Output | show esis statistics on page 64 |
| Output Fields | Table 19 on page 64 describes the output fields for the show esis statistics command. Output fields are listed in the approximate order in which they appear. |

Table 19: 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. |

```

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 20 on page 65 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 20: IP Multicast Operational Mode Commands

| Task | Command |
|--------------------------------------------------------------------------|--------------------------------------------------|
| 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 multicast scope. | <code>clear multicast scope</code> |
| Clear multicast snooping statistics. | <code>clear multicast snooping statistics</code> |
| Clear multicast statistics. | <code>clear multicast statistics</code> |
| Clear multicast sessions. | <code>clear multicast sessions</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> |

Table 20: IP Multicast Operational Mode Commands *(continued)*

| Task | Command |
|-------------------------------------------------------------------------------------------------------------|--------------------------------------------|
| Clear the Protocol Independent Multicast (PIM) join and prune states. | <code>clear pim join</code> |
| Clear PIM register message counters. | <code>clear pim register</code> |
| Clear PIM statistics. | <code>clear pim statistics</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> |
| 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 configuration information about IP multicast flow maps. | <code>show multicast flow-map</code> |
| Display multicast interface bandwidth information. | <code>show multicast interface</code> |
| Display multicast network configuration. | <code>show multicast moinfo</code> |
| Display entries in the multicast next-hop table. | <code>show multicast next-hops</code> |

Table 20: IP Multicast Operational Mode Commands *(continued)*

| Task | Command |
|-----------------------------------------------------------------------------------------------------------|-------------------------------------------------|
| 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> |
| Display administratively scoped addresses. | <code>show multicast scope</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 PIM neighbors. | <code>show pim neighbors</code> |
| Display rendezvous points. | <code>show pim rps</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>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Clear Internet Group Management Protocol (IGMP) group members. |
| Options | <p>none—Clear all IGMP members on all interfaces and for all address ranges.</p> <p>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.</p> <p>interface <i>interface-name</i>—(Optional) Clear all IGMP group members 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> |
| Required Privilege Level | clear |
| Related Topics | show igmp group show igmp interface |
| List of Sample Output | clear igmp membership on page 69 clear igmp membership interface on page 70 clear igmp membership group on page 70 |
| Output Fields | See show igmp group for an explanation of output fields. |
| 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 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 on all logical systems.</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 Topics | show igmp snooping membership |
| List of Sample Output | clear igmp snooping membership on page 72 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| clear igmp snooping membership | user@host> clear igmp snooping membership |

clear igmp snooping statistics

| | |
|---------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | clear igmp snooping statistics <instance <i>instance-name</i> > <interface <i>interface-name</i> > <learning-domain (all <i>learning-domain-name</i>)> <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced in JUNOS Release 8.5. |
| Description | Clear IP IGMP snooping statistics. |
| Options | <p>none—Clear IGMP snooping statistics for all supported address families on all interfaces on all logical systems.</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 Topics | show igmp snooping statistics |
| List of Sample Output | clear igmp snooping statistics on page 73 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| clear igmp snooping statistics | user@host> clear igmp snooping statistics |

clear igmp statistics

| | |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | clear igmp statistics <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Clear Internet Group Management Protocol (IGMP) statistics. |
| Options | none—Clear IGMP statistics on all interfaces. interface <i>interface-name</i> —(Optional) Clear IGMP statistics 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. |
| Required Privilege Level | clear |
| Related Topics | show igmp statistics |
| List of Sample Output | clear igmp statistics on page 74 |
| Output Fields | Seeshow igmp statistics for an explanation of output fields. |
| 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        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

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      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            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 Release 7.4. |
| Description | Clear Multicast Listener Discovery (MLD) group membership. |
| Options | <p>none—Clear all MLD memberships.</p> <p><i>group group-name</i>—(Optional) Clear MLD membership for the specified group.</p> <p><i>interface interface-name</i>—(Optional) Clear MLD group membership for the specified interface.</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 | view |
| Related Topics | show mld group |
| List of Sample Output | clear mld membership on page 76 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| clear mld membership | user@host> clear mld membership |

clear mld statistics

| | |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | clear mld statistics <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Clear Multicast Listener Discovery (MLD) statistics. |
| Options | <p>none—(Same as logical-system all) Clear MLD statistics for all interfaces on all logical systems.</p> <p>interface <i>interface-name</i>—(Optional) Clear MLD statistics 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 | clear |
| Related Topics | show mld statistics |
| List of Sample Output | clear mld statistics on page 77 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| 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 Release 7.4. |
| 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 Topics | show msdp source-active |
| List of Sample Output | clear msdp cache on page 78 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| clear msdp cache | user@host> clear msdp cache |

clear msdp statistics

| | |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | clear 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 Release 7.4. |
| Description | Clear Multicast Source Discovery Protocol (MSDP) peers statistics. |
| Options | <p>none—Clear MSDP statistics for all peers.</p> <p>instance <i>instance-name</i>—(Optional) Clear statistics for the specified instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>peer <i>peer-address</i>—(Optional) Clear the statistics for the specified peer.</p> |
| Required Privilege Level | clear |
| Related Topics | show msdp statistics |
| List of Sample Output | clear msdp statistics on page 79 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| clear msdp statistics | user@host> clear msdp statistics |

clear multicast scope

| | |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | clear multicast scope <inet inet6> <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced in JUNOS Release 7.6. |
| Description | Clear IP multicast scope statistics. |
| Options | <p>none—(Same as <i>logical-system</i> all) Clear multicast scope statistics on all logical systems.</p> <p>inet—(Optional) Clear multicast scope statistics for IPv4 family addresses.</p> <p>inet6—(Optional) Clear multicast scope statistics for IPv6 family addresses.</p> <p>interface <i>interface-name</i>—(Optional) Clear multicast scope 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 Topics | show multicast scope |
| List of Sample Output | clear multicast scope on page 80 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| clear multicast scope | user@host> clear multicast scope |

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 Release 8.5. |
| Description | Clear IP multicast snooping statistics. |
| Options | <p>none—Clear multicast snooping statistics for all supported address families on all interfaces on all logical systems.</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 Topics | show multicast snooping statistics |
| List of Sample Output | clear multicast snooping statistics on page 81 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| clear multicast snooping statistics | user@host> clear multicast snooping statistics |

clear multicast bandwidth-admission

Syntax clear multicast bandwidth-admission
 <group *group-address*>
 <inet | inet6>
 <instance *instance-name*>
 <interface *interface-name*>
 <source *source-address*>

Release Information Command introduced in JUNOS Release 8.3.

Description Reapply IP multicast bandwidth admissions.

Options none—Reapply multicast bandwidth admissions for all IPv4 forwarding entries in the master routing instance.

group group-address—(Optional) Reapply multicast bandwidth admissions for the specified group.

inet—(Optional) Reapply multicast bandwidth admission settings for IPv4 flows.

inet6—(Optional) Reapply multicast bandwidth admission settings for IPv6 flows.

instance instance-name—(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.

interface interface-name—(Optional) Examines the corresponding outbound interface in the relevant entries and acts as follows:

- 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.

To manually reject previously admitted outbound interfaces, you must specify the interface.

source source-address—(Optional) Use with the **group** option to reapply multicast bandwidth admission settings for the specified (source, group) entry.

Required Privilege Level clear

Related Topics show multicast interface

List of Sample Output clear multicast bandwidth-admission on page 83

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

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

clear multicast sessions

| | |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | clear multicast sessions <logical-system (all <i>logical-system-name</i>)> < <i>regular-expression</i> > |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Clear IP multicast sessions. |
| Options | <p>none—(Same as <i>logical-system</i> all) Clear multicast sessions on all logical systems.</p> <p><i>logical-system</i> (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 Topics | show multicast sessions |
| List of Sample Output | clear multicast sessions on page 84 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| clear multicast sessions | user@host> clear multicast sessions |

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>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Clear IP multicast statistics. |
| Options | <p>none—Clear multicast statistics for all supported address families on all interfaces on all logical systems.</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 Topics | show multicast statistics |
| List of Sample Output | clear multicast statistics on page 85 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| clear multicast statistics | user@host> clear multicast statistics |

clear pgm negative-acknowledgments

| | |
|-------------------------------------------|-------------------------------------------------------------------------------------------|
| Syntax | clear pgm negative-acknowledgments |
| Release Information | Command introduced before JUNOS 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 Topics | show pgm negative-acknowledgments |
| List of Sample Output | clear pgm negative-acknowledgments on page 86 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| clear pgm negative-acknowledgments | user@host> clear pgm negative-acknowledgments |

clear pgm source-path-messages

| | |
|---------------------------------------|---------------------------------------------------------------------------------------|
| Syntax | clear pgm source-path-messages |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Clear Pragmatic General Multicast (PGM) source-path messages. |
| Options | This command has no options. |
| Required Privilege Level | clear |
| Related Topics | show pgm source-path-messages |
| List of Sample Output | clear pgm source-path-messages on page 87 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| 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 Release 7.4. |
| Description | Clear Pragmatic General Multicast (PGM) statistics. |
| Options | This command has no options. |
| Required Privilege Level | clear |
| Related Topics | show pgm statistics |
| List of Sample Output | clear pgm statistics on page 88 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| 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)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| 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 on all logical systems.</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 Topics | show pim join |
| List of Sample Output | clear pim join on page 89 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| clear pim join | user@host> clear pim join |

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>)> |
| Release Information | Command introduced in JUNOS Release 7.6. |
| Description | Clear Protocol Independent Multicast (PIM) register message counters. |
| Options | <p>none—Clear PIM register message counters for all family addresses, instances, and interfaces on all logical systems.</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 Topics | show pim statistics |
| List of Sample Output | clear pim register on page 90 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| clear pim register | user@host> clear pim register |

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>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Clear Protocol Independent Multicast (PIM) statistics. |
| Options | <p>none—Clear PIM statistics for all family addresses, instances, and interfaces on all logical systems.</p> <p>inet inet6—(Optional) Clear PIM statistics for IPv4 or IPv6 family addresses, respectively.</p> <p>instance <i>instance-name</i>—(Optional) Clear statistics for a specific PIM-enabled routing instance.</p> <p>interface <i>interface-name</i>—(Optional) Clear PIM statistics for a specific interface.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Additional Information | The clear pim statistics command cannot be used to clear the PIM statistics on a backup Routing Engine when nonstop active routing is enabled. |
| Required Privilege Level | clear |
| Related Topics | show pim statistics |
| List of Sample Output | clear pim statistics on page 91 |
| Output Fields | See show pim statistics for an explanation of output fields. |
| 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
...

```

show dvmrp interfaces

| | |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show dvmrp interfaces <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS 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 on all logical systems.</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 94 |
| Output Fields | Table 21 on page 93 describes the output fields for the show dvmrp interfaces command. Output fields are listed in the approximate order in which they appear. |

Table 21: 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. |

```
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 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 on all logical systems.</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 96 |
| Output Fields | Table 22 on page 95 describes the output fields for the show dvmrp neighbors command. Output fields are listed in the approximate order in which they appear. |

Table 22: 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>major</i> <i>minor</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 prunes and grafts. ■ 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. |

```
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 Release 7.4. |
| Description | Display information about Distance Vector Multicast Routing Protocol (DVMRP) prefixes. |
| Options | <p>none—Display standard information about all DVMRP prefixes on all logical systems.</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 | <p>show dvmrp prefix on page 98</p> <p>show dvmrp prefix brief on page 98</p> <p>show dvmrp prefix detail on page 98</p> |
| Output Fields | Table 23 on page 97 describes the output fields for the show dvmrp prefix command. Output fields are listed in the approximate order in which they appear. |

Table 23: 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 prunes 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 prunes, in seconds. | detail |

```

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 Release 7.4. |
| Description | Display information about active Distance Vector Multicast Routing Protocol (DVMRP) prunes. |
| Options | <p>none—Display received and transmitted DVMRP prune information on all logical systems.</p> <p>all—(Optional) Display information about all received and transmitted prunes.</p> <p>rx—(Optional) Display information about received prunes.</p> <p>tx—(Optional) Display information about transmitted prunes.</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 99 |
| Output Fields | Table 24 on page 99 describes the output fields for the <code>show dvmrp prunes</code> command. Output fields are listed in the approximate order in which they appear. |

Table 24: 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. |

```

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>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display Internet Group Management Protocol (IGMP) group membership information. |
| Options | <p>none—Display standard information about membership for all IGMP groups on all logical systems.</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 Topics | clear igmp membership |
| List of Sample Output | <p>show igmp group (Include Mode) on page 102</p> <p>show igmp group (Exclude Mode) on page 102</p> <p>show igmp group brief on page 103</p> <p>show igmp group detail on page 103</p> |
| Output Fields | Table 25 on page 101 describes the output fields for the show igmp group command. Output fields are listed in the approximate order in which they appear. |

Table 25: 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 router 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 |
| 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 |

Table 25: show igmp group Output Fields (continued)

| Field Name | Field Description | Level of Output |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| 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 |

```

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>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display information about Internet Group Management Protocol (IGMP)-enabled interfaces. |
| Options | <p>none—Display standard information about all IGMP-enabled interfaces on all logical systems.</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 Topics | clear igmp membership |
| List of Sample Output | <p>show igmp interface on page 105</p> <p>show igmp interface brief on page 106</p> <p>show igmp interface detail on page 106</p> |
| Output Fields | Table 26 on page 104 describes the output fields for the show igmp interface command. Output fields are listed in the approximate order in which they appear. |

Table 26: show igmp interface Output Fields

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

Table 26: show igmp interface Output Fields (continued)

| Field Name | Field Description | Level of Output |
|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| 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 mode | 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. | 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 (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 |

```

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
                        Interface: so-1/0/0.0
                        Querier: 10.111.10.1
                        State:      Up Timeout:  None Version:  2 Groups:    2
                        Interface: so-1/0/1.0

```

```

Querier: 10.111.20.1
State:      Up Timeout:   None Version:  2 Groups:    4
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 105.

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 105.

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 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 Topics | <p>show igmp snooping membership</p> <p>show igmp snooping statistics</p> |
| List of Sample Output | show igmp snooping interface on page 108 |
| Output Fields | Table 27 on page 107 lists the output fields for the show igmp snooping interface command. Output fields are listed in the approximate order in which they appear. |

Table 27: 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 |
| 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 |

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

| Field Name | Field Description | Level of Output |
|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| immediate-leave | State of immediate leave: On or Off | All levels |
| router-interface | Router interfaces that are part of this learning domain. | 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 |
| 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 |

```

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

```

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 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 Topics | <p>show igmp snooping interface</p> <p>show igmp snooping statistics</p> <p>clear igmp snooping membership</p> |
| List of Sample Output | <p>show igmp snooping membership on page 110</p> <p>show igmp snooping membership (Exclude Mode) on page 110</p> <p>show igmp snooping membership interface ge-0/1/2.200 on page 111</p> <p>show igmp snooping membership vlan-id 100 on page 111</p> |
| Output Fields | Table 28 on page 109 lists the output fields for the show igmp snooping membership command. Output fields are listed in the approximate order in which they appear. |

Table 28: 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 |
| Interface | Interface on which this router is a proxy. | detail |
| Group | Multicast group address in the membership database. | All levels |

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

| Field Name | Field Description | Level of Output |
|------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| 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 |

show igmp snooping membership user@host> **show igmp snooping membership**
Instance: bridge-domain bar

```

Learning-Domain: default
Interface: ge-0/1/2.200
  Group: 225.1.1.1
    Source: 0.0.0.0
    Timeout: 398 Type: Static
  Group: 232.1.1.1
    Source: 192.168.1.1
    Timeout: 0 Type: Static
Interface: ge-0/1/3.200
  Group: 225.1.1.2
    Source: 0.0.0.0
    Timeout: 393 Type: Dynamic

```

show igmp snooping membership (Exclude Mode) user@host> **show igmp snooping membership**
Instance: vpls

```

Learning-Domain: default
Interface: ge-2/0/4.1000
  Group: 225.10.10.10
    Group mode: Exclude
    Source: 0.0.0.0
    Last reported by: 100.3.3.2
    Group timeout: 200 Type: Dynamic
  Group: 225.11.11.11
    Group mode: Exclude
    Source: 0.0.0.0
    Last reported by: 100.3.4.2
    Group timeout: 158 Type: Dynamic

```

```
show igmp snooping      user@host> show igmp snooping membership interface ge-0/1/2.200  
membership interface   Instance: bridge-domain bar  
ge-0/1/2.200           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      user@host> show igmp snooping membership vlan-id 100  
membership vlan-id 100 Instance: bridge-domain bar  
                          IGMP-snooping not running in vlan 100
```

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 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 Topics | show igmp snooping interface show igmp snooping membership clear igmp snooping statistics |
| List of Sample Output | show igmp snooping statistics on page 113 |
| Output Fields | Table 29 on page 112 lists the output fields for the show igmp snooping statistics command. Output fields are listed in the approximate order in which they appear. |

Table 29: 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 29: 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 |

```

show igmp snooping      user@host> show igmp snooping statistics
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 <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display Internet Group Management Protocol (IGMP) statistics. |
| Options | <p>none—Display IGMP statistics for all interfaces on all logical systems.</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 Topics | clear igmp statistics |
| List of Sample Output | <p>show igmp statistics on page 116</p> <p>show igmp statistics interface on page 117</p> |
| Output Fields | Table 30 on page 115 describes the output fields for the show igmp statistics command. Output fields are listed in the approximate order in which they appear. |

Table 30: 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 30: 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. |

```

show igmp statistics  user@host> show igmp statistics
IGMP packet statistics for all interfaces
IGMP Message type      Received      Sent  Rx errors
Membership Query        8883         459      0
V1 Membership Report    0            0      0
DVMRP                   0            0      0
PIM V1                  0            0      0
Cisco Trace             0            0      0
V2 Membership Report    0            0      0

```

| | | | |
|-------------------------------------|------|---|---|
| Group Leave | 0 | 0 | 0 |
| Mtrace Response | 0 | 0 | 0 |
| Mtrace Request | 0 | 0 | 0 |
| Domain Wide Report | 0 | 0 | 0 |
| V3 Membership Report | 0 | 0 | 0 |
| Other Unknown types | | | 0 |
| IGMP v3 unsupported type | | | 0 |
| IGMP v3 source required for SSM | | | 0 |
| IGMP v3 mode not applicable for SSM | | | 0 |
| IGMP Global Statistics | | | |
| Bad Length | 0 | | |
| Bad Checksum | 0 | | |
| Bad Receive If | 0 | | |
| Rx non-local | 1227 | | |
| Timed out | 0 | | |
| Rejected Report | 0 | | |
| Total Interfaces | 2 | | |

```

show igmp statistics user@host> show igmp statistics interface fe-1/0/1.0
interface           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 Release 7.4. |
| Description | Display information about Multicast Listener Discovery (MLD) group membership. |
| Options | <p>none—Display standard information about all MLD groups on all logical systems.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p><i>group-name</i>—(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 Topics | clear mld membership |
| List of Sample Output | <p>show mld group (Include Mode) on page 119</p> <p>show mld group (Exclude Mode) on page 119</p> <p>show mld group brief on page 120</p> <p>show mld group detail (Include Mode) on page 120</p> <p>show mld group detail (Exclude Mode) on page 120</p> |
| Output Fields | Table 31 on page 118 describes the output fields for the show mld group command. Output fields are listed in the approximate order in which they appear. |

Table 31: show mld group Output Fields

| Field Name | Field Description | Level of Output |
|------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Interface | Name of the interface that received the MLD membership report; <i>local</i> 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: <i>Include</i> or <i>Exclude</i> . | All levels |
| Last reported by | Address of the host that last reported membership in this group. | All levels |
| 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 |

Table 31: show mld group Output Fields (continued)

| Field Name | Field Description | Level of Output |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| 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 |

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 119 `show mld group (Exclude Mode)` on page 119.

**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 Release 7.4. |
| Description | Display information about Multicast Listener Discovery (MLD)-enabled interfaces. |
| Options | <p>none—Display standard information about all MLD-enabled interfaces on all logical systems.</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 Topics | clear mld membership |
| List of Sample Output | <p>show mld interface on page 124</p> <p>show mld interface brief on page 124</p> <p>show mld interface detail on page 124</p> |
| Output Fields | Table 32 on page 122 describes the output fields for the show mld interface command. Output fields are listed in the approximate order in which they appear. |

Table 32: 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 |
| 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 , 2 , or 3. | All levels |
| Groups | Number of groups on the interface. | All levels |

Table 32: show mld interface Output Fields (continued)

| Field Name | Field Description | Level of Output |
|-----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Passive mode | State of the passive mode option: <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. | All levels |
| OIF map | Name of the OIF map associated to the interface. | All levels |
| Passive mode | State of the passive mode option: <ul style="list-style-type: none"> ■ On—Indicates that the router can run MLD on the interface but not send or receive control traffic such as MLD reports, queries, and leaves. ■ Off—Indicates that the router can run MLD on the interface and send or receive control traffic such as MLD reports, queries, and leaves. | 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 |
| 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 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 |
| 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 |

show mld interface

```

user@host> show mld interface
Interface: fe-0/0/0
  Querier: None
  State: Up      Timeout:      0    Version: 1    Groups:      0
Interface: at-0/3/1.0
  Querier: 8038::c0a8:c345
  State: Up      Timeout:      None  Version: 1    Groups:      0
Interface: fe-1/0/1.0
  Querier: ::192.168.195.73
  State: Up      Timeout:      None  Version: 1    Groups:      3
  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 124.

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 124.

show mld statistics

| | |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show mld statistics <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display information about Multicast Listener Discovery (MLD) statistics. |
| Options | none—Display MLD statistics for all interfaces on all logical systems. 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 Topics | clear mld statistics |
| List of Sample Output | show mld statistics on page 126 show mld statistics interface on page 127 |
| Output Fields | Table 33 on page 125 describes the output fields for the show mld statistics command. Output fields are listed in the approximate order in which they appear. |

Table 33: 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. |
| MLD Message type | Summary of MLD statistics. <ul style="list-style-type: none"> ■ Listener Query (v1/v2)—Number of membership queries sent and received. ■ Listener Report (v1)—Number of version 1 membership reports sent and received. ■ Listener Done (v1/v2)—Number of Listener Done messages sent and received. ■ Listener Report (v2)—Number of version 2 membership reports sent and received. ■ Other Unknown types—Number of unknown message types received. ■ MLD v2 source required for SSM—Number of MLD version 2 messages received that contained no source. ■ MLD v2 mode not applicable for SSM—Number of MLD version 2 messages received that did not contain a mode applicable for source-specific multicast (SSM). |

Table 33: show mld statistics Output Fields (continued)

| Field Name | Field Description |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MLD Global Statistics | Summary of MLD statistics for all interfaces. <ul style="list-style-type: none"> ■ Bad Length—Number of messages received with length errors so severe that further classification could not occur. ■ Bad Checksum—Number of messages received with an invalid IP checksum. No further classification was performed. ■ Bad Receive If—Number of messages received on an interface not enabled for MLD. ■ Rx non-local—Number of messages received from nonlocal senders. ■ Timed out—Number of groups that timed out as a result of not receiving an explicit leave message. ■ Rejected Report—Number of reports dropped because of the MLD group policy. ■ Total Interfaces—Number of interfaces configured to support IGMP. |

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
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 user@host> show mld statistics interface fe-1/0/1.0
interface 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
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 Release 7.4. |
| Description | Display Multicast Source Discovery Protocol (MSDP) information. |
| Options | <p>none—Display standard MSDP information for all routing instances on all logical systems.</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 Topics | <p>show msdp source</p> <p>show msdp source-active</p> <p>show msdp statistics</p> |
| List of Sample Output | <p>show msdp on page 129</p> <p>show msdp brief on page 129</p> <p>show msdp detail on page 129</p> |
| Output Fields | Table 34 on page 128 describes the output fields for the show msdp command. Output fields are listed in the approximate order in which they appear. |

Table 34: 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 |
| Peer-Group | Peer group name. | All levels |

Table 34: show msdp Output Fields (continued)

| Field Name | Field Description | Level of Output |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| 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 |

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 129.

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 | show msdp source <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> <source-address> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display multicast sources learned from Multicast Source Discovery Protocol (MSDP). |
| Options | <p>none—Display standard MSDP source information for all routing instances on all logical systems.</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>source-address—(Optional) IP address and optional prefix length. Display information for the specified source address only.</p> |
| Required Privilege Level | view |
| Related Topics | show msdp show msdp source-active show msdp statistics |
| List of Sample Output | show msdp source on page 131 |

Output Fields Table 35 on page 131 describes the output fields for the `show msdp source` command. Output fields are listed in the approximate order in which they appear.

Table 35: 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. |

```

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 | Command introduced before JUNOS Release 7.4. |
| 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 on all logical systems.</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 Topics | <pre>show msdp show msdp source show msdp statistics</pre> |
| List of Sample Output | <pre>show msdp source-active on page 133 show msdp source-active brief on page 133 show msdp source-active detail on page 133</pre> |
| Output Fields | Table 36 on page 133 describes the output fields for the <code>show msdp source-active</code> command. Output fields are listed in the approximate order in which they appear. |

Table 36: show msdp source-active Output Fields

| Field Name | Field Description |
|----------------|-------------------------------------------------------------------|
| Group address | Multicast address of the group. |
| Source address | IP address of the source. |
| Peer address | IP address of the peer. |
| Originator | Address of the rendezvous point (RP) that originated the message. |
| Flags | Flags: Accept, Reject, or Filtered. |

```

show msdp      user@host> show msdp source-active
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 133.

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 133.

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 Release 7.4. |
| Description | Display statistics about Multicast Source Discovery Protocol (MSDP) peers. |
| Options | <p>none—Display statistics about all MSDP peers for all routing instances on all logical systems.</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 Topics | clear msdp statistics |
| List of Sample Output | show msdp statistics on page 135 |
| Output Fields | Table 37 on page 134 describes the output fields for the show msdp statistics command. Output fields are listed in the approximate order in which they appear. |

Table 37: 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. |
| Peer | Address of peer. |
| 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. |

Table 37: show msdp statistics Output Fields *(continued)*

| Field Name | Field Description |
|-------------------------------|-------------------------------------------------------------------------|
| 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. |
| Keepalive messages sent | Number of keepalive messages sent. |
| Keepalive messages received | Number of keepalive messages received. |
| Unknown messages received | Number of unknown messages received. |
| Error messages received | Number of error messages received. |

show msdp statistics

```

user@host> show msdp statistics
Global active source exceeded: 0

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
Keepalive messages sent: 17
Keepalive messages received: 19
Unknown messages received: 0
Error messages received: 0

```

show multicast flow-map

| | |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show multicast flow-map |
| Release Information | Command introduced in JUNOS Release 8.2. |
| Description | Display configuration information about IP multicast flow maps. |
| Options | This command has no options. |
| Required Privilege Level | view |
| List of Sample Output | show multicast flow-map on page 136 show multicast flow-map detail on page 136 |
| Output Fields | Table 38 on page 136 describes the output fields for the show multicast flow-map command. Output fields are listed in the approximate order in which they appear. |

Table 38: 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 to 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 |

```

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

```

```

show multicast flow-map  user@host> show multicast flow-map detail
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 interface

| | |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show multicast interface <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced in JUNOS Release 8.3. |
| Description | Display bandwidth information about IP multicast interfaces. |
| 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 multicast interface on page 139 |
| Output Fields | Table 39 on page 138 describes the output fields for the show multicast interface command. Output fields are listed in the approximate order in which they appear. |

Table 39: show multicast interface Output Fields

| Field Name | Field Description |
|---------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Interface | Name of the multicast interface. |
| Maximum bandwidth (bps) | Maximum bandwidth setting, in bits per second, for this interface. |
| Remaining bandwidth (bps) | Amount of bandwidth, in bits per second, remaining on the interface. |
| Mapped bandwidth deduction (bps) | Amount of bandwidth, in bits per second, used by any flows that are mapped to the interface. NOTE: Adding the mapped bandwidth deduction value to the local bandwidth deduction value results in the total deduction value for the interface. |
| Local bandwidth deduction (bps) | Amount of bandwidth, in bits per second, used by any mapped flows that are traversing the interface. NOTE: Adding the mapped bandwidth deduction value to the local bandwidth deduction value results in the total deduction value for the interface. |
| Reverse OIF mapping | State of the reverse OIF mapping feature (on or off). |
| Reverse OIF mapping no QoS adjustment | State of the no QoS adjustment feature (on or off) for interfaces that are using reverse OIF mapping. |
| Leave timer | Amount of time a mapped interface remains active after the last mapping ends. |

Table 39: show multicast interface Output Fields *(continued)*

| Field Name | Field Description |
|-------------------|-------------------------------------------------------------------------------------------------------|
| No QoS adjustment | State (on) of the no QoS adjustment feature when this feature is enabled. |
| | NOTE: This field does not appear in the output when the no QoS adjustment feature is disabled. |

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 mrinfo

| | |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show multicast mrinfo <host> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display configuration information about IP multicast networks, including neighboring multicast router addresses. |
| Options | <p>none—Display configuration information about all multicast networks.</p> <p>host—(Optional) Display configuration information about a particular host. Replace <i>host</i> with a hostname or IP address.</p> |
| Required Privilege Level | view |
| List of Sample Output | show multicast mrinfo on page 141 |
| Output Fields | Table 40 on page 140 describes the output fields for the show multicast mrinfo command. Output fields are listed in the approximate order in which they appear. |

Table 40: show multicast mrinfo Output Fields

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

```
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 | show multicast next-hops <brief detail> <identifier <i>identifier-number</i> > <inet inet6> <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display the entries in the IP multicast next-hop table. |
| Options | <p>none—Display standard information about all entries in the multicast next-hop table for all supported address families on all logical systems.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>identifier <i>identifier-number</i>—(Optional) Show a particular next hop by ID number. The range of values is 1 through 65,535.</p> <p>inet inet6—(Optional) Display entries for IPv4 or IPv6 family addresses, respectively.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | <p>show multicast next-hops on page 143</p> <p>show multicast next-hops brief on page 143</p> <p>show multicast next-hops detail on page 143</p> |
| Output Fields | Table 41 on page 142 describes the output fields for the show multicast next-hops command. Output fields are listed in the approximate order in which they appear. |

Table 41: show multicast next-hops Output Fields

| Field Name | Field Description |
|----------------------|---------------------------------------------------------------------------------------------------------|
| ID | Next-hop identifier of the prefix. The identifier is returned by the router's Packet Forwarding Engine. |
| Refcnt | 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. |

```

show multicast      user@host> show multicast next-hops
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

                     Family: INET6

```

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 143.

show multicast next-hops detail The output for the `show multicast next-hops detail` command is identical to that for the `show multicast next-hops` command. For sample output, see `show multicast next-hops` on page 143.

show multicast pim-to-igmp-proxy

| | |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show multicast pim-to-igmp-proxy <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced in JUNOS Release 9.6. |
| Description | Display configuration information about PIM-to-IGMP message translation, also known as PIM-to-IGMP proxy. |
| 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 multicast pim-to-igmp-proxy on page 144 |
| Output Fields | Table 42 on page 144 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 42: show multicast pim-to-igmp-proxy Output Fields

| Field Name | Field Description |
|-----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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. |

show multicast pim-to-igmp-proxy

```

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

```

show multicast pim-to-ml-d-proxy

| | |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show multicast pim-to-ml-d-proxy <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced in JUNOS Release 9.6. |
| Description | Display configuration information about PIM-to-MLD message translation, also known as PIM-to-MLD proxy. |
| 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 multicast pim-to-ml-d-proxy on page 145 |
| Output Fields | Table 43 on page 145 describes the output fields for the show multicast pim-to-ml-d-proxy command. Output fields are listed in the order in which they appear. |

Table 43: show multicast pim-to-ml-d-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: <i>enabled</i> or <i>disabled</i> . |
| <i>interface-name</i> | Name of upstream interface (no more than two allowed) on which PIM-to-MLD message translation is configured. |

show multicast pim-to-ml-d-proxy user@host> **show multicast pim-to-ml-d-proxy**
Proxy state: enabled
ge-0/5/0.1
ge-0/5/0.2

show multicast route

| | |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | <pre>show multicast route <brief detail extensive> <active all inactive> <group group> <inet inet6> <instance instance name> <logical-system (all logical-system-name)> <regular-expression> <source-prefix source-prefix></pre> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display the entries in the IP multicast forwarding table. You can display similar information with the show route table inet.1 command. |
| Options | <p>none—Display standard information about all entries in the multicast forwarding table for all routing instances on all logical systems.</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 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-style regular expression.</p> <p>source-prefix source-prefix—(Optional) Display the cache entries for a particular source prefix.</p> |
| Required Privilege Level | view |
| List of Sample Output | <p>show multicast route on page 147</p> <p>show multicast route brief on page 148</p> <p>show multicast route detail on page 148</p> <p>show multicast route extensive on page 148</p> |
| Output Fields | Table 44 on page 147 describes the output fields for the show multicast route command. Output fields are listed in the approximate order in which they appear. |

Table 44: show multicast route Output Fields

| Field Name | Field Description | Level of Output |
|----------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| Address 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 |
| Upstream interface | Name of the interface on which the packet with this source prefix is expected to arrive. | All levels |
| Downstream interface list | A list of interface names to which the packet with this source prefix is forwarded to. | All levels |
| 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. | 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 <code>show multicast nexthops</code> command. | detail extensive |
| Upstream protocol | Protocol running on the interface on which the packet with this source prefix is expected to arrive. | 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 <code>never</code> indicates a permanent forwarding entry. | extensive |
| Wrong incoming interface notifications | Number of times that the upstream interface was not available. | extensive |

```

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 brief The output for the show multicast route brief command is identical to that for the show multicast route command. For sample output, see show multicast route on page 147.

```
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: 0 kbps, 0 pps, 38 packets
  Next-hop ID: 262143
  Upstream protocol: PIM

Family: INET6
```

```
user@host> show multicast route extensive
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, 46454 packets
  Next-hop ID: 262142
  Upstream protocol: PIM
  Route state: Active
  Forwarding state: Forwarding
  Cache lifetime/timeout: 360 seconds
  Wrong incoming interface notifications: 0
```

```

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
  Route state: Active
  Forwarding state: Forwarding
  Cache lifetime/timeout: 348 seconds
  Wrong incoming interface notifications: 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
  Session description: Administratively Scoped
  Statistics: 0 kbps, 0 pps, 40 packets
  Next-hop ID: 262143
  Upstream protocol: PIM
  Route state: Active
  Forwarding state: Forwarding
  Cache lifetime/timeout: 360 seconds
  Wrong incoming interface notifications: 1

Family: INET6

```

show multicast rpf

| | |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show multicast rpf <inet inet6> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> <prefix> <summary> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display information about multicast reverse-path-forwarding (RPF) calculations. |
| Options | <p>none—Display RPF calculation information for all supported address families on all logical systems.</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><i>prefix</i>—(Optional) Display the RPF calculation information for the specified prefix.</p> <p>summary—(Optional) Display summary of all multicast RPF information.</p> |
| Required Privilege Level | view |
| List of Sample Output | <p>show multicast rpf on page 151</p> <p>show multicast rpf inet6 on page 152</p> <p>show multicast rpf prefix on page 153</p> <p>show multicast rpf summary on page 153</p> |

Output Fields Table 45 on page 151 describes the output fields for the **show multicast rpf** command. Output fields are listed in the approximate order in which they appear.

Table 45: 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. |
| Neighbor | Upstream RPF neighbor. |

```

show multicast rpf user@host> show multicast rpf

Multicast RPF table: inet.0, 12 entries

0.0.0.0/0
  Protocol: Static

10.255.14.132/32
  Protocol: Direct
  Interface: lo0.0

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

127.0.0.1/32
Inactive172.16.0.0/12
Protocol: Static
Interface: fxp0.0
Neighbor: 192.168.14.254

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

192.168.14.0/24
Protocol: Direct
Interface: fxp0.0

192.168.14.132/32
Protocol: Local

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

```

```

192.168.195.22/32
Protocol: Local

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

```

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

Multicast RPF table: inet6.0, 12 entries

```

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

::10.255.245.91/128
  Protocol: IS-IS
  Interface: so-1/1/1.0
  Neighbor: fe80::2a0:a5ff:fe28:2e8c

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

::192.168.195.22/128
  Protocol: Local

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

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

::192.168.195.77/128
  Protocol: Local

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

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

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

ff02::2/128
  Protocol: PIM

ff02::d/128
  Protocol: PIM

```

```
show multicast rpf prefix user@host> show multicast rpf ff02::/16  
  
Multicast RPF table: inet6.0, 13 entries  
  
ff02::2/128  
    Protocol: PIM  
  
ff02::d/128  
    Protocol: PIM  
  
...
```

```
show multicast rpf summary user@host> show multicast rpf summary  
  
Multicast RPF table: inet.0, 16 entries  
Multicast RPF table: inet6.0, 12 entries
```

show multicast scope

| | |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show multicast scope <inet inet6> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| 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 on all logical systems.</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 | <p>show multicast scope on page 155</p> <p>show multicast scope inet on page 155</p> <p>show multicast scope inet6 on page 155</p> |
| Output Fields | Table 46 on page 154 describes the output fields for the <code>show multicast scope</code> command. Output fields are listed in the approximate order in which they appear. |

Table 46: 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. |


```

show multicast scope user@host> show multicast scope

```

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

```

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

```

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

```

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

```

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

show multicast sessions

| | |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show multicast sessions <brief detail extensive> <logical-system (all <i>logical-system-name</i>)> < <i>regular-expression</i> > |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display information about announced IP multicast sessions. |
| Options | <p>none—Display standard information about all multicast sessions for all routing instances on all logical systems.</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 | <p>show multicast sessions on page 157</p> <p>show multicast sessions regular-expression detail on page 157</p> |
| Output Fields | Table 47 on page 156 describes the output fields for the show multicast sessions command. Output fields are listed in the approximate order in which they appear. |

Table 47: show multicast sessions Output Fields

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

```

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
regular-expression detail    SDP Version: 0  Originated by: -@128.223.83.33
                                Session: NASA TV (MPEG-1)
                                Description: NASA television in MPEG-1 format, provided by Private University.
                                Please contact the UO if you have problems with this feed.
                                Email: Your Name Here <multicast@lists.private.edu>
                                Phone: Your Name Here <888/555-1212>
                                Bandwidth: AS:1000
                                Start time: permanent
                                Stop time: none
                                Attribute: type:broadcast
                                Attribute: tool:IP/TV Content Manager 3.4.14
                                Attribute: live:capture:1
                                Attribute: x-iptv-capture:mp1s
                                Media: video 54302 RTP/AVP 32 31 96 97
                                Connection Data: 224.2.231.45 ttl 127
                                Attribute: quality:8
                                Attribute: framerate:30
                                Attribute: rtpmap:96 WBIH/90000
                                Attribute: rtpmap:97 MP4V-ES/90000
                                Attribute: x-iptv-svr:video 128.223.91.191 live
                                Attribute: 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 route

| | |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | 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> > |
| Release Information | Command introduced in JUNOS 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 | show multicast snooping route bridge-domain on page 159 |
| Output Fields | Table 48 on page 158 describes the output fields for the show multicast snooping route command. Output fields are listed in the approximate order in which they appear. |

Table 48: show multicast snooping route Output Fields

| Field Name | Field Description | Level of Output |
|------------|------------------------------------------------------------|-----------------|
| Family | IPv4 address family (INET) or IPv6 address family (INET6). | All levels |

Table 48: show multicast snooping route Output Fields *(continued)*

| Field Name | Field Description | Level of Output |
|------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| 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 |

```

show multicast snooping user@host> show multicast snooping route bridge-domain br-dom-1 extensive
route bridge-domain      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 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 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 on all logical systems.</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 Topics | clear multicast snooping statistics |
| List of Sample Output | show multicast snooping statistics on page 162 |
| Output Fields | Table 49 on page 160 describes the output fields for the show multicast snooping statistics command. Output fields are listed in the approximate order in which they appear. |

Table 49: 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. |
| 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. |

Table 49: show multicast snooping statistics Output Fields *(continued)*

| Field Name | Field Description |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| In Kbytes | Total accumulated incoming packets (in KB) since the last time the <code>clear multicast snooping statistics</code> command was issued. |
| Out Kbytes | Total accumulated outgoing packets (in KB) since the last time the <code>clear multicast snooping statistics</code> 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 <code>clear multicast snooping statistics</code> command was issued. |
| Out packets | Total number of outgoing packets since the last time the <code>clear multicast snooping statistics</code> command was issued. |

```

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 Release 7.4. |
| Description | Display IP multicast statistics. |
| Options | <p>none—Display multicast statistics for all supported address families for all routing instances on all logical systems.</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 Topics | clear multicast statistics |
| List of Sample Output | show multicast statistics on page 165 |
| Output Fields | Table 50 on page 163 describes the output fields for the show multicast statistics command. Output fields are listed in the approximate order in which they appear. |

Table 50: 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. |
| Resolve No Route | Number of resolve requests that were ignored because there was no route to the source. |

Table 50: show multicast statistics Output Fields (continued)

| Field Name | Field Description |
|------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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. |

```

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>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| 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 on all logical systems.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>inet inet6—(Optional) Display usage information for IPv4 or IPv6 family addresses, respectively.</p> <p>instance <i>instance-name</i>—(Optional) Display information about the most active DVMRP or PIM groups for a specific multicast instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | <p>show multicast usage on page 167</p> <p>show multicast usage brief on page 167</p> <p>show multicast usage instance on page 167</p> <p>show multicast usage detail on page 167</p> |
| Output Fields | Table 51 on page 166 describes the output fields for the show multicast usage command. Output fields are listed in the approximate order in which they appear. |

Table 51: show multicast usage Output Fields

| Field Name | Field Description |
|------------|---------------------------------------------------------------------------------------------------|
| Instance | Name of the routing instance. (Displayed when multicast is configured within a routing instance.) |
| Group | Group address. |
| Sources | Number of sources. |
| Packets | Number of packets that have been forwarded to this prefix. |
| Bytes | Amount of memory used. |
| Prefix | IP address. |

Table 51: show multicast usage Output Fields (continued)

| Field Name | Field Description |
|------------|-----------------------------|
| /len | Prefix length. |
| Groups | Number of multicast groups. |

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 167.

**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 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 169 |
| Output Fields | Table 52 on page 168 describes the output fields for the <code>show pgm negative-acknowledgments</code> command. Output fields are listed in the approximate order in which they appear. |

Table 52: 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. |

```
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 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 170 |
| Output Fields | Table 53 on page 170 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 53: 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. |

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 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 173 |
| Output Fields | Table 54 on page 171 describes the output fields for the show pgm statistics command. Output fields are listed in the approximate order in which they appear. |

Table 54: 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 54: 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). |

show pgm statisticsuser@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 TSU 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>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| 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 on all logical systems.</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 | <p>show pim bootstrap on page 174</p> <p>show pim bootstrap instance on page 175</p> |
| Output Fields | Table 55 on page 174 describes the output fields for the show pim bootstrap command. Output fields are listed in the approximate order in which they appear. |

Table 55: show pim bootstrap Output Fields

| Field Name | Field Description |
|---------------|----------------------------------------------------------------------------------------------|
| Instance | Name of the routing instance. |
| BSR | Bootstrap router. |
| Pri | Priority of the router to be elected to be the bootstrap router. |
| Local address | Local router's address. |
| Pri | Local router's address priority to be elected as the bootstrap router. |
| State | Local router's election state: Candidate , Elected , or Ineligible . |
| Timeout | How long until the local router declares the bootstrap router to be unreachable, in seconds. |

```

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> > <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| 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 all routing instances on all logical systems.</p> <p>inet inet6—(Optional) Display interface information for IPv4 or IPv6 family addresses, respectively.</p> <p>instance <i>instance-name</i>—(Optional) Display information about interfaces 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 | <p>show pim interfaces on page 177</p> <p>show pim interfaces inet on page 178</p> <p>show pim interfaces inet6 on page 178</p> |
| Output Fields | Table 56 on page 176 describes the output fields for the show pim interfaces command. Output fields are listed in the approximate order in which they appear. |

Table 56: 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 56: 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"> ■ Sparse—In sparse mode, routers must join and leave multicast groups explicitly. Upstream routers do not forward multicast traffic to this router unless this router has sent an explicit request (using a join message) to receive multicast traffic. ■ Dense—Unlike sparse mode, where data is forwarded only to routers 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). ■ 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. |
| 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"> ■ DR—Designated router. ■ NotDR—Not the designated router. ■ P2P—Point to point. |
| Count | Number of neighbors that have been seen on the interface. |
| DR address | Address of the designated router. |

show pim interfaces user@host> **show pim interfaces**
 Instance: PIM.master

| Name | Stat | Mode | IP | V | State | Count | DR address |
|------------|------|--------|----|---|-------|-------|--------------------------|
| fxp0.0 | Up | Sparse | 4 | 2 | DR | 1 | 192.68.12.51 |
| fxp1.0 | Up | Sparse | 4 | 2 | NotDR | 1 | 192.68.12.98 |
| fxp2.0 | Up | Sparse | 4 | 2 | DR | 0 | 10.1.1.1 |
| gre.0 | Up | Sparse | 4 | 2 | P2P | 0 | |
| lo0.0 | Up | Sparse | 4 | 2 | DR | 0 | 127.0.0.1 |
| pimd.32768 | Up | Sparse | 4 | 2 | P2P | 0 | |
| sr0.0 | Down | Sparse | 4 | 2 | P2P | 0 | |
| fxp0.0 | Up | Sparse | 6 | 2 | DR | 0 | fec0::192.68.12.51 |
| fxp1.0 | Up | Sparse | 6 | 2 | DR | 0 | fec0::192.68.12.97 |
| fxp2.0 | Up | Sparse | 6 | 2 | DR | 0 | fe80::2a0:c9ff:fe69:eb5f |

show pim interfaces inet user@host> **show pim interfaces inet**
 Instance: PIM.master

| Name | Stat | Mode | IP | V | State | Count | DR | address |
|------------|------|--------|----|---|-------|-------|--------------|---------|
| fxp0.0 | Up | Sparse | 4 | 2 | DR | 1 | 192.68.12.51 | |
| fxp1.0 | Up | Sparse | 4 | 2 | NotDR | 1 | 192.68.12.98 | |
| fxp2.0 | Up | Sparse | 4 | 2 | DR | 0 | 10.1.1.1 | |
| gre.0 | Up | Sparse | 4 | 2 | P2P | 0 | | |
| lo0.0 | Up | Sparse | 4 | 2 | DR | 0 | 127.0.0.1 | |
| pimd.32768 | Up | Sparse | 4 | 2 | P2P | 0 | | |
| sr0.0 | Down | Sparse | 4 | 2 | P2P | 0 | | |

show pim interfaces inet6 user@host> **show pim interfaces inet6**
 Instance: PIM.master

| Name | Stat | Mode | IP | V | State | Count | DR | address |
|--------|------|--------|----|---|-------|-------|--------------------------|---------|
| fxp0.0 | Up | Sparse | 6 | 2 | DR | 0 | fec0::192.68.12.51 | |
| fxp1.0 | Up | Sparse | 6 | 2 | DR | 0 | fec0::192.68.12.97 | |
| fxp2.0 | Up | Sparse | 6 | 2 | DR | 0 | fe80::2a0:c9ff:fe69:eb5f | |

show pim join

| | |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show pim join <brief detail extensive> <inet inet6> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> <range> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display information about Protocol Independent Multicast (PIM) groups. |
| Options | <p>none—Display the standard information about PIM groups for all supported family addresses for all routing instances on all logical systems.</p> <p>brief detail extensive—(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 Topics | clear pim join |
| List of Sample Output | <p>show pim join on page 181</p> <p>show pim join instance on page 181</p> <p>show pim join detail on page 181</p> <p>show pim join extensive on page 182</p> <p>show pim join instance extensive on page 183</p> |
| Output Fields | Table 57 on page 179 describes the output fields for the show pim join command. Output fields are listed in the approximate order in which they appear. |

Table 57: show pim join Output Fields

| Field Name | Field Description |
|------------|----------------------------------------------------------|
| Instance | Name of the routing instance. |
| Family | Name of the address family: inet (IPv4) or inet6 (IPv6). |
| R | Rendezvous Point Tree |

Table 57: show pim join Output Fields (continued)

| Field Name | Field Description |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| S | Sparse |
| W | Wildcard |
| Group | Group address. |
| Source | Multicast source: <ul style="list-style-type: none"> ■ * (wildcard value) ■ <i>ipv4-address</i> ■ <i>ipv6-address</i> |
| RP | Rendezvous point for the PIM group. |
| Flags | PIM flags: <ul style="list-style-type: none"> ■ 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. |
| 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). |
| Upstream neighbor | Information about the upstream neighbor: Direct , Local , Unknown , or a specific IP address. |
| 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 joins nor prunes toward the RP, because this router is the rendezvous point. ■ Local Source—Sending neither joins nor prunes toward the source, because the source is locally attached to this router. ■ Prune to RP—Sending a prune to the rendezvous point. ■ Prune to Source—Sending a prune to the source. |
| Downstream neighbors | Information about downstream interfaces: <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. |
| Assert Timeout | Length of time between assert cycles on downstream interface. Not displayed if assert timer is null. |

Table 57: show pim join Output Fields *(continued)*

| Field Name | Field Description |
|------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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, Timeout is Infinity. |

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

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

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

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

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

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

```

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

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

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

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

```

show pim join extensive

```

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

Group: 239.1.1.1
  Source: *
  RP: 10.255.14.144
  Flags: sparse,rptree,wildcard
  Upstream interface: Local
  Upstream neighbor: Local
  Upstream state: Local RP
  Downstream neighbors:
    Interface: so-1/0/0.0
      10.111.10.2 State: Join Flags: SRW Timeout: 174
    Interface: mt-1/1/0.32768
      10.10.47.100 State: Join Flags: SRW Timeout: Infinity

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
  Downstream neighbors:
    Interface: so-1/0/0.0
      10.111.10.2 State: Join Flags: S Timeout: 174
    Interface: mt-1/1/0.32768
      10.10.47.100 State: Join Flags: S Timeout: Infinity

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
  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

```

```
Interface: mt-1/1/0.32768
10.10.47.100 State: Join Flags: S Timeout: Infinity
```

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

**show pim join instance
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
Downstream neighbors:
Interface: mt-1/1/0.32768
10.10.47.101 State: Join Flags: SRW Timeout: 156
```

```
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
```

```
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
```

show pim mdt

| | |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show pim mdt instance <i>instance-name</i> <brief detail extensive> <incoming outgoing> <logical-system (all <i>logical-system-name</i>)> <range> |
| Release Information | Command introduced before JUNOS 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 | <p>show pim mdt instance on page 185</p> <p>show pim mdt instance detail on page 185</p> <p>show pim mdt instance extensive on page 186</p> <p>show pim mdt instance incoming on page 186</p> <p>show pim mdt instance outgoing on page 186</p> <p>show pim mdt instance (SSM Mode) on page 186</p> |
| Output Fields | Table 58 on page 184 describes the output fields for the show pim mdt command. Output fields are listed in the approximate order in which they appear. |

Table 58: 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 |
| Default group address | Default multicast group address using this tunnel. | All levels |

Table 58: show pim mdt Output Fields (continued)

| Field Name | Field Description | Level of Output |
|----------------------------|-------------------------------------------------------------------------------------------------------|-----------------|
| Default source address | Default multicast source address using this tunnel. | All levels |
| Default tunnel interface | Default multicast tunnel interface. | All levels |
| C-Group | Customer-facing multicast group address using this tunnel. | detail |
| C-Source | IP address of the multicast source in the customer's address space. | 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 |

```

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

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

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

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

| | | | |
|-----------------|------------------|-----------------|-----------------------|
| 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

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 mvpn

| | |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show pim mvpn <logical-system (all logical-system-name)> |
| Release Information | Command introduced in JUNOS 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 187 |
| Output Fields | Table 59 on page 187 describes the output fields for the show pim mvpn command. Output fields are listed in the approximate order in which they appear. |

Table 59: 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 | Mode the tunnel is operating in: PIM-SSM or PIM-ASM. Mode the tunnel is operating in: PIM-ASM, PIM-SSM, NGEN-PMSI or None. If NGEN-PMSI is displayed, enter the show mvpn instance command for more information. | All levels |

```

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 | show pim neighbors <brief detail> <inet inet6> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| 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 all routing instances on all logical systems.</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 <i>instance-name</i>—(Optional) Display information about neighbors for the specified 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 | <p>show pim neighbors on page 189</p> <p>show pim neighbors brief on page 190</p> <p>show pim neighbors instance on page 190</p> <p>show pim neighbors detail on page 190</p> <p>show pim neighbors detail (with BFD) on page 190</p> |
| Output Fields | Table 60 on page 188 describes the output fields for the show pim neighbors command. Output fields are listed in the approximate order in which they appear. |

Table 60: 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 router. | 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 |

Table 60: show pim neighbors Output Fields (continued)

| Field Name | Field Description | Level of Output |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| 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 <i>dd:hh:mm:ss ago</i> for less than a week and <i>nwnd:hh:mm:ss ago</i> 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- 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 <i>delay nnn ms override nnnn ms</i> . | detail |
| 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 |

```

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 189.

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: fe-3/0/2.0
 Address: 192.168.195.37, IPv4, PIM v2, Mode: Sparse
 Hello Option Holdtime: 65535 seconds
 Hello Option DR Priority: 1
 Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
 Join Suppression supported

| Rx Join: Group | Source | Timeout |
|----------------|----------------|---------|
| 225.1.1.1 | 192.168.195.78 | 0 |
| 225.1.1.1 | | 0 |

Interface: lo0.0
 Address: 10.255.245.91, IPv4, PIM v2, Mode: Sparse
 Hello Option Holdtime: 65535 seconds
 Hello Option DR Priority: 1
 Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
 Join Suppression supported

Interface: pd-6/0/0.32768
 Address: 0.0.0.0, IPv4, PIM v2, Mode: Sparse
 Hello Option Holdtime: 65535 seconds
 Hello Option DR Priority: 0
 Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
 Join Suppression supported

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

 Address: 192.168.11.2, IPv4, PIM v2
 BFD: Enabled, Operational state is up
 Hello Default Holdtime: 105 seconds 104 remaining
 Hello Option DR Priority: 1
 Hello Option Generation ID: 1907549685
 Hello Option LAN Prune Delay: delay 500 ms override 2000 ms

Interface: fe-1/0/1.0
 Address: 192.168.12.1, IPv4, PIM v2
 BFD: Disabled
 Hello Default Holdtime: 105 seconds 80 remaining
 Hello Option DR Priority: 1

```
Hello Option Generation ID: 1971554705  
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

show pim rps

| | |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | <pre>show pim rps <brief detail extensive> <group-address> <inet inet6> <instance instance-name> <logical-system (all logical-system-name)></pre> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display information about Protocol Independent Multicast (PIM) rendezvous points (RPs). |
| Options | <p>none—Display standard information about PIM RPs for all groups and family addresses for all routing instances on all logical systems.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>group-address—(Optional) Display the RPs for a particular group. If you specify a group address, the output lists the router that is the RP for that group.</p> <p>inet inet6—(Optional) Display information for IPv4 or IPv6 family addresses, respectively.</p> <p>instance instance-name—(Optional) Display information about RPs for a specific PIM-enabled routing instance.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | <p>show pim rps on page 194</p> <p>show pim rps brief on page 194</p> <p>show pim rps instance on page 195</p> <p>show pim rps extensive on page 195</p> <p>show pim rps extensive (PIM Anycast RP in Use) on page 195</p> |
| Output Fields | Table 61 on page 192 describes the output fields for the show pim rps command. Output fields are listed in the approximate order in which they appear. |

Table 61: show pim rps Output Fields

| Field Name | Field Description | Level of Output |
|------------|------------------------------------------------------------------------|-----------------|
| Instance | Name of the routing instance. | All levels |
| Family | Name of the address family: inet (IPv4) or inet6 (IPv6). | All levels |
| RP address | Address of the rendezvous point. | All levels |

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

| Field Name | Field Description | Level of Output |
|------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| 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 router 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 |
| Time Active | How long the RP has been active, in the format <i>hh:mm:ss</i> . | detail extensive |
| Device Index | Index value of the order in which the JUNOS Software finds and initializes the interface. | detail extensive |
| Subunit | Logical unit number of the interface. | detail extensive |
| Interface | Either the encapsulation or the de-encapsulation logical interface, depending on whether this router is a designated router (DR) facing an RP router, or is the local RP, respectively. | detail extensive |
| Group Ranges | Addresses of groups that this RP spans. | detail extensive |
| Active groups using RP | Number of groups currently using this RP. | detail extensive |
| total | Total number of active groups for this RP. | detail extensive |

Table 61: show pim rps Output Fields (continued)

| Field Name | Field Description | Level of Output |
|--------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Register State for RP | <p>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 router is a designated router facing an RP router, or is the local RP, respectively: ■ First Hop—PIM-designated router 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: <ul style="list-style-type: none"> 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 |
| 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 router 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 |

```

show pim rps      user@host> show pim rps
                    Instance: PIM.master
                    Address family INET
                    RP address      Type      Holdtime Timeout Groups Group prefixes
                    10.255.14.144  static      0       None      1 224.0.0.0/4

                    Address family INET6

```

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 194.


```

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  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  user@host> 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 source

| | |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show pim source <brief detail> <inet inet6> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> <source-prefix> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| 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 on all logical systems.</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 198</p> <p>show pim source brief on page 198</p> <p>show pim source detail on page 198</p> |
| Output Fields | Table 62 on page 197 describes the output fields for the show pim source command. Output fields are listed in the approximate order in which they appear. |

Table 62: show pim source Output Fields

| Field Name | Field Description |
|--------------------|-----------------------------------------------------------------------|
| Instance | Name of the routing instance. |
| RPF Address | 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. |
| Neighbor address | Address of the RPF neighbor used to reach the source address. |

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 198.

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 | show pim statistics <instance <i>instance-name</i> > <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display Protocol Independent Multicast (PIM) statistics. |
| Options | <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 Topics | clear pim statistics |
| List of Sample Output | show pim statistics on page 200 |
| Output Fields | Table 63 on page 199 describes the output fields for the show pim statistics command. Output fields are listed in the approximate order in which they appear. |

Table 63: show pim statistics Output Fields

| Field Name | Field Description |
|-------------------------------------------|-------------------------------------------------------------------|
| 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. |
| PIM statistics summary for all interfaces | Summary for all interfaces. |

show pim statisticsuser@host> **show pim statistics**

PIM statistics on all interfaces:

| PIM message type | Received | Sent | Rx errors |
|-----------------------|----------|-------|-----------|
| V2 Hello | 0 | 0 | 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 | 2102 | 4203 | 0 |
| V1 Register | 0 | 0 | 0 |
| V1 Register Stop | 0 | 0 | 0 |
| V1 Join Prune | 14153 | 13074 | 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 |

PIM statistics summary for all interfaces:

| | |
|----------------------------------|------|
| Hello dropped on neighbor policy | 35 |
| 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 | 1998 |
| Rx V1 Require V2 | 0 |
| Rx Register not RP | 0 |
| RP Filtered Source | 0 |
| Unknown Reg Stop | 0 |
| Rx Join/Prune no state | 1034 |
| 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 Register filtering drop | 0 |
| Tx Register filtering drop | 0 |
| Rx RP no state | 0 |
| Rx aggregate | 0 |
| Rx malformed packet | 0 |
| No RP | 0 |
| No route upstream | 0 |
| RP mismatch | 0 |
| RPF neighbor unknown | 0 |

show sap listen

| | |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show sap listen <brief detail> <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS 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 for all logical systems.</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 | <p>show sap listen on page 201</p> <p>show sap listen brief on page 201</p> <p>show sap listen detail on page 201</p> |
| Output Fields | Table 64 on page 201 describes the output fields for the show sap listen command. Output fields are listed in the approximate order in which they appear. |

Table 64: 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 |

| | |
|-------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| show sap listen | <pre>user@host> show sap listen Group address Port 224.2.127.254 9875 239.255.255.255 9875</pre> |
| 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 201. |
| 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 201. |

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 Release 7.4. |
| 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 202 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| test msdp dependent-peers | user@host> test msdp dependent-peers 10.0.0.1/24 |

Chapter 6

IPv6 Operational Mode Commands

Table 65 on page 203 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 65: IPv6 Operational Mode Commands

| Task | Command |
|-------------------------------------------|---------------------------------|
| Clear IPv6 neighbor cache information. | clear ipv6 neighbors |
| Clear IPv6 router advertisement counters. | clear ipv6 router-advertisement |
| Display neighbor discovery information. | show ipv6 neighbors |
| Display router advertisement information. | show ipv6 router-advertisement |



NOTE: For information about how to configure IPv6 parameters, see the *JUNOS Routing Protocols Configuration Guide*.

clear ipv6 neighbors

| | |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | clear ipv6 neighbors <all host <i>hostname</i> > |
| Release Information | Command introduced before JUNOS Release 7.4. |
| 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 Topics | show ipv6 neighbors |
| List of Sample Output | clear ipv6 neighbors on page 204 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| 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 Release 7.4. |
| Description | Clear IPv6 router advertisement counters. |
| Options | <p>none—Clear IPv6 router advertisement counters for all interfaces on all logical systems.</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 Topics | show ipv6 router-advertisement |
| List of Sample Output | clear ipv6 router-advertisement on page 205 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| clear ipv6 router-advertisement | user@host> clear ipv6 router-advertisement |

show ipv6 neighbors

| | |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show ipv6 neighbors |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display information about the IPv6 neighbor cache. |
| Options | This command has no options. |
| Required Privilege Level | view |
| Related Topics | clear ipv6 neighbors |
| List of Sample Output | show ipv6 neighbors on page 206 show ipv6 neighbors on page 206 |
| Output Fields | Table 66 on page 206 describes the output fields for the show ipv6 neighbors command. Output fields are listed in the approximate order in which they appear. |

Table 66: show ipv6 neighbors Output Fields

| Field Name | Field Description |
|-------------------|---------------------------------------------------------------------------------------------------------------|
| IPv6 Address | Name of 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 router: yes or no . |
| Secure | Whether this entry was created using the Secure Neighbor Discovery (SEND) protocol: yes or no . |
| Interface | Name of the interface. |

| | |
|----------------------------|---------------------------------------------------------------------------------------------|
| show ipv6 neighbors | user@host> show ipv6 neighbors |
| | IPv6 Address Linklayer Address State Exp Rtr Interface |
| | fe80::2a0:c9ff:fe5b:4c1e 00:a0:c9:5b:4c:1e reachable 15 yes fxp0.0 |
| show ipv6 neighbors | user@host > show ipv6 neighbors |
| | IPv6 Address Linklayer Address State Exp Rtr Secure |
| | Interface |
| | fe80::14fb:5dcf:54bd:ff76 00:90:69:a0:a8:bc stale 1113 yes yes |
| | ge-3/2/0.0 |

show ipv6 router-advertisement

| | |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show ipv6 router-advertisement <conflicts> <interface <i>interface</i> > <logical-system (all <i>logical-system-name</i>)> <prefix <i>prefix/prefix length</i> > |
| Release Information | Command introduced before JUNOS Release 7.4. |
| 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 on all logical systems.</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 Topics | clear ipv6 router-advertisement |
| List of Sample Output | <p>show ipv6 router-advertisement on page 208</p> <p>show ipv6 router-advertisement conflicts on page 208</p> <p>show ipv6 router-advertisement prefix on page 209</p> |
| Output Fields | Table 67 on page 207 describes the output fields for the show ipv6 router-advertisement command. Output fields are listed in the approximate order in which they appear. |

Table 67: show ipv6 router-advertisement Output Fields

| Field Name | Field Description |
|---------------------|-----------------------------------------------------------------------------|
| Interface | Name of the interface. |
| Advertisements sent | Number of router advertisements sent and elapsed time since they were sent. |
| Solicits received | Number of solicitation messages received. |

Table 67: show ipv6 router-advertisement Output Fields *(continued)*

| Field Name | Field Description |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| Advertisements received | Number of router advertisements received. |
| Advertisements from | Names of interfaces from which router advertisements have been received and 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). |

```

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 68 on page 211 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 68: IS-IS Operational Mode Commands

| Task | Command |
|----------------------------------------------------------------------------------------|--------------------------------------|
| Remove adjacencies. | clear isis adjacency |
| Remove database entries. | clear isis database |
| Reset IS-IS dynamic overload bit. | clear isis overload |
| Set IS-IS traffic statistics to zero. | clear isis statistics |
| Display adjacent routers. | show isis adjacency |
| Display authentication statistics. | show isis authentication |
| Display information about the level of backup coverage available for protected routes. | show isis backup coverage |
| Display information about MPLS LSPs designated as backup paths. | show isis backup label-switched-path |
| Display SPF calculations for backup paths. | show isis backup spf results |
| Display database entries. | show isis database |
| Display hostname mapping. | show isis hostname |
| Display the status of interfaces on which IS-IS is running. | show isis interface |
| Display IS-IS overview information. | show isis overview |
| Display IS-IS routing table entries. | show isis route |
| Display SPF calculations. | show isis spf |
| Display IS-IS traffic statistics. | show isis statistics |



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” on page 321. For information about monitoring Bidirectional Forwarding Detection (BFD) sessions for IS-IS clients, see “BFD Operational Mode Commands” on page 15. 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 *instance-name*>
 <interface *interface-name*>
 <logical-system (all | *logical-system-name*)>
 <neighbor>

Release Information Command introduced before JUNOS Release 7.4.

Description Remove entries from the Intermediate System-to-Intermediate System (IS-IS) adjacency database.

Options none—Remove all entries from the adjacency database.

instance instance-name—(Optional) Clear all adjacencies for the specified routing instance only.

interface interface-name—(Optional) Clear all adjacencies for the specified interface only.

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

neighbor—(Optional) Clear adjacencies for the specified neighbor only.

Required Privilege Level clear

Related Topics show isis adjacency

List of Sample Output clear isis adjacency on page 213

Output Fields See show isis adjacency for an explanation of output fields.

clear isis adjacency The following sample output displays IS-IS adjacency database information before and after the clear isis adjacency command is entered:

```
user@host> show isis adjacency
IS-IS adjacency database:
Interface      System          L State      Hold (secs) SNPA
so-1/0/0.0     karaku1         3 Up          26
so-1/1/3.0     1921.6800.5080 3 Up          23
so-5/0/0.0     1921.6800.5080 3 Up          19
```

```
user@host> clear isis adjacency karaku1
```

```
user@host> show isis adjacency
IS-IS adjacency database:
Interface      System          L State      Hold (secs) SNPA
so-1/0/0.0     karaku1         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 *instance-name*>
 <logical-system (all | *logical-system-name*)>
 <purge>

Release Information Command introduced before JUNOS Release 7.4.
 purge option introduced in JUNOS Release 9.0.

Description Remove the entries from the Intermediate System-to-Intermediate System (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 (LSPs) rather than the local deletion of entries from the IS-IS link-state database.



CAUTION: In a production network, the **purge** command option may cause short-term network-wide traffic disruptions. Use with caution!

Options none—Remove all entries from the IS-IS link-state database for all routing instances on all logical systems.

entries—(Optional) Name of the database entry.

instance *instance-name*—(Optional) Clear all entries for the specified routing instance.

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

purge—(Optional) Discard all entries in the IS-IS link-state database.

Required Privilege Level clear

Related Topics show isis database

List of Sample Output clear isis database on page 214

Output Fields See show isis database for an explanation of output fields.

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 <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Reset the Intermediate System-to-Intermediate System (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. |
| Options | none—Reset the IS-IS dynamic overload bit on all logical systems. 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 Topics | show isis database |
| List of Sample Output | clear isis overload on page 216 |
| Output Fields | See show isis database for an explanation of output fields. |

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>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Set statistics about Intermediate System-to-Intermediate System (IS-IS) traffic to zero. |
| Options | <p>none—Set IS-IS traffic statistics to zero for all routing instances on all logical systems.</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 Topics | show isis statistics |
| List of Sample Output | clear isis statistics on page 217 |
| Output Fields | See show isis statistics for an explanation of output fields. |
| clear isis statistics | The following sample output displays IS-IS statistics before and after the clear isis statistics command is entered: |

```
user@host> show isis statistics
IS-IS statistics for merino:
```

| PDU type | Received | Processed | Drops | Sent | Rexmit |
|----------|----------|-----------|-------|--------|--------|
| LSP | 12793 | 12793 | 0 | 8666 | 719 |
| IIH | 116751 | 116751 | 0 | 118834 | 0 |
| CSNP | 203956 | 203956 | 0 | 204080 | 0 |
| PSNP | 7356 | 7350 | 6 | 8635 | 0 |
| Unknown | 0 | 0 | 0 | 0 | 0 |
| Totals | 340856 | 340850 | 6 | 340215 | 719 |

```
Total packets received: 340856 Sent: 340934
```

```
SNP queue length:          0 Drops:          0
LSP queue length:          0 Drops:          0
```

```
SPF runs:                  1064
Fragments rebuilt:         1087
LSP regenerations:         436
Purges initiated:          0
```

```
user@host> clear isis statistics
```

```
user@host> show isis statistics
IS-IS statistics for merino:
```

| PDU type | Received | Processed | Drops | Sent | Rexmit |
|----------|----------|-----------|-------|------|--------|
|----------|----------|-----------|-------|------|--------|

| | | | | | |
|---------|---|---|---|---|---|
| LSP | 0 | 0 | 0 | 0 | 0 |
| IIH | 3 | 3 | 0 | 3 | 0 |
| CSNP | 2 | 2 | 0 | 4 | 0 |
| PSNP | 0 | 0 | 0 | 0 | 0 |
| Unknown | 0 | 0 | 0 | 0 | 0 |
| Totals | 5 | 5 | 0 | 7 | 0 |

Total packets received: 5 Sent: 7

| | | | |
|-------------------|---|--------|---|
| SNP queue length: | 0 | Drops: | 0 |
| LSP queue length: | 0 | Drops: | 0 |

| | |
|--------------------|---|
| SPF runs: | 0 |
| Fragments rebuilt: | 0 |
| LSP regenerations: | 0 |
| Purges initiated: | 0 |

show isis adjacency

| | |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show isis adjacency <brief detail extensive> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display information about Intermediate System-to-Intermediate System (IS-IS) neighbors. |
| Options | <p>none—Display standard information about IS-IS neighbors for all routing instances on all logical systems.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>instance <i>instance-name</i>—(Optional) Display adjacencies for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| Related Topics | clear isis adjacency |
| List of Sample Output | <p>show isis adjacency on page 221</p> <p>show isis adjacency brief on page 221</p> <p>show isis adjacency detail on page 221</p> <p>show isis adjacency extensive on page 222</p> |
| Output Fields | Table 69 on page 219 describes the output fields for the show isis adjacency command. Output fields are listed in the approximate order in which they appear. |

Table 69: show isis adjacency Output Fields

| Field Name | Field Description | Level of Output |
|------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Interface | Interface through which the neighbor is reachable. | All levels |
| 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 |

Table 69: show isis adjacency Output Fields (continued)

| Field Name | Field Description | Level of Output |
|------------------------------------|--------------------------------------------------------------------------------------------------------------------------|------------------|
| 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: L1 = Level 1 router; L2 = Level 2 router; L1/L2 = 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 not 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 the interface in the router's outbound LSPs. | detail extensive |
| IP addresses | IP address of this neighbor. | detail extensive |

Table 69: 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 |

show isis adjacency user@host> **show isis adjacency**

```
Interface          System      L State      HoId (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 221.

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 <brief detail extensive> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced in JUNOS Release 7.5. |
| Description | Display information about Intermediate System-to-Intermediate System (IS-IS) authentication. |
| Options | <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>instance <i>instance-name</i>—(Optional) Display SPF calculations 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 224 |
| Output Fields | Table 70 on page 223 describes the output fields for the show isis authentication command. Output fields are listed in the approximate order in which they appear. |

Table 70: 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. |
| CSN Auth | Complete sequence number authentication type. |
| PSN Auth | Partial sequence number authentication type. |
| L1 LSP Authentication | Layer 1 link-state PDU authentication type. |
| L2 LSP Authentication | Layer 2 link-state PDU authentication type. |

```
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 backup coverage

| | |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show isis backup coverage <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced in JUNOS Release 9.5. |
| Description | Display information about the level of backup coverage available for all the nodes and prefixes in the network. |
| Options | <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 Topics | show isis backup label-switched-path |
| List of Sample Output | show isis backup coverage on page 225 |
| Output Fields | Table 71 on page 225 lists the output fields for the show isis backup coverage command. Output fields are listed in the approximate order in which they appear. |

Table 71: 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. |
| IPv4 Unicast | Percentage of IPv4 unicast routes that are protected through backup coverage. |
| IPv6 Unicast | 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. |

```

show isis backup coverage  user@host> show isis backup coverage
Backup Coverage:
  Topology      Level1  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>)> |
| Release Information | Command introduced in JUNOS Release 9.5. |
| Description | Display information about MPLS label-switched-paths (LSPs) designated as backup routes for IS-IS routes. |
| Options | logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | view |
| Related Topics | show isis backup coverage |
| List of Sample Output | show isis backup label-switched-path on page 227 |
| Output Fields | Table 72 on page 227 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 72: 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 router 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. |
| 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. |

```

show isis backup      user@host> show isis backup label-switched-path
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 | 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-multicast ipv6-multicast ipv4-unicast unicast)> |
| Release Information | Command introduced in JUNOS 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 shortest-path-first (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 Topics | show isis backup coverage |
| List of Sample Output | show isis backup spf results on page 229 |
| Output Fields | Table 73 on page 228 lists the output fields for the show isis backup spf results command. Output fields are listed in the approximate order in which they appear. |

Table 73: 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. |

Table 73: show isis backup spf results Output Fields (continued)

| Field Name | Field Description |
|-----------------|----------------------------------------------------------------------------------------------------|
| 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 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. |

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:
```

```
kobuk.00, Address 0x8d85600
```

```
Primary next-hop: ge-0/2/0.0, camaro, SNPA: 0:90:69:f:62:fa
```

```
Primary next-hop: so-0/1/2.0, crater
```

```
Primary next-hop: ge-0/2/0.0, camaro, SNPA: 0:90:69:f:62:fa
```

```
Primary next-hop: so-0/1/2.0, crater
```

```
Root: crater, Metric: 10
```

```
Not eligible, Reason: Primary next-hop multipath
```

```
Root: camaro, Metric: 10
```

```
Not eligible, Reason: Primary next-hop multipath
```

```
Root: olympic, Metric: 25
```

```
Not eligible, Reason: Primary next-hop multipath
```

```
glacier.00, Address 0x8d85200
```

```
Primary next-hop: so-0/1/2.0, crater
```

```
Primary next-hop: so-0/1/2.0, crater
```

```
Root: crater, Metric: 10
```

```
Not eligible, Reason: Primary next-hop link fate sharing
```

```
Root: olympic, Metric: 15
```

```
Eligible, Backup next-hop: ge-0/2/0.0, camaro, SNPA: 0:90:69:f:62:fa
```

```
Eligible, Backup next-hop: so-1/0/2.0, olympic
```

```
Eligible, Backup next-hop: ge-0/2/0.0, camaro, SNPA: 0:90:69:f:62:fa
```

```
Eligible, Backup next-hop: so-1/0/2.0, olympic
```

```
Root: camaro, Metric: 20
```

```
Eligible, Backup next-hop: ge-0/2/0.0, camaro, SNPA: 0:90:69:f:62:fa
```

```
Eligible, Backup next-hop: so-1/0/2.0, olympic
```

```
Eligible, Backup next-hop: ge-0/2/0.0, camaro, SNPA: 0:90:69:f:62:fa
```

```
Eligible, Backup next-hop: so-1/0/2.0, olympic
```

```
olympic.00, Address 0x8d00c00
```

```
Primary next-hop: so-1/0/2.0, olympic
```

```
Primary next-hop: so-1/0/2.0, olympic
```

```
Root: olympic, Metric: 0
```

```
Not eligible, Reason: Primary next-hop link fate sharing
```

```
Root: crater, Metric: 20
```

```

    track-item: olympic.00-00
    track-item: banff.00-00
    Not eligible, Reason: Path loops
Root: camaro, Metric: 20
    track-item: olympic.00-00
    track-item: banff.00-00
    Not eligible, Reason: Path loops
camaro.00, Address 0x8d85a00
Primary next-hop: ge-0/2/0.0, camaro, SNPA: 0:90:69:f:62:fa
Primary next-hop: ge-0/2/0.0, camaro, SNPA: 0:90:69:f:62:fa
Root: camaro, Metric: 0
    Not eligible, Reason: Primary next-hop link fate sharing
Root: crater, Metric: 20
    track-item: camaro.00-00
    track-item: banff.00-00
    Not eligible, Reason: Path loops
Root: olympic, Metric: 20
    track-item: camaro.00-00
    track-item: banff.00-00
    Not eligible, Reason: Path loops
crater.00, Address 0x8d85000
Primary next-hop: so-0/1/2.0, crater
Primary next-hop: so-0/1/2.0, crater
Root: crater, Metric: 0
    Not eligible, Reason: Primary next-hop link fate sharing
Root: camaro, Metric: 20
    track-item: crater.00-00
    track-item: banff.00-00
    Not eligible, Reason: Path loops
Root: olympic, Metric: 20
    track-item: crater.00-00
    track-item: banff.00-00
    Not eligible, Reason: Path loops
5 nodes

```

show isis database

| | |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show isis database <brief detail extensive> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display the entries in the Intermediate System-to-Intermediate System (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 on all logical systems.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</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> |
| Required Privilege Level | view |
| Related Topics | clear isis database |
| List of Sample Output | <p>show isis database on page 232</p> <p>show isis database brief on page 233</p> <p>show isis database detail on page 233</p> <p>show isis database extensive on page 235</p> <p>show isis database extensive (CLNS) on page 236</p> |
| Output Fields | Table 74 on page 231 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 74: show isis database Output Fields

| Field Name | Field Description | Level of Output |
|-----------------|-------------------------------------------------------|------------------|
| 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 |
| IP prefix | Prefix advertised by this link-state PDU. | detail extensive |

Table 74: show isis database Output Fields (continued)

| Field Name | Field Description | Level of Output |
|-------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| 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 router. ■ Remaining Lifetime—Remaining lifetime of the link-state PDU, in seconds. | extensive |
| Packet | <ul style="list-style-type: none"> ■ LSP ID—The identifier for the link-state packet. ■ Length—Packet length. ■ Lifetime—Remaining lifetime, in seconds. ■ Checksum—The checksum of the LSP. ■ Sequence—The sequence number of the LSP. Every time the LSP 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 router can reach. ■ Speaks—Supported routing protocols. ■ IP router id—ID of the router (usually the IP address). ■ IP address—IPv4 address. ■ Hostname—Assigned name of the router. ■ IP prefix—IP prefix of the router. ■ Metric—IS-IS metric that measures the cost of the adjacency between the originating router and the advertised router. ■ IP extended prefix—Extended IP prefix of the router. ■ IS neighbor—Directly attached neighbor's name and metric. ■ IS extended neighbor—Directly attached neighbor's name, metric, and IP address. | extensive |

```

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 232.

show isis database detail user@host> **show isis database detail**
IS-IS level 1 link-state database:

```

kobuk.00-00 Sequence: 0x3, Checksum: 0x3167, Lifetime: 1048 secs
  IS neighbor: glacier.00          Metric:      10
  IP prefix: 10.255.70.103/32      Metric:      0 Internal Up
  IP prefix: 43.1.1.0/24          Metric:     10 Internal Up
  V6 prefix: abcd::10:255:70:103/128 Metric:      0 Internal Up

camaro.00-00 Sequence: 0x5, Checksum: 0x770e, Lifetime: 1082 secs
  IS neighbor: ranier.00          Metric:      10
  IS neighbor: glacier.02        Metric:      10
  IP prefix: 10.255.71.52/32      Metric:      0 Internal Up
  IP prefix: 23.1.1.0/24         Metric:     10 Internal Up
  IP prefix: 34.1.1.0/24         Metric:     10 Internal Up
  V6 prefix: abcd::10:255:71:52/128 Metric:      0 Internal Up

ranier.00-00 Sequence: 0x4, Checksum: 0xaa95, Lifetime: 1082 secs
  IS neighbor: camaro.00          Metric:      10
  IS neighbor: badlands.00       Metric:      10
  IP prefix: 10.255.71.241/32     Metric:      0 Internal Up
  IP prefix: 11.1.1.0/24         Metric:     10 Internal Up
  IP prefix: 23.1.1.0/24         Metric:     10 Internal Up
  V6 prefix: abcd::10:255:71:241/128 Metric:      0 Internal Up

glacier.00-00 Sequence: 0x4, Checksum: 0x206f, Lifetime: 1080 secs
  IS neighbor: kobuk.00          Metric:      10
  IS neighbor: glacier.02        Metric:      10
  IP prefix: 10.255.71.242/32     Metric:      0 Internal Up
  IP prefix: 34.1.1.0/24         Metric:     10 Internal Up
  IP prefix: 43.1.1.0/24         Metric:     10 Internal Up
  V6 prefix: abcd::10:255:71:242/128 Metric:      0 Internal Up

glacier.02-00 Sequence: 0x1, Checksum: 0xd141, Lifetime: 1080 secs
  IS neighbor: camaro.00          Metric:       0
  IS neighbor: glacier.00        Metric:       0

badlands.00-00 Sequence: 0x3, Checksum: 0x87a2, Lifetime: 1084 secs
  IS neighbor: ranier.00          Metric:      10
  IP prefix: 10.255.71.244/32     Metric:      0 Internal Up
  IP prefix: 11.1.1.0/24         Metric:     10 Internal Up
  V6 prefix: abcd::10:255:71:244/128 Metric:      0 Internal Up

```

IS-IS level 2 link-state database:

```

kobuk.00-00 Sequence: 0x6, Checksum: 0x8d6b, Lifetime: 1088 secs
  IS neighbor: glacier.00                      Metric: 10
  IP prefix: 10.255.70.103/32                  Metric: 0 Internal Up
  IP prefix: 10.255.71.52/32                   Metric: 20 Internal Up
  IP prefix: 10.255.71.241/32                  Metric: 30 Internal Up
  IP prefix: 10.255.71.242/32                  Metric: 10 Internal Up
  IP prefix: 10.255.71.244/32                  Metric: 40 Internal Up
  IP prefix: 11.1.1.0/24                      Metric: 40 Internal Up
  IP prefix: 23.1.1.0/24                      Metric: 30 Internal Up
  IP prefix: 34.1.1.0/24                      Metric: 20 Internal Up
  IP prefix: 43.1.1.0/24                      Metric: 10 Internal Up
  V6 prefix: abcd::10:255:70:103/128          Metric: 0 Internal Up

camaro.00-00 Sequence: 0x9, Checksum: 0x877b, Lifetime: 1092 secs
  IS neighbor: ranier.00                      Metric: 10
  IS neighbor: glacier.02                    Metric: 10
  IP prefix: 10.255.70.103/32                  Metric: 20 Internal Up
  IP prefix: 10.255.71.52/32                   Metric: 0 Internal Up
  IP prefix: 10.255.71.241/32                  Metric: 10 Internal Up
  IP prefix: 10.255.71.242/32                  Metric: 10 Internal Up
  IP prefix: 10.255.71.244/32                  Metric: 20 Internal Up
  IP prefix: 11.1.1.0/24                      Metric: 20 Internal Up
  IP prefix: 23.1.1.0/24                      Metric: 10 Internal Up
  IP prefix: 34.1.1.0/24                      Metric: 10 Internal Up
  IP prefix: 43.1.1.0/24                      Metric: 20 Internal Up
  V6 prefix: abcd::10:255:71:52/128           Metric: 0 Internal Up

ranier.00-00 Sequence: 0x8, Checksum: 0x855d, Lifetime: 1094 secs
  IS neighbor: camaro.00                      Metric: 10
  IS neighbor: badlands.00                   Metric: 10
  IP prefix: 10.255.70.103/32                  Metric: 30 Internal Up
  IP prefix: 10.255.71.52/32                   Metric: 10 Internal Up
  IP prefix: 10.255.71.241/32                  Metric: 0 Internal Up
  IP prefix: 10.255.71.242/32                  Metric: 20 Internal Up
  IP prefix: 10.255.71.244/32                  Metric: 10 Internal Up
  IP prefix: 11.1.1.0/24                      Metric: 10 Internal Up
  IP prefix: 23.1.1.0/24                      Metric: 10 Internal Up
  IP prefix: 34.1.1.0/24                      Metric: 20 Internal Up
  IP prefix: 43.1.1.0/24                      Metric: 30 Internal Up
  V6 prefix: abcd::10:255:71:241/128          Metric: 0 Internal Up

glacier.00-00 Sequence: 0x7, Checksum: 0xf892, Lifetime: 1089 secs
  IS neighbor: kobuk.00                      Metric: 10
  IS neighbor: glacier.02                    Metric: 10
  IP prefix: 10.255.70.103/32                  Metric: 10 Internal Up
  IP prefix: 10.255.71.52/32                   Metric: 10 Internal Up
  IP prefix: 10.255.71.241/32                  Metric: 20 Internal Up
  IP prefix: 10.255.71.242/32                  Metric: 0 Internal Up
  IP prefix: 10.255.71.244/32                  Metric: 30 Internal Up
  IP prefix: 11.1.1.0/24                      Metric: 30 Internal Up
  IP prefix: 23.1.1.0/24                      Metric: 20 Internal Up
  IP prefix: 34.1.1.0/24                      Metric: 10 Internal Up
  IP prefix: 43.1.1.0/24                      Metric: 10 Internal Up
  V6 prefix: abcd::10:255:71:242/128          Metric: 0 Internal Up

glacier.02-00 Sequence: 0x1, Checksum: 0xd141, Lifetime: 1080 secs
  IS neighbor: camaro.00                      Metric: 0
  IS neighbor: glacier.00                    Metric: 0

```



```

badlands.00-00 Sequence: 0x6, Checksum: 0x562, Lifetime: 1096 secs
  IS neighbor: ranier.00 Metric: 10
  IP prefix: 10.255.70.103/32 Metric: 40 Internal Up
  IP prefix: 10.255.71.52/32 Metric: 20 Internal Up
  IP prefix: 10.255.71.241/32 Metric: 10 Internal Up
  IP prefix: 10.255.71.242/32 Metric: 30 Internal Up
  IP prefix: 10.255.71.244/32 Metric: 0 Internal Up
  IP prefix: 11.1.1.0/24 Metric: 10 Internal Up
  IP prefix: 23.1.1.0/24 Metric: 20 Internal Up
  IP prefix: 34.1.1.0/24 Metric: 30 Internal Up
  IP prefix: 43.1.1.0/24 Metric: 40 Internal Up
  V6 prefix: abcd::10:255:71:244/128 Metric: 0 Internal Up

```

show isis database extensive user@host> **show isis database extensive isis2**
 IS-IS level 1 link-state database:

IS-IS level 2 link-state database:

```

isis2.00-00 Sequence: 0x82, Checksum: 0x6cc3, Lifetime: 1126 secs
  IS neighbor: isis1.00 Metric: 10
  IS neighbor: isis3.00 Metric: 10
  IP prefix: 10.255.245.202/32 Metric: 0 Internal
  IP prefix: 192.168.36.0/29 Metric: 10 Internal
  IP prefix: 192.168.36.16/30 Metric: 10 Internal
  IP prefix: 192.168.36.24/30 Metric: 10 Internal

```

```

Header: LSP ID: isis2.00-00, Length: 234 bytes
  Allocated length: 234 bytes, Router ID: 10.255.245.202
  Remaining lifetime: 1126 secs, Level: 2, Interface: 4
  Estimated free bytes: 0, Actual free bytes: 0
  Aging timer expires in: 1126 secs
  Protocols: IP, IPv6

Packet: LSP ID: isis2.00-00, Length: 234 bytes, Lifetime : 1198 secs
  Checksum: 0x6cc3, Sequence: 0x82, 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: 47.0005.80ff.f800.0000.0108.0001 (13)
  Speaks: IP
  Speaks: IPv6
  IP router id: 10.255.245.202
  IP address: 10.255.245.202
  Hostname: isis2
  IS neighbor: isis3.00, Internal, Metric: default 10
  IS neighbor: isis1.00, Internal, Metric: default 10
  IS neighbor: isis3.00, Metric: default 10
    IP address: 192.168.36.25
    Neighbor's IP address: 192.168.36.26
  IS neighbor: isis1.00, Metric: default 10
    IP address: 192.168.36.18
    Neighbor's IP address: 192.168.36.17
  IP prefix: 10.255.245.202/32, Internal, Metric: default 0
  IP prefix: 192.168.36.0/29, Internal, Metric: default 10
  IP prefix: 192.168.36.24/30, Internal, Metric: default 10
  IP prefix: 192.168.36.16/30, Internal, Metric: default 10
  IP prefix: 10.255.245.202/32 metric 0 up
    6 bytes of subtlvs
    Administrative tag 1: 1000
  IP prefix: 192.168.36.0/29 metric 10 up

```

```

IP prefix: 192.168.36.24/30 metric 10 up
IP prefix: 192.168.36.16/30 metric 10 up
No queued transmissions

```

**show isis database
extensive (CLNS)**

```

user@host> show isis database extensive
IS-IS level 1 link-state database:
isis2.00-00 Sequence: 0x1256, Checksum: 0x53da, Lifetime: 582 secs
IS neighbor: pro1-a.02 Metric: 10
ES neighbor: toothache Metric: 0
ES neighbor: 1921.6800.4002 Metric: 10
IP prefix: 192.168.37.64/29 Metric: 10 Internal Up

Header: LSP ID: toothache.00-00, Length: 140 bytes
Allocated length: 284 bytes, Router ID: 0.0.0.0
Remaining lifetime: 582 secs, Level: 1, Interface: 66
Estimated free bytes: 144, Actual free bytes: 144
Aging timer expires in: 582 secs
Protocols: IP, CLNS

Packet: LSP ID: toothache.00-00, Length: 140 bytes, Lifetime : 1199 secs
Checksum: 0x53da, Sequence: 0x1256, 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: 47.0005.80ff.f800.0000.0108.0001 (13)
Speaks: CLNP
Speaks: IP
Hostname: toothache
IP address: 192.168.37.69
IP extended prefix: 192.168.37.64/29 metric 10 up
IP prefix: 192.168.37.64/29, Internal, Metric: default 10, Up
IS neighbor: pro1-a.02, Internal, Metric: default 10
IS extended neighbor: pro1-a.02, Metric: default 10
ES neighbor TLV: Internal, Metric: default 0
ES: toothache
ES neighbor TLV: Internal, Metric: default 10
ES: 1921.6800.4002
No queued transmissions

```

show isis hostname

| | |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show isis hostname <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display Intermediate System-to-Intermediate System (IS-IS) hostname database information. |
| Options | none—Display IS-IS hostname database information on all logical systems. 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 hostname on page 237 |
| Output Fields | Table 75 on page 237 describes the output fields for the show isis hostname command. Output fields are listed in the approximate order in which they appear. |

Table 75: show isis hostname Output Fields

| Field Name | Field Description |
|------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| System Id | System identifier mapped to the hostname. |
| Hostname | Hostname mapped to the system identifier. |
| Type | Type of mapping between system identifier and hostname. <ul style="list-style-type: none"> ■ Dynamic—Hostname mapping determined as described in RFC 2763, <i>Dynamic Hostname Exchange Mechanism for IS-IS</i>. ■ Static—Hostname mapping configured by user. |

```

show isis hostname user@host> show isis hostname
IS-IS hostname database:
System Id      Hostname
1921.6800.4201 isis1
1921.6800.4202 isis2
1921.6800.4203 isis3
Type
Dynamic
Static
Dynamic

```

show isis interface

| | |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show isis interface <brief detail extensive> <interface-name> <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display status information about Intermediate System-to-Intermediate System (IS-IS)-enabled interfaces. |
| Options | <p>none—Display standard information about all IS-IS-enabled interfaces on all logical systems.</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 <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 isis interface on page 239</p> <p>show isis interface brief on page 240</p> <p>show isis interface detail on page 240</p> <p>show isis interface extensive on page 240</p> |
| Output Fields | Table 76 on page 238 describes the output fields for the show isis interface command. Output fields are listed in the approximate order in which they appear. |

Table 76: show isis interface Output Fields

| Field Name | Field Description | Level of Output |
|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <i>interface-name</i> | Name of the interface. | detail |
| Designated router | A router 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 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 |

Table 76: show isis interface Output Fields (continued)

| Field Name | Field Description | Level of Output |
|------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------|------------------|
| LSP interval | Interval between link-state PDUs sent from the interface. | detail |
| Sysid | System identifier. | detail |
| Interface | Interface through which the adjacency is made. | 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. | All levels |
| Level 1 DR | Level 1 designated intermediate system. | brief |
| Level 2 DR | Level 2 designated intermediate system. | brief |
| L1/L2 Metric | Interface's metric for Level 1 and Level 2. If there is no information, the metric is 0. | none specified |
| Adjacency advertisement: Advertise | This router has signaled not 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 router's outbound LSPs. | detail extensive |
| Adjacencies | Number of adjacencies established on this interface. | detail |
| Priority | Priority value for this interface. | detail |
| Metric | Metric value for this interface. | detail |
| Hello(s) | Interface's hello interval. | detail |
| Hold(s) | Interface's hold time. | detail |
| 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 number of seconds remaining. | extensive |

```

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 239.

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 overview

| | |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show isis overview <logical-system (all <i>logical-system-name</i>)> <instance <i>instance-name</i> > |
| Release Information | Command introduced in JUNOS Release 8.5. |
| Description | Display Intermediate System-to-Intermediate System (IS-IS) overview information. |
| Options | <p>none—Display standard overview information about IS-IS for all routing instances on all logical systems.</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>instance <i>instance-name</i>—(Optional) Display overview information for the specified routing instance.</p> |
| Required Privilege Level | view |
| List of Sample Output | show isis overview on page 242 |
| Output Fields | Table 77 on page 241 lists the output fields for the show isis overview command. Output fields are listed in the approximate order in which they appear. |

Table 77: show isis overview Output Fields

| Field Name | Field Description |
|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------|
| instance | The IS-IS routing instance. |
| Router ID | Router ID of the router. |
| Adjacency holddown | Adjacency holddown capability: enabled or disabled . |
| Maximum Areas | Maximum number of IS-IS areas advertised by router. |
| 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 calculations. |
| SPF holddown | Delay before performing additional Shortest Path First (SPF) calculations after the maximum number of consecutive SPF calculations is reached. |
| SPF rapid runs | Maximum number of Shortest Path First 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 . |

Table 77: show isis overview Output Fields (continued)

| Field Name | Field Description |
|-------------------------------|-------------------------------------------------------------------------------------------------------------------------|
| 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 | OSI CLNP Protocol capability is enabled. (J Series routers only) |
| 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 is enabled | Narrow metrics capability is enabled. |

show isis overview user@host> **show isis overview**

Sample Output

```
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 <i>instance-name</i>] [logical-system (all <i>logical-system-name</i>)] [topology (ipv4-multicast ipv6-multicast ipv6-unicast unicast)]</pre> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display the routes in the Intermediate System-to-Intermediate System (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 on all logical systems.</p> <p><i>destination</i>—(Optional) Destination address for the route.</p> <p>inet inet6—(Optional) Display inet (IPv4) or inet6 (IPv6) routes, respectively.</p> <p>instance <i>instance-name</i>—(Optional) Display routes 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>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 | <pre>show isis route logical-system on page 245 show isis route (CLNS) on page 246</pre> |
| Output Fields | Table 78 on page 245 describes the output fields for the <code>show isis route</code> command. Output fields are listed in the approximate order in which they appear. |

Table 78: show isis route Output Fields

| Field Name | Field Description |
|-----------------|--------------------------------------------------------------------------------------------------------------------------------------------|
| Current version | Number of the current version of the IS-IS routing table. |
| 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. |

```

show isis route      user@host> show isis route logical-system ls1
logical-system      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 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)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display information about Intermediate System-to-Intermediate System (IS-IS) shortest-path-first (SPF) calculations. |
| Options | <p>brief—Display an overview of SPF calculations.</p> <p>log—Display the log of SPF calculations.</p> <p>results—Display the results 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>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 (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 | <p>show isis spf brief on page 248</p> <p>show isis spf log on page 249</p> <p>show isis spf results on page 250</p> <p>show isis spf results (CLNS) on page 251</p> |
| Output Fields | Table 79 on page 247 describes the output fields for the show isis spf command. Output fields are listed in the approximate order in which they appear. |

Table 79: show isis spf Output Fields

| Field Name | Field Description |
|------------|----------------------------|
| Node | System ID of a node. |
| Metric | Metric to the node. |
| Interface | Interface of the next hop. |
| Via | System ID of the next hop. |

Table 79: show isis spf Output Fields (continued)

| Field Name | Field Description |
|----------------|-----------------------------------------------------------------------------------------|
| 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. |

show isis spf brief

```
user@host> show isis spf brief logical-system ls1
```

```
IS-IS level 1 SPF results:
```

| Node | Metric | Interface | Via | SNPA |
|---------|--------|------------|------|------------------|
| scat.00 | 10 | ge-1/1/0.0 | scat | 0:90:69:a6:48:9d |
| fix.02 | 10 | | | |
| fix.00 | 0 | | | |
| 3 nodes | | | | |

```
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 | | | | |

```
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 |
| fix.02 | 10 | | | |
| fix.00 | 0 | | | |
| 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 | |
| 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 | |
| fix.00 | 0 | | | |
| 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 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:
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
          20       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
          20       gr-0/2/0.0     h
          30       8009:3::a09:3200/126
skag.02   20       gr-0/2/0.0     h
          20       gr-0/2/0.0     h
h.00      10       gr-0/2/0.0     h
          20       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>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display statistics about Intermediate System-to-Intermediate System (IS-IS) traffic. |
| Options | none—Display IS-IS traffic statistics for all routing instances on all logical systems. 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 Topics | clear isis statistics |
| List of Sample Output | show isis statistics on page 253 |
| Output Fields | Table 80 on page 252 describes the output fields for the show isis statistics command. Output fields are listed in the approximate order in which they appear. |

Table 80: show isis statistics Output Fields

| Field Name | Field Description |
|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PDU type | Protocol data unit type: <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; that is, 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 router, in turn, forwards the missing link-state PDU to the requesting router. ■ 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. |

Table 80: show isis statistics Output Fields *(continued)*

| Field Name | Field Description |
|-----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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 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. |

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 81 on page 255 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 81: 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 Release 9.6. |
| Description | On MX Series routers, clear information regarding all Link Layer Discovery Protocol (LLDP) neighbors or LLDP neighbors of the specified interface. |
| Options | interface <i>interface-name</i> —(Optional) Clear the LLDP neighbors on the specified interface. |
| Required Privilege Level | clear |
| Related Topics | clear lldp statistics |
| List of Sample Output | clear lldp statistics on page 256 |
| 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. |
| clear lldp statistics | user@host> clear lldp statistics user@host> clear lldp statistics interface ge-0/2/0 |

clear lldp statistics

| | |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | clear lldp neighbor <interface <i>interface-name</i> > |
| Release Information | Command introduced in JUNOS Release 9.6. |
| Description | On MX Series routers, clear all Link Layer Discovery Protocols (LLDP) statistics or LLDP statistics associated with the specified interface. |
| Options | interface <i>interface-name</i> —(Optional) Clear LLDP statistics on the specified interface. |
| Required Privilege Level | clear |
| Related Topics | clear lldp neighbor |
| List of Sample Output | clear lldp neighbor on page 257 |
| 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. |
| clear lldp neighbor | user@host> clear lldp neighbors user@host> clear lldp neighbors interface ge-0/2/2 |

show lldp

| | |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show lldp <detail> |
| Release Information | Command introduced in JUNOS Release 9.6. |
| Description | On MX 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 259 show lldp detail on page 259 |
| Output Fields | Table 82 on page 258 describes the output fields for the show lldp command. Output fields are listed in the approximate order in which they appear. |

Table 82: 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 | List of LLDP interfaces, showing status (Enabled or Disabled) and Neighbor count (detail only). |
| LLDP basic TLVs supported | List of basic LLDP TLVs supported by this device (detail only). |
| LLDP 802 TLVs supported | List of IEEE 802.1 LLDP TLVs supported by this device (detail only). |


```

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

```

```

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 Release 9.6. |
| Description | On MX 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 261 |
| Output Fields | Table 83 on page 260 describes the output fields for the show lldp local-information command. Output fields are listed in the approximate order in which they appear. |

Table 83: 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. |
| Interface ID | List of local interface identifiers. |
| Interface Description | List of local interface descriptions. |
| Status | List of interface conditions: UP or DOWN. |

```

show lldp      user@host> show lldp local-information
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 | show lldp neighbors <interface <i>interface-name</i> > |
| Release Information | Command introduced in JUNOS Release 9.6. |
| Description | On MX Series routers, display information about LLDP neighbors. |
| Options | interface <i>interface-name</i> —(Optional) Display the neighbor information about a particular physical interface. |
| Required Privilege Level | view |
| Related Topics | clear lldp neighbor |
| List of Sample Output | show lldp neighbors on page 264 show lldp neighbors interface ge-0/0/4 on page 264 |
| Output Fields | Table 84 on page 262 describes the output fields for the show lldp neighbors command. Output fields are listed in the approximate order in which they appear. |

Table 84: show lldp neighbors Output Fields

| Field Name | Field Description |
|------------------------------------|------------------------------------------------------------------------------------------------------------------------|
| LLDP Remote Devices Information | Information about remote devices. |
| LocalInterface | List of local interfaces for which neighbor information is available. |
| ChassisId | List of chassis identifiers for neighbors. |
| PortInfo | List of port information gathered from neighbors. This could be the port identifier or port description. |
| SysName | List of system names gathered from neighbors. |
| LLDP Neighbor Information | Information about both local and neighbor systems on the interface (appears when the interface option is used). |
| Local Information | Information about local systems on the interface (appears when the interface option is used). |
| 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). |

Table 84: show lldp neighbors Output Fields *(continued)*

| Field Name | Field Description |
|----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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). |

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 |

show lldp neighbors interface ge-0/0/4 user@host> **show lldp neighbors interface ge-0/0/4**
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 Release 9.6. |
| Description | On MX 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 266 |
| Output Fields | Table 85 on page 265 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 85: 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. |

```
show lldp user@host> show lldp remote-global-statistics  
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 Release 9.6. |
| Description | On MX 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 Topics | clear lldp statistics |
| List of Sample Output | show lldp statistics on page 268 show lldp statistics interface ge-0/1/1 on page 268 |
| Output Fields | Table 84 on page 262 describes the output fields for the show lldp statistics command. Output fields are listed in the approximate order in which they appear. |

Table 86: show lldp statistics Output Fields

| Field Name | Field Description |
|--------------|-------------------------------------------------------------------------------------------|
| Interface | Interface name. |
| 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. |

```

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

```

```

show lldp statistics   user@host> show lldp statistics interface ge-0/1/1
interface ge-0/1/1 Interface Received Transmitted Unknown-TLVs With-Errors Discarded
-----
ge-0/1/1  544      540          0           0           0

```

Chapter 9

OSPF Operational Mode Commands

Table 87 on page 269 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 87: 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 overload</code> |
| Clear OSPF statistics. | <code>clear (ospf ospf3) statistics</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> |
| 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” on page 321. For information about monitoring Bidirectional Forwarding Detection (BFD) sessions for OSPF clients, see “BFD Operational Mode Commands” on page 15. For information about how to configure OSPF, see the *JUNOS Routing Protocols Configuration Guide*.

clear (ospf | ospf3) database

Syntax clear (ospf | ospf3) database
 <advertising-router (*router-id* | self)>
 <area *area-id*>
 <asbrsummary>
 <external>
 <instance *instance-name*>
 <inter-area-prefix>
 <inter-area-router>
 <intra-area-prefix>
 <link-local>
 <logical-system (all | *logical-system-name*)>
 <lsa-id *lsa-id*>
 <netsummary >
 <network>
 <nssa>
 <opaque-area>
 <purge>
 <realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>
 <router>

Release Information Command introduced before JUNOS Release 7.4.
 advertising-router *router-id*, area *area-id*, asbrsummary, external, inter-area-prefix, inter-area-router, intra-area-prefix, link-local, lsa-id *lsa-id*, netsummary, network, nssa, opaque-area, and router options added in JUNOS Release 8.3. You must use the **purge** command with these options.
 realm option added in JUNOS Release 9.2.
 advertising-router (*router-id* | self) option added in JUNOS Release 9.5.

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 **purge** 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.



NOTE: You must use the **purge** command with any of the options added in JUNOS Release 8.3. You can use advertising-router *router-id*, area *area-id*, instance *instance-name*, logical-system *logical-system-name*, and lsa-id *lsa-id* with the options to further define the LSAs you want to discard.

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 router or by this router.

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—(OSPF3 only) (Optional) Discard interarea prefix LSAs.

inter-area-router—(OSPF3 only) (Optional) Discard interarea router LSAs.

intra-area-prefix—(OSPF3 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*)—(ospf3 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 Topics show ospf database
show ospf3 database

List of Sample Output clear ospf database on page 273

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

clear ospf database user@host> **clear ospf database**

clear (ospf | ospf3) io-statistics

| | |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | clear (ospf ospf3) statistics <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Clear Open Shortest Path First (OSPF) input and output statistics. |
| Options | none—Clear OSPF input and output statistics on all logical systems. 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 274 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| clear ospf io-statistics | user@host> clear ospf io-statistics |

clear (ospf | ospf3) neighbor

| | |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | clear (ospf ospf3) neighbor <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> <neighbor> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)> |
| Release Information | Command introduced before JUNOS Release 7.4. realm option introduced in JUNOS Release 9.2. |
| Description | Tear down Open Shortest Path First (OSPF) neighbor connections. |
| Options | <p>none—Tear down OSPF connections with all neighbors for all routing instances on all logical systems.</p> <p>instance <i>instance-name</i>—(Optional) Tear down neighbor connections 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 the state of the specified neighbor only.</p> <p>realm (ipv4-multicast ipv4-unicast ipv6-multicast)—(Optional) (OSPF3 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.</p> |
| Required Privilege Level | clear |
| Related Topics | show (ospf ospf3) neighbor |
| List of Sample Output | clear ospf neighbor on page 275 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| clear ospf neighbor | user@host> clear ospf neighbor |

clear ospf overload

| | |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | clear ospf overload <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Clear the Open Shortest Path First version 2 (OSPFv2) overload bit and rebuild link-state advertisements (LSAs). |
| Options | <p>none—Clear the overload bit and rebuild LSAs for all routing instances on all logical systems.</p> <p>instance <i>instance-name</i>—(Optional) Clear the overload bit and rebuild LSAs 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 |
| List of Sample Output | clear ospf overload on page 276 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| clear ospf overload | user@host> clear ospf overload |

clear (ospf | ospf3) statistics

| | |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | clear (ospf ospf3) statistics <logical-system (all <i>logical-system-name</i>)> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)> |
| Release Information | Command introduced before JUNOS Release 7.4. realm option introduced in JUNOS Release 9.2. |
| Description | Clear Open Shortest Path First (OSPF) statistics. |
| Options | <p>none—Clear OSPF statistics on all logical systems.</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) (OSPF3 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 Topics | show (ospf ospf3) statistics |
| List of Sample Output | clear ospf statistics on page 277 |
| Output Fields | See show (ospf ospf3) statistics for an explanation of output fields. |

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 |

```
LSAs retransmitted: 3, 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 |

LSAs retransmitted: 0, last 5 seconds: 0

Flood queue depth: 0

Total rexmit entries: 0, db summaries: 0, lsreq entries: 0

Receive errors:

None

show ospf database

Syntax show ospf database
 <brief | detail | extensive | summary>
 <advertising-router (*address* | self)>
 <area *area-id*>
 <asbrsummary>
 <external>
 <instance *instance-name*>
 <link-local>
 <logical-system (all | *logical-system-name*)>
 <lsa-id *lsa-id*>
 <netsummary>
 <network>
 <nssa>
 <opaque-area>
 <router>

Release Information Command introduced before JUNOS Release 7.4.
 advertising-router self (*address* | self) option introduced in JUNOS Release 9.5.

Description Display the entries in the Open Shortest Path First version 2 (OSPFv2) link-state database, which contains data about link-state advertisement (LSA) packets.

Options none—Display standard information about entries in the OSPFv2 link-state database for all routing instances on all logical systems.

brief | detail | extensive | summary—(Optional) Display the specified level of output.

advertising-router (*address* | self)—(Optional) Display the LSAs advertised either by a particular router or by this router.

area *area-id*—(Optional) Display the LSAs in a particular area.

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 Topics clear (ospf | ospf3) database

List of Sample Output show ospf database on page 281
 show ospf database brief on page 282
 show ospf database detail on page 282
 show ospf database extensive on page 283
 show ospf database summary on page 285

Output Fields Table 88 on page 280 describes the output fields for the `show ospf database` command. Output fields are listed in the approximate order in which they appear.

Table 88: 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 router. | All levels |
| Adv Rtr | Address of the router that sent the advertisement. | All levels |
| Seq | Link sequence number of the advertisement. | All levels |
| Age | Time elapsed since the LSA was originated, in seconds. | 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 router that generated the LSP. ■ link count—Number of links in the advertisement. ■ id—ID of a router or subnet on the link. ■ data—For stub networks, the subnet mask. Otherwise, the IP address of the router that generated the LSP. ■ type —Type of link. It can be <code>PointToPoint</code>, <code>Transit</code>, <code>Stub</code>, or <code>Virtual</code>. ■ 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 |

Table 88: show ospf database Output Fields (continued)

| Field Name | Field Description | Level of Output |
|-------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| 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 time | 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 |
| 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 |

```

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    0x800000002    172    0x20 0x6d72    28
NSSA  *33.1.1.1           10.255.71.242    0x800000001    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 281.

```

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    0x800000002  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    0x800000002  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    0x800000002  218  0x20 0x6d72  28
  mask 255.255.255.0
  TOS 0x0, metric 1
Summary  *24.1.1.0          10.255.71.242    0x800000002  223  0x20 0x607e  28
  mask 255.255.255.0
  TOS 0x0, metric 1
NSSA     *33.1.1.1           10.255.71.242    0x800000002  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    0x800000004  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    0x800000003  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    0x800000002  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    0x800000001  263  0x20 0xfeec  28
  mask 255.255.255.0
  TOS 0x0, metric 1
Summary  *24.1.1.0          10.255.71.242    0x800000002  223  0x20 0x607e  28
  mask 255.255.255.0
  TOS 0x0, metric 1
NSSA     *33.1.1.1           10.255.71.242    0x800000001  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    0x800000003  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 | 0xfec | 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 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>

Release Information Command introduced before JUNOS Release 7.4.
 realm option introduced in JUNOS Release 9.2.
 advertising-router (*address* | self) option introduced in JUNOS Release 9.5.

Description Display the entries in the Open Shortest Path First version 3 (OSPFv3) link-state database, which contains data about link-state advertisement (LSA) packets.

Options none—Display standard information about all entries in the OSPFv3 link-state database on all logical systems.

brief | detail | extensive | summary—(Optional) Display the specified level of output.

advertising-router (*address* | self)—(Optional) Display the LSAs advertised either by a particular router or by this router.

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 Topics clear (ospf | ospf3) database

List of Sample Output show ospf3 database brief on page 291
show ospf3 database extensive on page 292
show ospf3 database summary on page 294

Output Fields Table 89 on page 287 lists the output fields for the **show ospf3 database** command. Output fields are listed in the approximate order in which they appear.

Table 89: 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 |
| 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 router. | brief detail extensive |
| Adv Rtr | Address of the router 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 |

Table 89: show ospf3 database Output Fields (continued)

| Field Name | Field Description | Level of Output |
|----------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------|
| 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 router 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 router interface depends on the interface's type: <ul style="list-style-type: none"> ■ PointToPoint (1)—Point-to-point connection to another router. ■ 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 router. | detail extensive |
| Nbr-if-id | Interface ID of the neighbor's interface for this router link. | detail extensive |
| Nbr-rtr-id | Router ID of the neighbor router (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 |
| 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 routers attached to the link. Only routers 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 |

Table 89: 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 |
| 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 router described by the LSA. | detail extensive |
| options | Optional capabilities supported by the router. | 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 |
| 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) | | |

Table 89: show ospf3 database Output Fields (continued)

| Field Name | Field Description | Level of Output |
|-----------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| 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 router 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 |
| 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 |

Table 89: show ospf3 database Output Fields (continued)

| Field Name | Field Description | Level of Output |
|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| 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 |
| <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 |

```

show ospf3 database user@host> show ospf3 database brief
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 | show (ospf ospf3) interface <brief detail extensive> <area <i>area-id</i> > <interface-name> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)> |
| Release Information | Command introduced before JUNOS Release 7.4. area option introduced in JUNOS Release 9.2. realm option introduced in JUNOS Release 9.2. |
| Description | Display the status of Open Shortest Path First (OSPF) interfaces. |
| Options | <p>none—Display standard information about the status of all OSPF interfaces for all routing instances on all logical systems</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>area <i>area-id</i>—(Optional) Display information about the interfaces that belong to the specified area.</p> <p><i>interface-name</i>—(Optional) Display information for the specified interface.</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) (OSPF3 only) 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.</p> |
| Required Privilege Level | view |
| List of Sample Output | <p>show ospf interface brief on page 298</p> <p>show ospf interface detail on page 298</p> <p>show ospf3 interface detail on page 298</p> <p>show ospf interface detail(When Multiarea Adjacency Is Configured) on page 299</p> <p>show ospf interface area <i>area-id</i> on page 300</p> <p>show ospf interface extensive (When Flooding Reduction Is Enabled) on page 300</p> |
| Output Fields | Table 90 on page 297 lists the output fields for the show (ospf ospf3) interface command. Output fields are listed in the approximate order in which they appear. |

Table 90: 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's maximum transmission unit (MTU). | detail extensive |
| Cost | Interface's 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 only for 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 |
| 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 |

Table 90: show (ospf | ospf3) interface Output Fields (continued)

| Field Name | Field Description | Level of Output |
|------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| Auth type | (OSPFv2) Authentication mechanism for sending and receiving OSPF protocol packets: <ul style="list-style-type: none"> ■ MD5—MD5 mechanism is configured in accordance with RFC 2328. ■ None—No authentication method is configured. ■ Password—Simple password (RFC 2328) is configured. | detail 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 router 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 |

```

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
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
1o0.0              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
1o0.0              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
so-0/0/0.0         PtToPt  0.0.0.0    0.0.0.0        0.0.0.0         1

  Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Topology default (ID 0) -> Cost: 1
so-0/0/0.0         PtToPt  0.0.0.0    0.0.0.0        0.0.0.0         0

  Type: P2P, Address: 192.168.37.46, Mask: 255.255.255.254, MTU: 4470, Cost: 1
  Adj count: 0, , Passive
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Topology default (ID 0) -> Passive, Cost: 1
so-1/0/0.0         PtToPt  0.0.0.0    0.0.0.0        0.0.0.0         1

  Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Topology default (ID 0) -> Cost: 1
so-1/0/0.0         PtToPt  0.0.0.0    0.0.0.0        0.0.0.0         0

  Type: P2P, Address: 192.168.37.54, Mask: 255.255.255.254, MTU: 4470, Cost: 1
  Adj count: 0, , Passive
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Topology default (ID 0) -> Passive, Cost: 1
so-0/0/0.0         PtToPt  1.1.1.1    0.0.0.0        0.0.0.0         1

  Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
  Adj count: 1, Secondary
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Topology default (ID 0) -> Cost: 1
so-1/0/0.0         PtToPt  1.1.1.1    0.0.0.0        0.0.0.0         1

  Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
  Adj count: 1, Secondary
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Topology default (ID 0) -> Cost: 1
so-0/0/0.0         PtToPt  2.2.2.2    0.0.0.0        0.0.0.0         1

```

```

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1, Secondary
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-1/0/0.0          PtToPt  2.2.2.2          0.0.0.0          0.0.0.0          1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1, Secondary
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1

```

```

show ospf interface area user@host> show ospf interface area 1.1.1.1
area-id
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 user@host> show ospf interface extensive
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

(When Flooding
Reduction Is Enabled)
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 | ospf3) io-statistics

| | |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show (ospf ospf3) io-statistics <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display Open Shortest Path First (OSPF) input and output statistics. |
| Options | none—Display OSPF input and output statistics on all logical systems. 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 Topics | clear (ospf ospf3) statistics |
| List of Sample Output | show ospf io-statistics on page 301 |
| Output Fields | Table 91 on page 301 lists the output fields for the show ospf io-statistics command. Output fields are listed in the approximate order in which they appear. |

Table 91: 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. |

```
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 | 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> > |
| Release Information | Command introduced before JUNOS Release 7.4. topology option introduced in JUNOS Release 9.0. realm option introduced in JUNOS Release 9.2. |
| Description | Display the entries in the Open Shortest Path First (OSPF) log of SPF calculations. |
| Options | none—Display entries in the OSPF log of SPF calculations for all routing instances on all logical systems. instance <i>instance-name</i> —(Optional) Display entries for the specified routing instance. logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system. topology <i>topology-name</i> —(Optional) Display entries for the specified topology. realm (ipv4-multicast ipv4-unicast ipv6-multicast)—(ospf3 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. |
| Required Privilege Level | view |
| List of Sample Output | show ospf log on page 303 show ospf log topology voice on page 303 |
| Output Fields | Table 92 on page 302 lists the output fields for the show (ospf ospf3) log command. Output fields are listed in the approximate order in which they appear. |

Table 92: show (ospf | ospf3) log Output Fields

| Field Name | Field Description |
|------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| When | Time, in weeks (w) and days (d), since the SPF calculation was made. |
| 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. |

```

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 show (ospf | ospf3) neighbor
 <brief | detail | extensive>
 <area *area-id*>
 <instance (all | *instance-name*)>
 <interface *interface-name*>
 <logical-system (all | *logical-system-name*)>
 <neighbor>
 <realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>

Release Information Command introduced before JUNOS Release 7.4.
 instance all option introduced in JUNOS Release 9.1.
 area option introduced in JUNOS Release 9.2.
 interface option introduced in JUNOS Release 9.2.
 realm option introduced in JUNOS Release 9.2.

Description Display information about Open Shortest Path First (OSPF) neighbors.

Options none—Display standard information about all OSPF neighbors for all routing instances on all logical systems.

brief | detail | extensive—(Optional) Display the specified level of output.

area *area-id*—(Optional) Display information about the OSPF neighbors for the specified area.

instance (all | *instance-name*)—(Optional) Display all OSPF interfaces for all routing instances or under the named routing instance.

interface *interface-name*—(Optional) Display information about OSPF neighbors for the specified logical interface.

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

neighbor—(Optional) Display information about the specified OSPF neighbor.

realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)—(Optional) (OSPF3 only) 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 Topics clear (ospf | ospf3) neighbor

List of Sample Output show ospf neighbor brief on page 306
 show ospf neighbor detail on page 306
 show ospf neighbor extensive on page 307
 show ospf3 neighbor detail on page 308
 show ospf neighbor area *area-id* on page 308

show ospf neighbor interface interface-name on page 308
 show ospf3 neighbor instance all (OSPFv3 Multiple Family Address Support Enabled) on page 308

Output Fields Table 93 on page 305 lists the output fields for the show (ospf | ospf3) neighbor command. Output fields are listed in the approximate order in which they appear.

Table 93: 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 |
| State | State of the neighbor: <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—Router 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 routers. The goal of this step is to determine which router is the master, and to determine the initial sequence number. ■ Full—Neighboring routers 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 may occur, for example, because the router 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 routers 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 |

Table 93: show (ospf | ospf3) neighbor Output Fields (continued)

| Field Name | Field Description | Level of Output |
|--------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| Link state retransmission list | Total number of link-state advertisements retransmitted. For extensive output only, the following information is also displayed: <ul style="list-style-type: none"> ■ Type—Type of link advertisement: ASBR, Sum, Extern, Network, NSSA, OpaqArea, Router, or Summary. ■ LSA ID—LSA identifier included in the advertisement. An asterisk preceding the identifier marks database entries that originated from the local router. ■ Adv rtr—Address of the router that sent the advertisement. ■ Seq—Link sequence number of the advertisement. | detail extensive |
| 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 |

```

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, retransmit DBD in 2 sec
  retransmit 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 <extensive> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> <neighbor> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)> |
| Release Information | Command introduced in JUNOS Release 7.4. realm statement introduced in JUNOS Release 9.2. |
| Description | Display Open Shortest Path First (OSPF) overview information. |
| Options | <p>none—Display standard information about all OSPF neighbors for all routing instances on all logical systems.</p> <p>extensive—(Optional) Display trace information in addition to standard information.</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>neighbor—(Optional) Display information about the specified OSPF neighbor.</p> <p>realm (ipv4-multicast ipv4-unicast ipv6-multicast)—(Optional) (OSPF3 only) Display information about 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 | show ospf overview on page 310 show ospf overview extensive on page 312 |
| Output Fields | Table 94 on page 309 lists the output fields for the show ospf overview command. Output fields are listed in the approximate order in which they appear. |

Table 94: show ospf overview Output Fields

| Field name | Field Description | Level of Output |
|-------------------|----------------------------|-----------------|
| Instance | The OSPF routing instance. | All levels |
| Router ID | Router ID of the router. | All levels |
| Route table index | Route table index. | All levels |

Table 94: show ospf overview Output Fields (continued)

| Field name | Field Description | Level of Output |
|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| 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 |
| 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 holddown timer begins. | |
| LSA refresh time | Refresh period for link-state advertisement. (in minutes) | 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 router as part of the topology. | All levels |
| Helper mode | 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 . | All levels |
| Area border routers | Number of area border routers. | All levels |
| Neighbors | Number of autonomous system boundary routers. | All levels |

```

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 extensive user@hostshow 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 show (ospf | ospf3) route
 <detail>
 <abr | asbr | extern | inter | intra>
 <instance *instance-name*>
 <logical-system (all | *logical-system-name*)>
 <realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>
 <topology *topology-name*>
 <transit>

Release Information Command introduced before JUNOS Release 7.4.
 topology option introduced in JUNOS Release 9.0.
 realm option introduced in JUNOS Release 9.2.

Description Display the entries in the Open Shortest Path First (OSPF) routing table.

Options none—Display standard information about all entries in the OSPF routing table for all routing instances on all logical systems and all topologies.

detail—(Optional) Display detailed information.

abr—(Optional) Display routes to area border routers.

asbr—(Optional) Display routes to autonomous system border routers.

extern—(Optional) Display external routes.

inter—(Optional) Display interarea routes.

intra—(Optional) Display intra-area routes.

instance *instance-name*—(Optional) Display entries for the specified routing instance.

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

realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)—(ospf3 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.

topology *topology-name*—(ospf only) (Optional) Display routes for a particular topology.

transit—(Optional) (OSPF3 only) Display OSPFv3 routes to pseudonodes.

Required Privilege Level view

List of Sample Output show ospf route on page 315
 show ospf route detail on page 315
 show ospf3 route on page 315
 show ospf3 route detail on page 316

show ospf route topology voice on page 317

Output Fields Table 95 on page 314 list the output fields for the `show (ospf | ospf3) route` command. Output fields are listed in the approximate order in which they appear.

Table 95: 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 |
| Route type | The type of router 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 |

Table 95: show (ospf | ospf3) route Output Fields (continued)

| Field Name | Field Description | Output Level |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------|
| optional-capability | Optional capabilities propagated in the router LSA. This field is in the output for router routes only (when Route Type is Area BR, AS BR, Area/AS BR, or Router), not for network routes. Three bits in this field are defined as follows: <ul style="list-style-type: none"> ■ 0x4 (V)—Router is at the end of a virtual active link. ■ 0x2 (E)—Router is an autonomous system boundary router. ■ 0x1 (B)—Router is an area border router. | detail |
| 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 |

```

show ospf route      user@host> show ospf route
Prefix                Path  Route  NH  Metric  NextHop  Nexthop
                    Type  Type  Type  Type      Interface
addr/label
10.255.71.12          Intra Router  IP    1      fe-0/0/2.0  192.16.22.86
10.255.71.13/32      Intra Network IP    0      1o0.0
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  Type      Interface  addr/
label
10.255.14.174          Inter AS BR  IP    210    t1-3/0/1.0
    area 0.0.0.2, origin 10.255.14.185
10.255.14.178          Intra Router  IP    200    t3-3/1/3.0
    area 0.0.0.2, origin 10.255.14.178, optional-capability 0x0
10.210.1.0/30          Intra Network IP    10     t3-3/1/2.0
    area 0.0.0.2, origin 10.255.14.172, priority medium
100.1.1.1/32           Inter Network IP    210    t1-3/0/1.0
    area 0.0.0.2, origin 10.255.14.185, priority low
112.3.1.0/24           Ext2  Network IP    0     t1-3/0/1.0
    area 0.0.0.0, origin 10.255.14.174, priority high
200.3.3.0/30           Inter Network IP    220    t1-3/0/1.0
    area 0.0.0.2, origin 10.255.14.185, priority high

```

```

show ospf3 route     user@host> show ospf3 route
Prefix                Path  Route  NH  Metric
                    type  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  Prefix
Metric  NextHop      Nexthop
                    Type  Type  Type  Interface  addr/label
10.255.245.1          Intra Router  IP    40     fxp1.1      192.168.36.17

```

```

    area 0.0.0.0, origin 10.255.245.1 optional-capability 0x0,
10.255.245.3      Intra AS BR      IP      1 fxp2.3      192.168.36.34
    area 0.0.0.0, origin 10.255.245.3 optional-capability 0x0,
10.255.245.1/32   Intra Network   IP      40 fxp1.1      192.168.36.17
    area 0.0.0.0, origin 10.255.245.1, priority high
10.255.245.2/32   Intra Network   IP      0 lo0.0
    area 0.0.0.0, origin 10.255.245.2, priority medium
10.255.245.3/32   Intra Network   IP      1 fxp2.3      192.168.36.34
    area 0.0.0.0, origin 10.255.245.3, priority low

                Intra Transit      IP      1
    NH-interface fe-0/0/2.0
192::168:222:84/126                Intra Network   IP      1
    NH-interface fe-0/0/2.0
abcd::71:12/128                    Intra Network   IP      0
    NH-interface lo0.0
abcd::71:13/128                    Intra Network   LSP      1
    NH-interface fe-0/0/2.0, NH-addr lsp-cd

```

show ospf3 route detailuser@host> **show ospf3 route detail**

| Prefix | Path | Route | NH |
|-------------------------------------------------------------|-------|------------|--------|
| Metric | | | |
| 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 | | | |

```

user@host show ospf route topology voice
show ospf route topology voice
Topology voice Route Table:
Prefix          Path    Route    NH    Metric  NextHop    Nexthop
                  Type    Type     Type                   Interface  addr/label
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)> |
| Release Information | Command introduced before JUNOS Release 7.4. realm statement introduced in JUNOS Release 9.2. |
| Description | Display OSPF statistics. |
| Options | <p>none—Display OSPF statistics for all routing instances on all logical systems.</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) (OSPF3 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 Topics | clear (ospf ospf3) statistics |
| List of Sample Output | show ospf statistics on page 319 |
| Output Fields | Table 96 on page 318 lists the output fields for the show (ospf ospf3) statistics command. Output fields are listed in the approximate order in which they appear. |

Table 96: 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. |
| LSAs retransmitted | Total number of link-state advertisements transmitted, and number retransmitted in the last 5 seconds. |
| Receive errors | Number and type of receive errors. |

show ospf statisticsuser@host> **show ospf statistics**

| Packet type | Total | | Last 5 seconds | |
|-------------|--------|----------|----------------|----------|
| | Sent | Received | Sent | Received |
| Hello | 505739 | 990495 | 4 | 5 |
| DbD | 20 | 26 | 0 | 0 |
| LSReq | 6 | 5 | 0 | 0 |
| LSUpdate | 27060 | 15319 | 0 | 0 |
| LSAck | 10923 | 52470 | 0 | 0 |

LSAs retransmitted: 16, last 5 seconds: 0

Receive errors:

862 no interface found

115923 no virtual link found

Chapter 10

Protocol-Independent Routing Operational Mode Commands

Table 97 on page 321 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 97: 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> |
| Display routes containing a specified AS path. | <code>show route aspath-regex</code> |

Table 97: Protocol-Independent Routing Operational Mode Commands *(continued)*

| Task | Command |
|-------------------------------------------------------------------------------------------------------|---------------------------------------------|
| 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 forwarding table. | <code>show route forwarding-table</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> |
| 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 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> |

Table 97: Protocol-Independent Routing Operational Mode Commands *(continued)*

| Task | Command |
|--------------------------------------------------------------------|------------------------------------------|
| 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>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display the distribution of autonomous system (AS) paths that the local router 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. |
| Options | <p>none—Display basic information about AS paths that the local router 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 |
| List of Sample Output | <p>show as-path on page 325</p> <p>show as-path detail on page 326</p> |
| Output Fields | Table 98 on page 324 lists the output fields for the show as-path command. Output fields are listed in the approximate order in which they appear. |

Table 98: show as-path Output Fields

| Field Name | Field Description | Level of Output |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Total AS paths | Total number of AS paths. | brief none |
| Bucket | Bucket value. This value represents a traffic classification on the interface. | All levels |
| Count | Path reference count. | 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 was 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—Router 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 |

Table 98: show as-path Output Fields (continued)

| Field Name | Field Description | Level of Output |
|-------------|--------------------------------------|-----------------|
| 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 |

```

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 9232 I
    14203 2914 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>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display autonomous system (AS) path domain information. |
| Options | <p>none—(Optional) Display AS path domain information for all routing instances on all logical systems.</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 as-path domain on page 329 |
| Output Fields | Table 99 on page 328 lists the output fields for the show as-path domain command. Output fields are listed in the approximate order in which they appear |

Table 99: 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 router. |
| Loops | How many times this AS number can appear in an AS path. |

```
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>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display autonomous system (AS) path summary information. |
| Options | <p>none—(Optional) Display AS path summary information for all routing instances on all logical systems.</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 as-path summary on page 330 |
| Output Fields | Table 100 on page 330 lists the output fields for the show as-path summary command. Output fields are listed in the approximate order in which they appear. |

Table 100: show as-path summary Output Fields

| Field Name | Field Description |
|---------------|--------------------------------------------------------------------------------|
| AS Path | AS path number. |
| Buckets | Bucket value. This value represents a traffic classification on the interface. |
| Max | Maximum limit for the number of AS numbers. |
| Min | Minimum limit for the number of AS numbers. |
| Avg | Average amount of AS numbers. |
| Std deviation | Standard deviation for the number of AS numbers. |

```

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 | show route <all> <destination-prefix> <logical-system (all <i>logical-system-name</i>)> <private> |
| Release Information | Command introduced before JUNOS Release 7.4. private option introduced in JUNOS Release 9.5. |
| Description | Display the active entries in the routing tables. |
| Options | <p>none—Display brief information about all active entries in the routing tables on all logical systems.</p> <p>all—(Optional) Display information about all routing tables, including private, or internal, routing tables.</p> <p><i>destination-prefix</i>—(Optional) Display active entries for the specified address or range of addresses.</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>private—(Optional) Display information only about all private, or internal, routing tables.</p> |
| Required Privilege Level | view |
| List of Sample Output | show route on page 334 show route destination-prefix on page 334 |
| Output Fields | Table 101 on page 331 describes the output fields for the <code>show route</code> command. Output fields are listed in the approximate order in which they appear. |

Table 101: 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> | 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 101: show route Output Fields (continued)

| Field Name | Field Description |
|---------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <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"> ■ <i>MPLS-label</i> (for example, 80001). ■ <i>interface-name</i> (for example, ge-1/0/2). ■ <i>neighbor-address:control-word-status:encapsulation type:vc-id :source</i> (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96): <ul style="list-style-type: none"> ■ <i>neighbor-address</i>—Address of the neighbor. ■ <i>control-word-status</i>—Whether the use of the control word has been negotiated for this virtual circuit: NoCtrlWord or CtrlWord. ■ <i>encapsulation type</i>— 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 ■ <i>vc-id</i>—Virtual circuit identifier. ■ <i>source</i>—Source of the advertisement: Local or Remote. |
| <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 <i>to</i> 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 Software 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> |
| <i>weeks:days 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). |
| <i>metric</i> | 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. |
| <i>localpref</i> | Local preference value included in the route. |
| <i>from</i> | Interface from which the route was received. |

Table 101: show route Output Fields (continued)

| Field Name | Field Description |
|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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 was originated:</p> <ul style="list-style-type: none"> ■ I—IIGP. ■ 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 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. |
| to | Next hop to the destination. An angle bracket (>) indicates that the route is the selected route. |
| 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 Multiprotocol Label Switching (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 router is performing unequal-cost load balancing. This information is available when you enable Border Gateway Protocol (BGP) multipath load balancing. ■ lsp-path-name—Name of the label-switched path (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). |

show route user@host> **show route**

```

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

red.inet.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.10.10.10/32  *[Direct/0] 01:08:46
                > via lo0.1
10.255.245.212/32 *[BGP/170] 00:01:40, 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
10.255.245.213/32 *[BGP/170] 00:40:47, localpref 100
                AS path: 100 I
                > to 100.1.1.1 via so-0/0/1.0

```

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 active-path

| | |
|-------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show route active-path <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced in JUNOS Release 8.0. |
| 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 | <p>none—Display all active 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 | <p>show route active-path on page 335</p> <p>show route active-path brief on page 335</p> <p>show route active-path detail on page 336</p> <p>show route active-path extensive on page 337</p> <p>show route active-path terse on page 338</p> |
| Output Fields | For information about output fields, see the <code>show route</code> command (Table 101 on page 331), the <code>show route detail</code> command (Table 104 on page 360), the <code>show route extensive</code> command (Table 110 on page 380), or the <code>show route terse</code> command (Table 118 on page 483). |
| show route active-path | <pre> 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 </pre> |
| show route active-path brief | The output for the <code>show route active-path brief</code> command is identical to that for the <code>show route active-path</code> command. For sample output, see <code>show route active-path</code> on page 335. |

```

show route active-path detail      user@host> show route active-path detail

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)

10.255.70.19/32 (1 entry, 1 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 3
    Next hop: via lo0.0, selected
    State: <Active Int>
    Local AS: 200
    Age: 21:37:10
    Task: IF
    Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3

    AS path: I

10.255.71.50/32 (1 entry, 1 announced)
  *IS-IS Preference: 15
    Level: 1
    Next hop type: Router, Next hop index: 397
    Next-hop reference count: 4
    Next hop: 100.1.2.1 via so-2/1/3.0, selected
    State: <Active Int>
    Local AS: 200
    Age: 21:31      Metric: 10
    Task: IS-IS
    Announcement bits (4): 0-KRT 2-IS-IS 5-Resolve tree 2 6-Resolve
tree 3
    AS path: I

100.1.2.0/24 (1 entry, 1 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 3
    Next hop: via so-2/1/3.0, selected
    State: <Active Int>
    Local AS: 200
    Age: 21:54
    Task: IF
    Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3

    AS path: I

100.1.2.2/32 (1 entry, 1 announced)
  *Local Preference: 0
    Next hop type: Local
    Next-hop reference count: 11
    Interface: so-2/1/3.0
    State: <Active NoReadvrt Int>
    Local AS: 200
    Age: 21:59
    Task: IF
    Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
    AS path: I

192.168.64.0/21 (1 entry, 1 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 3
    Next hop: via fxp0.0, selected

```

```

State: <Active Int>
Local AS: 200
Age: 21:37:10
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

```

```

192.168.70.19/32 (1 entry, 1 announced)
  *Local Preference: 0
    Next hop type: Local
    Next-hop reference count: 11
    Interface: fxp0.0
    State: <Active NoReadvrt Int>
    Local AS: 200
    Age: 21:37:10
    Task: IF
    Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
    AS path: I

```

**show route active-path
extensive**

```

user@host> show route active-path extensive

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
10.255.70.19/32 (1 entry, 1 announced)
TSI:
IS-IS level 1, LSP fragment 0
IS-IS level 2, LSP fragment 0
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 3
    Next hop: via lo0.0, selected
    State: <Active Int>
    Local AS: 200
    Age: 21:39:47
    Task: IF
    Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3
    AS path: I

10.255.71.50/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.255.71.50/32 -> {100.1.2.1}
IS-IS level 2, LSP fragment 0
  *IS-IS Preference: 15
    Level: 1
    Next hop type: Router, Next hop index: 397
    Next-hop reference count: 4
    Next hop: 100.1.2.1 via so-2/1/3.0, selected
    State: <Active Int>
    Local AS: 200
    Age: 24:08      Metric: 10
    Task: IS-IS
    Announcement bits (4): 0-KRT 2-IS-IS 5-Resolve tree 2 6-Resolve
tree 3
    AS path: I

100.1.2.0/24 (1 entry, 1 announced)
TSI:
IS-IS level 1, LSP fragment 0
IS-IS level 2, LSP fragment 0
  *Direct Preference: 0
    Next hop type: Interface

```

```

Next-hop reference count: 3
Next hop: via so-2/1/3.0, selected
State: <Active Int>
Local AS: 200
Age: 24:31
Task: IF
Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3

AS path: I

100.1.2.2/32 (1 entry, 1 announced)
*Local Preference: 0
Next hop type: Local
Next-hop reference count: 11
Interface: so-2/1/3.0
State: <Active NoReadvrt Int>
Local AS: 200
Age: 24:36
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

192.168.64.0/21 (1 entry, 1 announced)
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 3
Next hop: via fxp0.0, selected
State: <Active Int>
Local AS: 200
Age: 21:39:47
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

192.168.70.19/32 (1 entry, 1 announced)
*Local Preference: 0
Next hop type: Local
Next-hop reference count: 11
Interface: fxp0.0
State: <Active NoReadvrt Int>
Local AS: 200
Age: 21:39:47
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

```

show route active-path terse user@host> **show route active-path terse**

```
inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
```

| A | Destination | P | Prf | Metric 1 | Metric 2 | Next hop | AS path |
|---|------------------|---|-----|----------|----------|-------------|---------|
| * | 10.255.70.19/32 | D | 0 | | | >100.0 | |
| * | 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 | show route advertising-protocol <i>protocol neighbor-address</i> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS 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><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 <p><i>neighbor-address</i>—Address of the neighboring router to which the route entry is being transmitted.</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 | 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. For more information, see the <i>JUNOS Routing Protocols Configuration Guide</i> . |
| Required Privilege Level | view |
| List of Sample Output | show route advertising-protocol bgp (Layer 3 VPN) on page 341 show route advertising-protocol bgp detail on page 342 show route advertising-protocol bgp detail (Layer 2 VPN) on page 342 show route advertising-protocol bgp detail (Layer 3 VPN) on page 342 |
| Output Fields | Table 102 on page 339 lists the output fields for the show route advertising-protocol command. Output fields are listed in the approximate order in which they appear. |

Table 102: 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 |

Table 102: show route advertising-protocol Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <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 (the routes are not used because of a routing policy) | All levels |
| Prefix | Destination prefix. | brief none |
| <i>destination-prefix</i> (entry , announced) | Destination prefix. The entry value is the number of routes for this destination, and the announced value is the number of routes being announced for this destination. | detail extensive |
| 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 Label Distribution Protocol (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 a Resource Reservation Protocol (RSVP) or a Label Distribution Protocol (LDP) label-switched path (LSP) tunnel. | detail extensive |
| Nexthop | Next hop to the destination. An angle bracket (>) indicates that the route is the selected route. | 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 102: 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 was originated:</p> <ul style="list-style-type: none"> ■ I—IIGP. ■ 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. | All levels |
| Communities | Community path attribute for the route. See Table 107 on page 367 for all possible values for this field. | 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 |

```

show route      user@host> show route advertising-protocol bgp 10.255.14.171
advertising-protocol bgp
(Layer 3 VPN)    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      user@host> show route advertising-protocol bgp 111.222.1.3 detail
advertising-protocol bgp
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
                  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

show route      user@host> show route advertising-protocol bgp 192.168.24.1 detail
advertising-protocol bgp
detail (Layer 2 VPN)
                  vpn-a.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
                  192.168.16.1:1:1/96 (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
                      Layer2-info: encaps:VLAN, control flags:, mtu:

show route      user@host> show route advertising-protocol bgp 10.255.14.176 detail
advertising-protocol bgp
detail (Layer 3 VPN)
                  vpna.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
                  * 10.49.0.0/30 (1 entry, 1 announced)
                    BGP group ibgp type Internal
                      Route Distinguisher: 10.255.14.174:2
                      VPN Label: 101264
                      Nexthop: Self
                      Localpref: 100
                      AS path: I
                      Communities: target:200:100
                      AttrSet AS: 100
                        Localpref: 100
                        AS path: I
                  ...

```

show route all

| | |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show route all <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display information about all routes in all routing tables, including private, or internal, tables. |
| 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 route all on page 343 |
| Output Fields | In JUNOS Release 9.5 and later, only the output fields for the show route all displays all routing tables, including private, or hidden, routing tables. The show route command output fields (see Table 101 on page 331) do not display entries for private, or hidden routing tables in JUNOS Release 9.5 and later. |

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 | show route aspath-regex <i>regular-expression</i> <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| 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>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 | <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. ■ *—Zero or more repetitions of an AS path term. ■ +—One or more repetitions of an AS path term. ■ ?—Zero or one repetition of an AS path term. ■ <i>aspath_term</i> <i>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 | <p>show route aspath-regex (Matching a Specific AS Number) on page 345</p> <p>show route aspath-regex (Matching Any Path with Two AS Numbers) on page 345</p> |
| Output Fields | For information about output fields, see the <code>show route</code> command (Table 101 on page 331). |

**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 | show route best <i>destination-prefix</i> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| 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 | <p><i>destination-prefix</i>—Address or range of addresses.</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 | <p>show route best on page 346</p> <p>show route best detail on page 347</p> <p>show route best extensive on page 347</p> <p>show route best terse on page 348</p> |
| Output Fields | For information about output fields, see the <code>show route</code> command (Table 101 on page 331), the <code>show route detail</code> command (Table 104 on page 360), the <code>show route extensive</code> command (Table 110 on page 380), or the <code>show route terse</code> command (Table 118 on page 483). |
| show route best | <pre> 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 </pre> |


```

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 347.

```

show route best terse  user@host> show route best 10.255.70.103 terse
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1   Metric 2   Next hop      AS path
* 10.255.70.103/32  0 10      2           >10.31.1.6
                                     so-0/3/0.0

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1   Metric 2   Next hop      AS path
* 10.255.70.103/32  R  7      2           >so-0/3/0.0

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1   Metric 2   Next hop      AS path
* 10.0.0.0/8        D  0           >fxp2.0
                    D  0           >fxp1.0

```

show route brief

| | |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show route brief <destination-prefix> <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display brief information about the active entries in the routing tables. |
| Options | <p>none—Display all active entries in the routing table on all logical systems.</p> <p><i>destination-prefix</i>—(Optional) Display active entries for the specified address or range of addresses.</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 | view |
| List of Sample Output | show route brief on page 349 |
| Output Fields | For information about output fields, see the <code>show route</code> command (Table 101 on page 331). |
| show route brief | <pre> 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 </pre> |

```
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 Release 7.4. |
| Description | Display circuit cross-connect (CCC) entries in the Multiprotocol Link Switching (MPLS) routing table. |
| Options | ccc—Name of an entry with a circuit cross-connect interface. 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 |
| Related Topics | show connections |
| List of Sample Output | show route ccc extensive on page 351 |
| Output Fields | For information about output fields, see the show route command (Table 101 on page 331), the show route detail command (Table 104 on page 360), the show route extensive command (Table 110 on page 380), or the show route terse command (Table 118 on page 483). |
| show route ccc extensive | <pre> 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 </pre> |

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>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| 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 Topics | show route detail |
| List of Sample Output | show route community on page 352 |
| Output Fields | For information about output fields, see the show route command (Table 101 on page 331), the show route detail command (Table 104 on page 360), the show route extensive command (Table 110 on page 380), or the show route terse command (Table 118 on page 483). |
| show route community | <pre> 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) </pre> |

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>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| 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 | <p><i>community-name</i>—Name of the community.</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 |
| List of Sample Output | show route community-name on page 353 |
| Output Fields | For information about output fields, see the show route command (Table 101 on page 331), the show route detail command (Table 104 on page 360), the show route extensive command (Table 110 on page 380), or the show route terse command (Table 118 on page 483). |
| show route community-name | <pre> 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) </pre> |

```

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>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display the Border Gateway Protocol (BGP) routes for which updates might have been reduced because of route flap damping. |
| Options | <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>suppressed—Display entries that have been suppressed and are no longer being installed into the forwarding table or exported by routing protocols.</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 |
| Related Topics | clear bgp damping show policy damping |
| List of Sample Output | show route damping decayed detail on page 358 show route damping history on page 358 show route damping history detail on page 358 |
| Output Fields | Table 103 on page 355 lists the output fields for the show route damping command. Output fields are listed in the approximate order in which they appear. |

Table 103: show route damping 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 |
| destinations | 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 ■ 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 |

Table 103: show route damping Output Fields (continued)

| Field Name | Field Description | Level of Output |
|-------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <i>destination-prefix</i> (entry, announced) | Destination prefix. The entry value is the number of routes for this destination, and the announced value is the number of routes being announced for this destination. | detail extensive |
| [<i>protocol,</i> <i>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 Software 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 router 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 Table 106 on page 365. | detail extensive |
| Local AS | AS number of the local router. | detail extensive |
| Peer AS | AS number of the peer router. | detail extensive |
| 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 <i>n</i> is an index used by Juniper Networks customer support only. | detail extensive |

Table 103: show route damping 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 was originated:</p> <ul style="list-style-type: none"> ■ I—IIGP. ■ 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 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. | 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 Table 107 on page 367 for all possible values for this field. | 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 |
| 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 |

Table 103: show route damping Output Fields (continued)

| Field Name | Field Description | Level of Output |
|--------------------|-------------------------------------------------------------------------------------------------------|-----------------|
| Preference will be | (suppressed keyword only) Preference value that will be applied to the route when it is again active. | All levels |

```

show route damping      user@host> show route damping decayed detail
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 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      user@host> show route damping history
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      user@host> show route damping history detail
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 <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| 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><i>destination-prefix</i>—(Optional) Display active entries for the specified address or range of addresses.</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 | view |
| List of Sample Output | show route detail on page 368 |
| Output Fields | Table 104 on page 360 describes the output fields for the show route detail command. Output fields are listed in the approximate order in which they appear. |

Table 104: 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> | <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) ■ hidden (routes that are not used because of a routing policy) |

Table 104: 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"> ■ <i>MPLS-label</i> (for example, 80001). ■ <i>interface-name</i> (for example, ge-1/0/2). ■ <i>neighbor-address:control-word-status:encapsulation type:vc-id:source</i> (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96). ■ <i>neighbor-address</i>—Address of the neighbor. ■ <i>control-word-status</i>—Whether the use of the control word has been negotiated for this virtual circuit: NoCtrlWord or CtrlWord. ■ <i>encapsulation type</i>—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 ■ <i>vc-id</i>—Virtual circuit identifier. ■ <i>source</i>—Source of the advertisement: Local or Remote. |
| label stacking | <p>(Next-to-the-last-hop router for MPLS only) Depth of the Multiprotocol Label Switching (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 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 Software 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 105 on page 364. |

Table 104: show route detail Output Fields (continued)

| Field Name | Field Description |
|---------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Next-hop reference count | Number of references made to the next hop. |
| 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 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 Multiprotocol Label Switching (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 router is performing unequal-cost load balancing. This information is available when you enable Border Gateway Protocol (BGP) multipath load balancing. |
| Label-switched-path <i>lsp-path-name</i> | Name of the label-switched path (LSP) used to reach the next hop. |
| Label operation | MPLS label and operation occurring at this router. 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 router 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 106 on page 365. |
| Local AS | AS number of the local router. |
| Age | How long the route has been known. |
| 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. |
| 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. |

Table 104: show route detail Output Fields (continued)

| Field Name | Field Description |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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 was 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. |
| 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 107 on page 367 for all possible values for this field. |
| Layer2-info: encaps | Layer 2 encapsulation (for example, VPLS). |
| control flags | Control flags: none or Site Down. |
| mtu | Maximum transmission unit (MTU) information. |
| Label-Base, range | First label in a block of labels and label block size. A remote PE router uses this first label when sending traffic toward the advertising PE router. |
| 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. |

Table 105 on page 364 describes all possible values for the **Next-hop Types** output field.

Table 105: Next-Hop Types Output Field Values

| Next-Hop Type | Description |
|--------------------------|---------------------------------------------------------------------------------------------|
| broadcast (bcast) | Broadcast next hop. |
| deny | Deny next hop. |
| 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 (mcr) | 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 (recv) | Receive. |
| reject (rjct) | Discard. An ICMP unreachable message was sent. |
| resolve (rslv) | Resolving 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 106 on page 365 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 106: 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. |
| Clone | Route is a clone. |
| Cisco Non-deterministic MED selection | Cisco nondeterministic MED is enabled and a path with a lower MED is available. |
| 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. |
| Local Preference | Path with a higher local preference value 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. |
| Martian | Route is a martian (ignored because it is obviously invalid). |
| MartianOK | Route exempt from martian filtering. |

Table 106: State Output Field Values *(continued)*

| Value | Description |
|--------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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 greater number of next hops is available. |
| Origin | Path with 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 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 107 on page 367 describes the possible values for the **Communities** output field.

Table 107: 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 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>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 <i>area-number:ospf-route-type:options</i> . |
| <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 <i>area-number:ospf-route-type:options</i> . |
| <i>target</i> | Defines which VPN the route participates in; target has the format <i>32-bit IP address:16-bit number</i> . 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. |

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

...

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
    Task: BGP_69.10.255.70.103+179
    Announcement bits (1): 0-green-l2vpn
    AS path: I
    Communities: target:11111:1 Layer2-info: encaps:VPLS,
    control flags:, mtu: 0
    Label-base: 800008, range: 8
    Localpref: 100
    Router ID: 10.255.70.103
    Primary Routing Table bgp.l2vpn.0

10.255.71.52:1:1:1/96 (1 entry, 1 announced)
  *L2VPN Preference: 170/-1
    Next-hop reference count: 5
    Protocol next hop: 10.255.71.52
    Indirect next hop: 0 -
    State: <Active Int Ext>
    Age: 1:31:40 Metric2: 1
    Task: green-l2vpn
    Announcement bits (1): 1-BGP.0.0.0.0+179
    AS path: I
    Communities: Layer2-info: encaps:VPLS, control flags:Site-Down,
    mtu: 0
    Label-base: 800016, range: 8, status-vector: 0x9F

10.255.71.52:1:5:1/96 (1 entry, 1 announced)
  *L2VPN Preference: 170/-101
    Next-hop reference count: 5
    Protocol next hop: 10.255.71.52
    Indirect next hop: 0 -
    State: <Active Int Ext>
    Age: 1:31:40 Metric2: 1
    Task: green-l2vpn
    Announcement bits (1): 1-BGP.0.0.0.0+179
    AS path: I
    Communities: Layer2-info: encaps:VPLS, control flags:, mtu: 0
    Label-base: 800008, range: 8, status-vector: 0x9F

...

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
  *L2CKT Preference: 7
    Next hop: via so-1/1/2.0 weight 1, selected

```

```
Label-switched-path my-lsp
Label operation: Push 100000[0]
Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
State: <Active Int>
Local AS: 99
Age: 10:21
Task: l2 circuit
Announcement bits (1): 0-LDP
AS path: I
VC Label 100000, MTU 1500, VLAN ID 512
```

show route exact

| | |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show route exact <i>destination-prefix</i> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display only the routes that exactly match the specified address or range of addresses. |
| Options | <p><i>destination-prefix</i>—Address or range of addresses.</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 | <p>show route exact on page 374</p> <p>show route exact detail on page 374</p> <p>show route exact extensive on page 375</p> <p>show route exact terse on page 375</p> |
| Output Fields | For information about output fields, see the <code>show route</code> command (Table 101 on page 331), the <code>show route detail</code> command (Table 104 on page 360), the <code>show route extensive</code> command (Table 110 on page 380), or the <code>show route terse</code> command (Table 118 on page 483). |
| show route exact | <pre> 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 </pre> |
| show route exact detail | <pre> 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 </pre> |

```

show route exact user@host> show route exact 207.17.136.0/24 extensive
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)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| 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>routing-table-name—(Optional) Display information about a particular routing table (for example, inet.0) for which policy-based export is currently enabled. (For information about the different types of routing tables, see the <i>JUNOS Routing Protocols Configuration Guide</i>.)</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 route export on page 377</p> <p>show route export detail on page 377</p> <p>show route export instance detail on page 377</p> |
| Output Fields | Table 108 on page 376 lists the output fields for the show route export command. Output fields are listed in the approximate order in which they appear. |

Table 108: show route export Output Fields

| Field Name | Field Description | Level of Output |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Table or <i>table-name</i> | Name of the routing tables that either import or export routes. | All levels |
| Routes | Number of routes exported from this table into other tables. If a particular route is exported to different tables, the counter will only increment by one. | brief none |
| Export | Whether the table is currently exporting routes to other tables: Y or N (Yes or No). | brief none |
| Import | Tables currently importing routes from the originator table. (Not displayed for tables that are not exporting any routes.) | detail |

Table 108: show route export Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| 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 instance.inet.0. ■ multicast—Indicates instance.inet.2. ■ unicast multicast—Indicates instance.inet.0 and instance.inet.2. | 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 |

```

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 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 on all logical systems.</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 | <p>show route export vrf-target on page 379</p> <p>show route export vrf-target community on page 379</p> <p>show route export vrf-target detail on page 379</p> |
| Output Fields | Table 109 on page 378 lists the output fields for the show route export vrf-target command. Output fields are listed in the approximate order in which they appear. |

Table 109: 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 |
| Export | Number of routing tables that are currently exporting routes with this target community. Omitted for tables that are not exporting routes. | brief none |

Table 109: show route export vrf-target Output Fields (continued)

| Field Name | Field Description | Level of Output |
|-----------------|-------------------------------------------------------------------------------------------------|-----------------|
| 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 |

```

show route export      user@host> show route export vrf-target
vrf-target           Route Target      Family      Import      Export
69:1                   inet      unicast      2           2
69:2                   inet      unicast      2           2

show route export      user@host> show route export vrf-target community target:69:1
vrf-target community Route Target      Family      Import      Export
69:1                   inet      unicast      2           2

show route export      user@host> show route export vrf-target detail
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 <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display extensive information about the active entries in the routing tables. |
| Options | <p>none—Display all active entries in the routing table on all logical systems.</p> <p><i>destination-prefix</i>—(Optional) Display active entries for the specified address or range of addresses.</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 | view |
| List of Sample Output | <p>show route extensive on page 385</p> <p>show route extensive (Access Route) on page 391</p> <p>show route extensive (Route Reflector) on page 391</p> |
| Output Fields | Table 110 on page 380 describes the output fields for the <code>show route extensive</code> command. Output fields are listed in the approximate order in which they appear. |

Table 110: 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> | <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). ■ hidden (routes that are not used because of a routing policy). |

Table 110: 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"> ■ <i>MPLS-label</i> (for example, 80001). ■ <i>interface-name</i> (for example, ge-1/0/2). ■ <i>neighbor-address:control-word-status:encapsulation type:vc-id:source</i> (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96). ■ <i>neighbor-address</i>—Address of the neighbor. ■ <i>control-word-status</i>—Whether the use of the control word has been negotiated for this virtual circuit: NoCtrlWord or CtrlWord. ■ <i>encapsulation type</i>—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. ■ <i>vc-id</i>—Virtual circuit identifier. ■ <i>source</i>—Source of the advertisement: Local or Remote. |
| TSI | Protocol header information. |
| label stacking | <p>(Next-to-the-last-hop router for MPLS only) Depth of the Multiprotocol Label Switching (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 Software 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 | (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. |
| Route Distinguisher | IP subnet augmented with a 64-bit prefix. |

Table 110: 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 Table 105 on page 364. |
| Next-hop reference count | Number of references made to the next hop. |
| 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 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 Multiprotocol Label Switching (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 router is performing unequal-cost load balancing. This information is available when you enable Border Gateway Protocol (BGP) multipath load balancing. |
| Label-switched-path <i>lsp-path-name</i> | Name of the label-switched path (LSP) used to reach the next hop. |
| Label operation | MPLS label and operation occurring at this router. 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 router 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 router. 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 Table 106 on page 365. |

Table 110: show route extensive Output Fields (continued)

| Field Name | Field Description |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Inactive reason | <p>If the route is inactive, the reason for its current state is indicated. Typical reasons include:</p> <ul style="list-style-type: none"> ■ Active preferred—Currently active route was selected over this route. ■ Always compare MED—Path with a lower multiple exit discriminator (MED) is available. ■ AS path—Shorter AS path is available. ■ Cisco Non-deterministic MED selection—Cisco nondeterministic MED is enabled and a path with a lower MED is available. ■ Cluster list length—Path with a shorter cluster list length is available. ■ Forwarding use only—Path is only available for forwarding purposes. ■ IGP metric—Path through the next hop with a lower IGP metric is available. ■ IGP metric type—Path with a lower OSPF link-state advertisement type is available. ■ Interior > Exterior > Exterior via Interior—Direct, static, IGP, or EBGp path is available. ■ Local preference—Path with a higher local preference value is available. ■ Next hop address—Path with a lower metric next hop is available. ■ No difference—Path from a neighbor with a lower IP address is available. ■ Not Best in its group—Occurs when multiple peers of the same external AS advertise the same prefix and are grouped together in the selection process. When this reason is displayed, an additional reason is provided (typically one of the other reasons listed). ■ Number of gateways—Path with a higher number of next hops is available. ■ Origin—Path with a lower origin code is available. ■ OSPF version—Path does not support the indicated OSPF version. ■ RIB preference—Route from a higher-numbered routing table is available. ■ Route distinguisher—64-bit prefix added to IP subnets to make them unique. ■ Route metric or MED comparison—Route with a lower metric or MED is available. ■ Route preference—Route with a lower preference value is available. ■ Router ID—Path through a neighbor with a lower ID is available. ■ Unusable path—Path is not usable because of one of the following conditions: the route is damped, the route is rejected by an import policy, or the route is unresolved. ■ Update source—Last tiebreaker is the lowest IP address value. |
| Local AS | Autonomous system (AS) number of the local router. |
| Age | How long the route has been known. |
| 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. |
| Task | Name of the protocol that has added the route. |
| Announcement bits | List of protocols that announce this route. <i>n-Resolve inet</i> indicates that the route is used for route resolution for next hops found in the routing table <i>n</i> is an index used by Juniper Networks Customer Support only. |

Table 110: show route extensive Output Fields (continued)

| Field Name | Field Description |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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 was 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 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. |
| AS path: I <Originator> | (For router 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 router reflected output only) Cluster ID sent by the route reflector. |
| Originator ID | (For router 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 Table 107 on page 367 for all possible values for this field. |
| Layer2-info: encaps | Layer 2 encapsulation (for example, VPLS). |
| control flags | Control flags: none or Site Down. |
| mtu | Maximum transmission unit (MTU) information. |
| Label-Base, range | First label in a block of labels and label block size. A remote PE router uses this first label when sending traffic toward the advertising PE router. |
| 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. |

Table 110: show route extensive Output Fields (continued)

| Field Name | Field Description |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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 via 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. |

```

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

...

800010 (1 entry, 1 announced)

TSI:
KRT in-kernel 800010 /36 -> {vt-3/2/0.32769}
    *VPLS    Preference: 7
             Next-hop reference count: 2
             Next hop: via vt-3/2/0.32769, selected
             Label operation: Pop
             State: <Active Int>
             Age: 1:31:53
             Task: Common L2 VC
             Announcement bits (1): 0-KRT
             AS path: I

vt-3/2/0.32769 (1 entry, 1 announced)
TSI:
KRT in-kernel vt-3/2/0.32769.0      /16 -> {indirect(1048574)}
    *VPLS    Preference: 7
             Next-hop reference count: 2
             Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
             Label-switched-path green-r1-r3
             Label operation: Push 800012, Push 100096(top)
             Protocol next hop: 10.255.70.103
             Push 800012
             Indirect next hop: 87272e4 1048574
             State: <Active Int>
             Age: 1:31:53    Metric2: 2
             Task: Common L2 VC
             Announcement bits (2): 0-KRT 1-Common L2 VC
             AS path: I
             Communities: target:11111:1 Layer2-info: encaps:VPLS,
             control flags:, mtu: 0
             Indirect next hops: 1
                 Protocol next hop: 10.255.70.103 Metric: 2
                 Push 800012
                 Indirect next hop: 87272e4 1048574
                 Indirect path forwarding next hops: 1
                     Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1
                     10.255.70.103/32 Originating RIB: inet.3
                     Metric: 2                                Node path count: 1
                     Forwarding nexthops: 1
                         Nexthop: 10.31.1.6 via ge-3/1/0.0

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

abcd::10:255:71:52/128 (1 entry, 0 announced)
    *Direct Preference: 0
             Next hop type: Interface
             Next-hop reference count: 1
             Next hop: via lo0.0, selected
             State: <Active Int>
             Local AS: 69

```

```

Age: 1:34:07
Task: IF
AS path: I

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:34:07
    Task: IF
    AS path: I

ff02::2/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::2/128 -> {}
  *PIM Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:34:08
    Task: PIM Recv6
    Announcement bits (1): 0-KRT
    AS path: I

ff02::d/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::d/128 -> {}
  *PIM Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:34:08
    Task: PIM Recv6
    Announcement bits (1): 0-KRT
    AS path: I

ff02::16/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::16/128 -> {}
  *MLD Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:34:06
    Task: MLD
    Announcement bits (1): 0-KRT
    AS path: I

private.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.16385, selected
    State: <Active NoReadvrt Int>
    Age: 1:34:07
    Task: IF

```

```

AS path: I

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

10.255.70.103:1:3:1/96 (1 entry, 1 announced)
  *BGP Preference: 170/-101
    Route Distinguisher: 10.255.70.103:1
    Next-hop reference count: 7
    Source: 10.255.70.103
    Protocol next hop: 10.255.70.103
    Indirect next hop: 2 no-forward
    State: <Secondary Active Int Ext>
    Local AS: 69 Peer AS: 69
    Age: 1:28:12 Metric2: 1
    Task: BGP_69.10.255.70.103+179
    Announcement bits (1): 0-green-l2vpn
    AS path: I
    Communities: target:11111:1 Layer2-info: encaps:VPLS,
    control flags:, mtu: 0
    Label-base: 800008, range: 8
    Localpref: 100
    Router ID: 10.255.70.103
    Primary Routing Table bgp.l2vpn.0

10.255.71.52:1:1:1/96 (1 entry, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 8699540
  *L2VPN Preference: 170/-1
    Next-hop reference count: 5
    Protocol next hop: 10.255.71.52
    Indirect next hop: 0 -
    State: <Active Int Ext>
    Age: 1:34:03 Metric2: 1
    Task: green-l2vpn
    Announcement bits (1): 1-BGP.0.0.0.0+179
    AS path: I
    Communities: Layer2-info: encaps:VPLS, control flags:Site-Down,
    mtu: 0
    Label-base: 800016, range: 8, status-vector: 0x9F

10.255.71.52:1:5:1/96 (1 entry, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 8699528
  *L2VPN Preference: 170/-101
    Next-hop reference count: 5
    Protocol next hop: 10.255.71.52
    Indirect next hop: 0 -
    State: <Active Int Ext>
    Age: 1:34:03 Metric2: 1
    Task: green-l2vpn
    Announcement bits (1): 1-BGP.0.0.0.0+179
    AS path: I
    Communities: Layer2-info: encaps:VPLS, control flags:, mtu: 0
    Label-base: 800008, range: 8, status-vector: 0x9F

...

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

TSI:

```

```

10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
  *L2CKT Preference: 7
    Next hop: via so-1/1/2.0 weight 1, selected
    Label-switched-path my-lsp
    Label operation: Push 100000[0]
    Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
    State: <Active Int>
    Local AS: 99
    Age: 10:21
    Task: 12 circuit
    Announcement bits (1): 0-LDP
    AS path: I
    VC Label 100000, MTU 1500, VLAN ID 512

```

**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

```

show route flow validation

| | |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show route flow validation <brief detail> <table <i>table-name</i> > < <i>ip-prefix</i> > <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display flow route information. |
| Options | <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>table <i>table-name</i>—(Optional) Name of the flow route table.</p> <p><i>ip-prefix</i>—(Optional) IP address for the flow route.</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 flow validation on page 393 |
| Output Fields | Table 111 on page 392 lists the output fields for the show route flow validation command. Output fields are listed in the approximate order in which they appear. |

Table 111: 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 |
| 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 |

```
show route flow      user@host> show route flow validation  
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 show route forwarding-table
 <detail | extensive | summary>
 <all>;
 <ccc interface-name>
 <destination>
 <family family | matching matching>
 <label name>
 <multicast>
 <table (default | routing-table-name)>
 <vpn vpn>

Syntax (Routing Matrix) show route forwarding-table
 <detail | extensive | summary>
 <all>;
 <ccc interface-name>
 <destination>
 <family family | matching matching>
 <label name>
 <lcc number>
 <multicast>
 <table routing-table-name>
 <vpn vpn>

Release Information Command introduced before JUNOS Release 7.4.
 all option introduced in JUNOS Release 9.6.

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. For more information, see the *JUNOS System Basics and Services Command Reference*.

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.

ccc interface-name—(Optional) Display route entries for the specified circuit cross-connect interface.

destination—(Optional) Destination prefix.

family family—(Optional) Display routing table entries for the specified family: inet, inet6, iso, mpls, tnp, unix, or vpls.

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.

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 | routing-table-name)—(Optional) Display route entries for all the routing tables in the main routing instance or for the specified routing table.

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 398
 show route forwarding-table detail on page 399
 show route forwarding-table destination extensive (Weights and Balances) on page 400
 show route forwarding-table extensive on page 400
 show route forwarding-table extensive (RPF) on page 401
 show route forwarding-table family mpls on page 402
 show route forwarding-table family vpls on page 402
 show route forwarding-table family vpls extensive on page 402
 show route forwarding-table vpn on page 404

Output Fields Table 112 on page 395 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 may 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 112: show route forwarding-table Output Fields

| Field Name | Field Description | Level of Output |
|----------------|--------------------------------------------------------------|------------------|
| 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 112: show route forwarding-table Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| Route Type (Type) | <p>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):</p> <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 | <p>Route type flags:</p> <ul style="list-style-type: none"> ■ none—No flags are enabled. ■ accounting—Route has accounting enabled. ■ cached—Cache route. ■ incoming-iface<i>interface-number</i>—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 112: 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. ■ 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 (mcr)—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 (recv)—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. | none detail extensive |
| Next-hop interface (Netif) | Interface used to reach the next hop. | none detail extensive |
| Weight | Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when Multiprotocol Label Switching (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 Border Gateway Protocol (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 |

```

show route forwarding-table user@host> show route forwarding-table
Routing table: inet
Internet:
Destination          Type RtRef Next hop          Type Index NhRef Netif
default              user   0                rjct   10    1
default              perm   0                rjct   12    2
2.2.2.0/24           ifdn   0 f.0.8.0            ucst   39    1
so-1/1/0.0
2.2.2.1/32           user   0                rjct   12    2
2.2.2.1/32           intf   0 2.2.2.1           locl   38    1
4.1.1.0/24           intf   0 f.0.8.0            ucst   43    1
so-1/0/0.0
4.1.1.2/32           intf   0 4.1.1.2           locl   42    1
5.1.1.0/24           user   0                ucst   44    4
so-1/0/0.0
10.255.245.220/32    intf   0 10.255.245.220     locl   25    1
10.255.245.245/32    user   0                ucst   44    4
so-1/0/0.0
10.255.245.246/32    user   0                ucst   44    4
so-1/0/0.0
192.16.0.0/12        user   1 192.168.5.254       ucst   29    11 fxp0.0
192.168.0.0/18       user   0 192.168.5.254       ucst   29    11 fxp0.0
192.168.0.0/20       user   0 192.168.5.254       ucst   29    11 fxp0.0
192.168.1.0/24       user   0 192.168.5.254       ucst   29    11 fxp0.0
192.168.5.0/24       intf   0                rslv   22    1 fxp0.0
192.168.5.0/32       dest   0 192.168.5.0         recv   20    1 fxp0.0
192.168.5.49/32      dest   0 0:0:c0:e8:69:db     ucst   31    1 fxp0.0
192.168.5.73/32      dest   0 0:a0:c9:85:c:44     ucst   32    1 fxp0.0
192.168.5.80/32      dest   0 0:d0:b7:1e:92:f2    ucst   30    1 fxp0.0
192.168.5.220/32     intf   0 192.168.5.220       locl   21    2

...

Routing table: iso
ISO:
Destination          Type RtRef Next hop          Type Index NhRef Netif
default              perm   0                rjct   27    1
47.0005.80ff.f800.0000.0108.0003.0102.5524.5220.00
intf   0                locl   28    1

Routing table: inet6
Internet6:
Destination          Type RtRef Next hop          Type Index NhRef Netif
default              perm   0                rjct   6    1
ff00::/8             perm   0                mdsc   4    1
ff02::1/128          perm   0 ff02::1            mcst   3    1

```

```

Routing table: ccc
MPLS:
Interface.Label      Type RtRef Next hop          Type Index NhRef Netif
default              perm  0                rjct 16      1
100004(top)fe-0/0/1.0

```

**show route
forwarding-table detail**

```

user@host> show route forwarding-table detail
Routing table: inet
Internet:
Destination          Type RtRef Next hop          Type Index NhRef Netif
default              user  2 0:90:69:8e:b1:1b ucst 132   4 fxp0.0
default              perm  0                rjct 14    1
10.1.1.0/24          intf  0 ff.3.0.21          ucst 322   1 so-5/3/0.0
10.1.1.0/32          dest  0 10.1.1.0          recv 324   1 so-5/3/0.0
10.1.1.1/32          intf  0 10.1.1.1          locl 321   1
10.1.1.255/32        dest  0 10.1.1.255        bcst 323   1 so-5/3/0.0
10.21.21.0/24        intf  0 ff.3.0.21          ucst 326   1 so-5/3/0.0
10.21.21.0/32        dest  0 10.21.21.0        recv 328   1 so-5/3/0.0
10.21.21.1/32        intf  0 10.21.21.1        locl 325   1
10.21.21.255/32      dest  0 10.21.21.255      bcst 327   1 so-5/3/0.0
127.0.0.1/32         intf  0 127.0.0.1          locl 320   1
172.17.28.19/32      clon  1 192.168.4.254      ucst 132   4 fxp0.0
172.17.28.44/32      clon  1 192.168.4.254      ucst 132   4 fxp0.0

```

...

```

Routing table: private1__inet
Internet:
Destination          Type RtRef Next hop          Type Index NhRef Netif
default              perm  0                rjct 46    1
10.0.0.0/8           intf  0                rslv 136   1 fxp1.0
10.0.0.0/32          dest  0 10.0.0.0          recv 134   1 fxp1.0
10.0.0.4/32          intf  0 10.0.0.4          locl 135   2
10.0.0.4/32          dest  0 10.0.0.4          locl 135   2

```

...

```

Routing table: iso
ISO:
Destination          Type RtRef Next hop          Type Index NhRef Netif
default              perm  0                rjct 38    1

```

```

Routing table: inet6
Internet6:
Destination          Type RtRef Next hop          Type Index NhRef Netif
default              perm  0                rjct 22    1
ff00::/8             perm  0                mdsc 21    1
ff02::1/128          perm  0 ff02::1          mcst 17    1

```

...

```

Routing table: mpls
MPLS:
Destination          Type RtRef Next hop          Type Index NhRef Netif
default              perm  0                rjct 28    1

```

show route forwarding-table destination extensive (Weights and Balances) user@host> **show route forwarding-table destination 3.4.2.1 extensive**
 Routing table: inet [Index 0]
 Internet:

```

Destination: 3.4.2.1/32
  Route type: user
  Route reference: 0
  Flags: sent to PFE
  Next-hop type: unicast
  Nexthop: 4.4.4.4
  Next-hop type: unicast
  Next-hop interface: so-1/1/0.0
  Nexthop: 145.12.1.2
  Next-hop type: unicast
  Next-hop interface: so-0/1/2.0
  Route interface-index: 0
  Index: 262143   Reference: 1
  Index: 335      Reference: 2
  Weight: 22      Balance: 3
  Index: 337      Reference: 2
  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
  Next-hop type: unicast
  Next-hop interface: fxp0.0
  Route interface-index: 0
  Index: 132      Reference: 4

```

```

Destination: default
  Route type: permanent
  Route reference: 0
  Flags: none
  Next-hop type: reject
  Route interface-index: 0
  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
  Next-hop type: local
  Route interface-index: 0
  Index: 320      Reference: 1

```

...

Routing table: private1__inet [Index 1]
 Internet:

```

Destination: default
  Route type: permanent
  Route reference: 0
  Flags: sent to PFE
  Next-hop type: reject
  Route interface-index: 0
  Index: 46       Reference: 1

```

```

Destination: 10.0.0.0/8
  Route type: interface
  Route reference: 0
  Flags: sent to PFE
  Next-hop type: resolve
  Next-hop interface: fxp1.0
  Route interface-index: 3
  Index: 136      Reference: 1

```

...

```

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          dym  0
default          perm  0
fe-0/1/0.0       dym  0
00:90:69:0c:20:1f/48      <<<<<Remote CE

                        dym  0                          indr  351      4
                        Push 800000, Push 100002(top)

so-0/0/0.0
00:90:69:85:b0:1f/48      <<<<<Local CE

                        dym  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                               Route interface-index: 72
Flags: sent to PFE
Next-hop type: flood                             Index: 289      Reference: 1
Next-hop type: unicast                           Index: 291      Reference: 3
Next-hop interface: fe-0/1/3.0
Next-hop type: unicast                           Index: 290      Reference: 3
Next-hop interface: fe-0/1/2.0

Destination: default
Route type: permanent
Route reference: 0                               Route interface-index: 0
Flags: none
Next-hop type: discard                           Index: 341      Reference: 1

```



```

Destination: fe-0/1/2.0
Route type: dynamic
Route reference: 0
Flags: sent to PFE
Next-hop type: flood
Next-hop type: indirect
Next-hop type: Push 800016
Next-hop interface: at-1/0/1.0
Next-hop type: indirect
Next hop: 10.31.3.2
Next-hop type: Push 800000
Next-hop interface: fe-0/1/1.0
Next-hop type: unicast
Next-hop interface: fe-0/1/3.0
Route interface-index: 69
Index: 293 Reference: 1
Index: 363 Reference: 4
Index: 301 Reference: 5
Index: 291 Reference: 3

Destination: fe-0/1/3.0
Route type: dynamic
Route reference: 0
Flags: sent to PFE
Next-hop type: flood
Next-hop type: indirect
Next-hop type: Push 800016
Next-hop interface: at-1/0/1.0
Next-hop type: indirect
Next hop: 10.31.3.2
Next-hop type: Push 800000
Next-hop interface: fe-0/1/1.0
Next-hop type: unicast
Next-hop interface: fe-0/1/2.0
Route interface-index: 70
Index: 292 Reference: 1
Index: 363 Reference: 4
Index: 301 Reference: 5
Index: 290 Reference: 3

Destination: 10:00:00:01:01:01/48
Route type: dynamic
Route reference: 0
Flags: sent to PFE, prefix load balance
Next-hop type: unicast
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
Route interface-index: 70
Index: 291 Reference: 3

Destination: 10:00:00:01:01:04/48
Route type: dynamic
Route reference: 0
Flags: sent to PFE, prefix load balance
Next-hop type: unicast
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
Route interface-index: 69
Index: 290 Reference: 3

Destination: 10:00:00:01:03:05/48
Route type: dynamic
Route reference: 0
Flags: sent to PFE, prefix load balance
Next-hop type: indirect
Next hop: 10.31.3.2
Next-hop type: Push 800000
Next-hop interface: fe-0/1/1.0
Route interface-index: 74
Index: 301 Reference: 5

```

```

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      Type Index NhRef Netif
  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
  100004(top) so-1/0/0.0 Push 100004, Push
  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 hidden

| | |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show route hidden <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display only hidden route information. A hidden route is unusable, even if it is the best path. |
| Options | <p>brief detail extensive terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Required Privilege Level | view |
| List of Sample Output | <p>show route hidden on page 405</p> <p>show route hidden detail on page 406</p> <p>show route hidden extensive on page 406</p> <p>show route hidden terse on page 406</p> |
| Output Fields | For information about output fields, see the <code>show route</code> command (Table 101 on page 331), the <code>show route detail</code> command (Table 104 on page 360), the <code>show route extensive</code> command (Table 110 on page 380), or the <code>show route terse</code> command (Table 118 on page 483). |
| show route hidden | <pre> 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 ... </pre> |

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 406.

show route hidden terse user@host> **show route hidden terse**

```

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
127.0.0.1/32      D   0                >100.0

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
10.5.5.5/32        B 170      100        Unusable     100 I
10.12.1.0/24       B 170      100        Unusable     100 I

```

```

10.12.80.4/30      B 170      100      Unusable      I

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

bgp.l3vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
10.4.4.4:4:10.5.5.5/32
                   B 170      100      Unusable      100 I
10.4.4.4:4:10.12.1.0/24
                   B 170      100      Unusable      100 I
10.4.4.4:4:10.12.80.4/30
                   B 170      100      Unusable      I

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

show route inactive-path

| | |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show route inactive-path <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| 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 | <p>show route inactive-path on page 408</p> <p>show route inactive-path detail on page 409</p> <p>show route inactive-path extensive on page 410</p> <p>show route inactive-path terse on page 410</p> |
| Output Fields | For information about output fields, see the <code>show route</code> command (Table 101 on page 331), the <code>show route detail</code> command (Table 104 on page 360), the <code>show route extensive</code> command (Table 110 on page 380), or the <code>show route terse</code> command (Table 118 on page 483). |

```

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 409.

show route inactive-path terse

```
user@host> show route inactive-path terse

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1   Metric 2   Next hop      AS path
  10.12.100.12/30   0  10           1             >so-0/3/0.0

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1   Metric 2   Next hop      AS path
  10.0.0.0/8        D   0             >fxp1.0

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1   Metric 2   Next hop      AS path
  10.12.80.0/30     B 170          100             >10.12.80.1    100 I

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

bgp.l3vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```


show route inactive-prefix

| | |
|---------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show route inactive-prefix <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| 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 | <p>show route inactive-prefix on page 411</p> <p>show route inactive-prefix detail on page 411</p> <p>show route inactive-prefix extensive on page 411</p> <p>show route inactive-prefix terse on page 412</p> |
| Output Fields | For information about output fields, see the <code>show route</code> command (Table 101 on page 331), the <code>show route detail</code> command (Table 104 on page 360), the <code>show route extensive</code> command (Table 110 on page 380), or the <code>show route terse</code> command (Table 118 on page 483). |
| show route inactive-prefix | <pre>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</pre> |
| show route inactive-prefix detail | <pre>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</pre> |
| show route inactive-prefix extensive | The output for the <code>show route inactive-prefix extensive</code> command is identical to that of the <code>show route inactive-path detail</code> command. For sample output, see <code>show route inactive-prefix detail</code> on page 411. |

```
show route      user@host> show route inactive-prefix terse
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                >1o0.0
```

show route instance

| | |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show route instance <brief detail summary> <instance-name> <logical-system (all <i>logical-system-name</i>)> <operational> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display routing instance information. |
| Options | <p>none—(Same as brief) Display standard information about all routing instances on all logical systems.</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><i>instance-name</i>—(Optional) Display information for a 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>operational—(Optional) Display operational routing instances.</p> |
| Required Privilege Level | view |
| List of Sample Output | <p>show route instance on page 414</p> <p>show route instance detail (Graceful Restart Complete) on page 414</p> <p>show route instance detail (Graceful Restart Incomplete) on page 416</p> <p>show route instance detail (VPLS Routing Instance) on page 418</p> <p>show route instance operational on page 418</p> <p>show route instance summary on page 418</p> |
| Output Fields | Table 113 on page 413 lists the output fields for the show route instance command. Output fields are listed in the approximate order in which they appear. |

Table 113: show route instance Output Fields

| Field Name | Field Description | Level of Output |
|----------------------------------|-------------------------------------------------------------------------------------------------------------------|--------------------------|
| Instance or <i>instance-name</i> | Name of the routing instance. | All levels |
| Operational Routing Instances | (operational keyword only) Names of all operational routing instances. | — |
| Type | Type of routing instance: forwarding , l2vpn , no-forwarding , vpls , 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 |

Table 113: show route instance Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------|
| 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. | none brief detail |
| 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:<i>protocol-name</i>—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 |

```

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 |

| | | | |
|--------|-------|-----------------|-------|
| L2VPN | l2vpn | BGP-L.iso.0 | 0/0/0 |
| | | BGP-L.mpls.0 | 4/0/0 |
| | | BGP-L.inet6.0 | 0/0/0 |
| | | L2VPN.inet.0 | 0/0/0 |
| | | L2VPN.iso.0 | 0/0/0 |
| | | L2VPN.inet6.0 | 0/0/0 |
| LDP | vrf | L2VPN.l2vpn.0 | 2/0/0 |
| | | LDP.inet.0 | 4/0/0 |
| | | LDP.iso.0 | 0/0/0 |
| | | LDP.mpls.0 | 0/0/0 |
| | | LDP.inet6.0 | 0/0/0 |
| OSPF | vrf | LDP.l2circuit.0 | 0/0/0 |
| | | 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 | show route label <i>label</i> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| 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 | <p>show route label on page 420</p> <p>show route label detail on page 421</p> <p>show route label extensive on page 421</p> <p>show route label terse on page 421</p> |
| Output Fields | For information about output fields, see the show route command (Table 101 on page 331), the show route detail command (Table 104 on page 360), the show route extensive command (Table 110 on page 380), or the show route terse command (Table 118 on page 483). The Ref Cnt field name in the show route label command is equivalent to the Reference Count field displayed in the show route command. |
| show route label | <pre> 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 </pre> |

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 421.

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 | show route label-switched-path <i>path-name</i> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display the routes used in a Multiprotocol Label Switching (MPLS) label-switched path (LSP). |
| Options | <p><i>path-name</i>—LSP tunnel name.</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 |
| List of Sample Output | show route label-switched-path on page 422 |
| Output Fields | For information about output fields, see the show route command (Table 101 on page 331), the show route detail command (Table 104 on page 360), the show route extensive command (Table 110 on page 380), or the show route terse command (Table 118 on page 483). |
| show route label-switched-path | <pre> 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 </pre> |

show route martians

| | |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show route martians <logical-system (all <i>logical-system-name</i>)> <table <i>routing-table-name</i> > |
| Release Information | Command introduced before JUNOS Release 7.4. |
| 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 on all logical systems.</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 only the martian entries associated with a particular routing table.</p> |
| Required Privilege Level | view |
| List of Sample Output | show route martians on page 423 |
| Output Fields | Table 114 on page 423 lists the output fields for the show route martians command. Output fields are listed in the approximate order in which they appear |

Table 114: 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. |

```

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
    128.0.0.0/16 orlonger -- disallowed
    191.255.0.0/16 orlonger -- disallowed
    192.0.0.0/24 orlonger -- disallowed
    223.255.255.0/24 orlonger -- disallowed
    240.0.0.0/4 orlonger -- disallowed

```

```
inet.1:
0.0.0.0/0 exact -- allowed
0.0.0.0/8 orlonger -- disallowed
127.0.0.0/8 orlonger -- disallowed
128.0.0.0/16 orlonger -- disallowed
191.255.0.0/16 orlonger -- disallowed
192.0.0.0/24 orlonger -- disallowed
223.255.255.0/24 orlonger -- disallowed
240.0.0.0/4 orlonger -- disallowed

....
```

show route next-hop

| | |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show route next-hop <i>next-hop</i> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display the entries in the routing table that are being sent to the specified next-hop address. |
| Options | <p>none—Display information about all entries in the routing table that are being sent to the specified next-hop address for routes on all logical systems.</p> <p><i>next-hop</i>—Next-hop address.</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 |
| List of Sample Output | <p>show route next-hop on page 425</p> <p>show route next-hop detail on page 426</p> <p>show route next-hop extensive on page 427</p> <p>show route next-hop terse on page 429</p> |
| Output Fields | For information about output fields, see the <code>show route</code> command (Table 101 on page 331), the <code>show route detail</code> command (Table 104 on page 360), the <code>show route extensive</code> command (Table 110 on page 380), or the <code>show route terse</code> command (Table 118 on page 483). |

```

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-kerne1 207.17.136.192/32 -> {192.168.71.254}
```

```

*Static Preference: 5
  Next-hop reference count: 22
  Next hop: 192.168.71.254 via fxp0.0, selected
  State: <Active NoReadvrt Int Ext>
  Local AS: 69
  Age: 2:02:28
  Task: RT
  Announcement bits (1): 0-KRT
  AS path: I

```

```
private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
```

```
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

```
mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
```

```
inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
```

```
private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

```
green.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
```

```
red.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route next-hop terse user@host> **show route next-hop 192.168.71.254 terse**

```
inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
```

```
Restart Complete
```

```
+ = Active Route, - = Last Active, * = Both
```

| A | Destination | P | Prf | Metric 1 | Metric 2 | Next hop | AS path |
|---|-------------------|---|-----|----------|----------|-----------------|---------|
| * | 10.10.0.0/16 | S | 5 | | | >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>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display the route entries in each routing table that are not associated with any community. |
| Options | <p>none—(Same as brief) Display the route entries in each routing table that are not associated with any community.</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 |
| List of Sample Output | <p>show route no-community on page 431</p> <p>show route no-community detail on page 432</p> <p>show route no-community extensive on page 432</p> <p>show route no-community terse on page 433</p> |
| Output Fields | For information about output fields, see the show route command (Table 101 on page 331), the show route detail command (Table 104 on page 360), the show route extensive command (Table 110 on page 380), or the show route terse command (Table 118 on page 483). |
| show route no-community | <pre> 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 </pre> |

```

> to 35.1.1.2 via ge-3/1/0.0
  via so-0/1/2.0
  via so-0/3/2.0
16.1.1.0/24    *[OSPF/10] 00:05:14, metric 2
> via so-0/1/2.0
.....

```

**show route
no-community detail**

```

user@host> show route no-community detail

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 38:08
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

10.209.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 38:08
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

....

```

**show route
no-community extensive**

```

user@host> show route no-community extensive

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 2:03:33
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

10.209.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.209.0.0/16 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 2:03:33
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

```

```

show route      user@host> show route no-community terse
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                >100.0
* 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>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| 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 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>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>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 | <p>show route output address on page 434</p> <p>show route output address detail on page 435</p> <p>show route output address extensive on page 435</p> <p>show route output address terse on page 436</p> <p>show route output interface on page 436</p> <p>show route output interface detail on page 436</p> <p>show route output interface extensive on page 437</p> <p>show route output interface terse on page 437</p> |
| Output Fields | For information about output fields, see the show route command (Table 101 on page 331), the show route detail command (Table 104 on page 360), the show route extensive command (Table 110 on page 380), or the show route terse command (Table 118 on page 483). |
| show route output address | <pre> 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 </pre> |


```

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 435.

**show route output
address terse**

```
user@host> show route output address 36.1.1.1 terse
```

```
inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
```

| A Destination | P | Prf | Metric 1 | Metric 2 | Next hop | AS path |
|---------------|---|-----|----------|----------|-------------|---------|
| * 36.1.1.0/24 | D | 0 | | | >so-0/1/2.0 | |
| | 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 436.

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       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 protocol

| | |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show route protocol <i>protocol</i> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. ospf2 and ospf3 options introduced in JUNOS Release 9.2. |
| Description | Display the route entries in the routing table that were learned from a particular protocol. |
| Options | <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 ■ 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 ■ 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

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

Required Privilege Level view

List of Sample Output

show route protocol access on page 439
 show route protocol access-internal extensive on page 439
 show route protocol bgp on page 440
 show route protocol bgp detail on page 440
 show route protocol bgp extensive on page 440
 show route protocol bgp terse on page 440
 show route protocol direct on page 441
 show route protocol l2circuit detail on page 441
 show route protocol l2vpn extensive on page 442
 show route protocol ldp on page 442
 show route protocol ldp extensive on page 443
 show route protocol ospf (Layer 3 VPN) on page 444
 show route protocol ospf detail on page 444
 show route protocol rip on page 445
 show route protocol rip detail on page 445
 show route protocol ripng table inet6 on page 445

Output Fields For information about output fields, see the show route command (Table 101 on page 331), the show route detail command (Table 104 on page 360), the show route extensive command (Table 110 on page 380), or the show route terse command (Table 118 on page 483).

show route protocol access

```
user@host> show route protocol access

inet.0: 30380 destinations, 30382 routes (30379 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

13.160.0.3/32      *[Access/13] 00:00:09
                  > to 13.160.0.2 via fe-0/0/0.0
13.160.0.4/32      *[Access/13] 00:00:09
                  > to 13.160.0.2 via fe-0/0/0.0
13.160.0.5/32      *[Access/13] 00:00:09
                  > to 13.160.0.2 via fe-0/0/0.0
```

show route protocol access-internal extensive

```
user@host> show route protocol access-internal 13.160.0.19 extensive

inet.0: 100020 destinations, 100022 routes (100019 active, 0 holddown, 1 hidden)
13.160.0.19/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 13.160.0.19/32 -> {13.160.0.2}
  *Access-internal Preference: 12
    Next-hop reference count: 200000
    Next hop: 13.160.0.2 via fe-0/0/0.0, selected
    State: <Active Int>

Age: 36
```

```
Task: RPD Unix Domain Server./var/run/rpd_serv.local
Announcement bits (1): 0-KRT
AS path: I
```

show route protocol bgp user@host> **show route protocol bgp 192.168.64.0/21**

```
inet.0: 24 destinations, 32 routes (23 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
192.168.64.0/21      [BGP/170] 00:04:33, localpref 100
                    AS path: 10023 21 I
                    > to 100.1.3.2 via ge-5/0/3.0, Push 100080
```

show route protocol bgp detail **show route protocol bgp 66.117.63.0/24 exact detail**

```
inet.0: 227318 destinations, 227319 routes (227305 active, 0 holddown, 13 hidden)
66.117.63.0/24 (1 entry, 1 announced)
  *BGP      Preference: 170/-101
            Next hop type: Indirect
            Next-hop reference count: 681816
            Source: 207.17.136.192
            Next hop type: Router, Next hop index: 324
            Next hop: 192.168.167.254 via fxp0.0, selected
            Protocol next hop: 207.17.136.29
            Indirect next hop: 8c7b09c 342
            State: <Active Int Ext
            Local AS: 200 Peer AS: 10458
            Age: 20:31:24 Metric2: 0
            Task: BGP_10458_10458.207.17.136.192+179
            Announcement bits (2): 0-KRT 3-Resolve tree 2
            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
            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: 24 destinations, 32 routes (23 active, 0 holddown, 1 hidden)
192.168.64.0/21 (2 entries, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 86f50a8
  BGP      Preference: 170/-101
            Next-hop reference count: 3
            Source: 100.1.3.2
            Next hop: 100.1.3.2 via ge-5/0/3.0, selected
            Label operation: Push 100080
            State: <Ext>
            Inactive reason: Route Preference
            Local AS: 21 Peer AS: 10023
            Age: 4:43
            Task: BGP_10023.100.1.3.2+4282
            AS path: 10023 21 I
            Route Label: 100080
            Localpref: 100
            Router ID: 100.1.3.2
```

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: 35 destinations, 35 routes (34 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

```
127.0.0.1/32      [Direct/0] 14:36:24
                  > via lo0.0
111.222.5.0/24    *[Direct/0] 14:36:24
                  > via fxp0.0
111.222.8.16/28   *[Direct/0] 14:36:24
                  > via at-5/3/0.0
111.222.8.100/30  *[Direct/0] 14:36:24
                  > via at-5/3/0.129
111.222.8.104/30  *[Direct/0] 14:36:24
                  > via at-5/3/0.128
111.222.8.161/32  *[Direct/0] 14:36:24
                  > via t3-5/2/0.0
111.222.8.163/32  *[Direct/0] 14:36:24
                  > via t3-5/2/1.0
...
```

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

```
47.0005.80ff.f800.0000.0108.0001.1921.6800.5081.00/160
                  *[Direct/0] 14:36:24
                  > via lo0.0
```

show route protocol l2circuit detail

user@router> show route protocol l2circuit detail

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
100000 (1 entry, 1 announced)

```
*L2CKT Preference: 7
  Next hop: via ge-2/0/0.0, selected
  Label operation: Pop      Offset: 4
  State: <Active Int>
  Local AS:    99
  Age: 9:52
  Task: Common L2 VC
  Announcement bits (1): 0-KRT
  AS path: I
```

ge-2/0/0.0 (1 entry, 1 announced)

```
*L2CKT Preference: 7
  Next hop: via so-1/1/2.0 weight 1, selected
  Label-switched-path my-lsp
  Label operation: Push 100000, Push 100000(top)[0] Offset: -4
  Protocol next hop: 10.245.255.63
  Push 100000 Offset: -4
  Indirect next hop: 86af0c0 298
  State: <Active Int>
  Local AS:    99
  Age: 9:52
  Task: Common L2 VC
  Announcement bits (2): 0-KRT 1-Common L2 VC
  AS path: I
```

```

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
  *L2CKT Preference: 7
    Next hop: via so-1/1/2.0 weight 1, selected
    Label-switched-path my-lsp
    Label operation: Push 100000[0]
    Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
    State: <Active Int>
    Local AS: 99
    Age: 10:21
    Task: l2 circuit
    Announcement bits (1): 0-LDP
    AS path: I
    VC Label 100000, MTU 1500, VLAN ID 512

```

**show route protocol
l2vpn extensive**

```

user@host> show route protocol l2vpn extensive

inet.0: 14 destinations, 15 routes (13 active, 0 holddown, 1 hidden)

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
800001 (1 entry, 1 announced)
TSI:
KRT in-kernel 800001 /36 -> {so-0/0/0.0}
  *L2VPN Preference: 7
    Next hop: via so-0/0/0.0 weight 49087 balance 97%, selected
    Label operation: Pop Offset: 4
    State: <Active Int>
    Local AS: 69
    Age: 7:48
    Task: Common L2 VC
    Announcement bits (1): 0-KRT
    AS path: I

so-0/0/0.0 (1 entry, 1 announced)
TSI:
KRT in-kernel so-0/0/0.0 /16 -> {indirect(288)}
  *L2VPN Preference: 7
    Next hop: via so-0/0/1.0, selected
    Label operation: Push 800000 Offset: -4
    Protocol next hop: 10.255.14.220
    Push 800000 Offset: -4
    Indirect next hop: 85142a0 288
    State: <Active Int>
    Local AS: 69
    Age: 7:48
    Task: Common L2 VC
    Announcement bits (2): 0-KRT 1-Common L2 VC
    AS path: I
    Communities: target:69:1 Layer2-info: encaps:PPP,
    control flags:2, mtu: 0

```

show route protocol ldp

```

user@host> show route protocol ldp

inet.0: 12 destinations, 13 routes (12 active, 0 holddown, 0 hidden)

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

```


+ = Active Route, - = Last Active, * = Both

```
192.168.16.1/32    *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Push 100000
192.168.17.1/32    *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0
```

private1___.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)

+ = Active Route, - = Last Active, * = Both

```
100064            *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Pop
100064(S=0)        *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Pop
100080            *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Swap 100000
```

show route protocol ldp extensive

user@host> show route protocol ldp extensive

```
192.168.16.1/32 (1 entry, 1 announced)
  State: <FlashAll>
  *LDP    Preference: 9
          Next-hop reference count: 3
          Next hop: via t1-4/0/0.0, selected
          Label operation: Push 100000
          State: <Active Int>
          Local AS: 65500
          Age: 1d 23:03:58      Metric: 1
          Task: LDP
          Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 2
          AS path: I
```

```
192.168.17.1/32 (1 entry, 1 announced)
  State: <FlashAll>
  *LDP    Preference: 9
          Next-hop reference count: 3
          Next hop: via t1-4/0/0.0, selected
          State: <Active Int>
          Local AS: 65500
          Age: 1d 23:03:58      Metric: 1
          Task: LDP
          Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 2
          AS path: I
```

private1___.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)

100064 (1 entry, 1 announced)

TSI:

KRT in-kernel 100064 /36 -> {t1-4/0/0.0}

```
*LDP    Preference: 9
          Next-hop reference count: 2
          Next hop: via t1-4/0/0.0, selected
          State: <Active Int>
          Local AS: 65500
          Age: 1d 23:03:58      Metric: 1
          Task: LDP
          Announcement bits (1): 0-KRT
          AS path: I
```

Prefixes bound to route: 192.168.17.1/32

100064(S=0) (1 entry, 1 announced)

TSI:

KRT in-kernel 100064 /40 -> {t1-4/0/0.0}

```
*LDP      Preference: 9
           Next-hop reference count: 2
           Next hop: via t1-4/0/0.0, selected
           Label operation: Pop
           State: <Active Int>
           Local AS: 65500
           Age: 1d 23:03:58      Metric: 1
           Task: LDP
           Announcement bits (1): 0-KRT
           AS path: I
```

100080 (1 entry, 1 announced)

TSI:

KRT in-kernel 100080 /36 -> {t1-4/0/0.0}

```
*LDP      Preference: 9
           Next-hop reference count: 2
           Next hop: via t1-4/0/0.0, selected
           Label operation: Swap 100000
           State: <Active Int>
           Local AS: 65500
           Age: 1d 23:03:58      Metric: 1
           Task: LDP
           Announcement bits (1): 0-KRT
           AS path: I
           Prefixes bound to route: 192.168.16.1/32
```

show route protocol ospf (Layer 3 VPN)

user@host> **show route protocol ospf**

inet.0: 40 destinations, 40 routes (39 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

```
10.39.1.4/30      *[OSPF/10] 00:05:18, metric 4
                  > via t3-3/2/0.0
10.39.1.8/30      [OSPF/10] 00:05:18, metric 2
                  > via t3-3/2/0.0
10.255.14.171/32  *[OSPF/10] 00:05:18, metric 4
                  > via t3-3/2/0.0
10.255.14.179/32  *[OSPF/10] 00:05:18, metric 2
                  > via t3-3/2/0.0
224.0.0.5/32      *[OSPF/10] 20:25:55, metric 1
```

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

```
10.39.1.16/30     [OSPF/10] 00:05:43, metric 1
                  > via so-0/2/2.0
10.255.14.173/32  *[OSPF/10] 00:05:43, metric 1
                  > via so-0/2/2.0
224.0.0.5/32      *[OSPF/10] 20:26:20, metric 1
```

show route protocol ospf detail

user@host> **show route protocol ospf detail**

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

```
10.39.1.16/30 (2 entries, 0 announced)
  OSPF      Preference: 10
            Nexthop: via so-0/2/2.0, selected
```

```

State: <Int>
Inactive reason: Route Preference
Age: 6:25      Metric: 1
Area: 0.0.0.0
Task: VPN-AB-OSPF
AS path: I
Communities: Route-Type:0.0.0.0:1:0

```

...

show route protocol rip

```

user@host> show route protocol rip
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32  * [RIP/100] 20:24:34, metric 2
                  > to 10.39.1.22 via t3-0/2/2.0
224.0.0.9/32     * [RIP/100] 00:03:59, metric 1

```

show route protocol rip detail

```

user@host> show route protocol rip detail
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32 (1 entry, 1 announced)
    *RIP      Preference: 100
              Nexthop: 10.39.1.22 via t3-0/2/2.0, selected
              State: <Active Int>
              Age: 20:25:02  Metric: 2
              Task: VPN-AB-RIPv2
              Announcement bits (2): 0-KRT 2-BGP.0.0.0.0+179
              AS path: I
              Route learned from 10.39.1.22 expires in 96 seconds

```

show route protocol ripng table inet6

```

user@host> show route protocol ripng table inet6
inet6.0: 4215 destinations, 4215 routes (4214 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

1111::1/128      * [RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::2/128      * [RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::3/128      * [RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::4/128      * [RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::5/128      * [RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::6/128      * [RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0

```

show route range

| | |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show route range <brief detail extensive terse> <destination-prefix> <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display routing table entries using a prefix range. |
| Options | <p>none—Display standard information about all routing table entries using a prefix range on all logical systems.</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><i>destination-prefix</i>—(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 446</p> <p>show route range destination-prefix on page 447</p> <p>show route range detail on page 447</p> <p>show route range extensive on page 448</p> <p>show route range terse on page 449</p> |
| Output Fields | For information about output fields, see the show route command (Table 101 on page 331), the show route detail command (Table 104 on page 360), the show route extensive command (Table 110 on page 380), or the show route terse command (Table 118 on page 483). |
| show route range | <pre> 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 </pre> |

```

...
> 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-kernel 10.10.0.0/16 -> {192.168.71.254}
    *Static Preference: 5
        Next-hop reference count: 22
        Next hop: 192.168.71.254 via fxp0.0, selected
        State: <Active NoReadvrt Int Ext>
        Age: 30:17
        Task: RT
        Announcement bits (1): 0-KRT
        AS path: I

10.209.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.209.0.0/16 -> {192.168.71.254}
    *Static Preference: 5
        Next-hop reference count: 22
        Next hop: 192.168.71.254 via fxp0.0, selected
        State: <Active NoReadvrt Int Ext>
        Age: 30:17
        Task: RT
        Announcement bits (1): 0-KRT
        AS path: I

10.255.71.14/32 (1 entry, 0 announced)
    *Direct Preference: 0
        Next hop type: Interface
        Next-hop reference count: 1
        Next hop: via lo0.0, selected
        State: <Active Int>
        Age: 30:17
        Task: IF
        AS path: I

172.16.0.0/12 (1 entry, 1 announced)
TSI:
KRT in-kernel 172.16.0.0/12 -> {192.168.71.254}
    *Static Preference: 5
        Next-hop reference count: 22
        Next hop: 192.168.71.254 via fxp0.0, selected
        State: <Active NoReadvrt Int Ext>
        Age: 30:17
        Task: RT
        Announcement bits (1): 0-KRT
        AS path: I

...
```

show route range terse user@host> **show route range terse**

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

| A Destination | P | Prf | Metric 1 | Metric 2 | Next hop | AS path |
|---------------------|---|-----|----------|----------|-----------------|---------|
| * 10.10.0.0/16 | S | 5 | | | >192.168.71.254 | |
| * 10.209.0.0/16 | S | 5 | | | >192.168.71.254 | |
| * 10.255.71.14/32 | D | 0 | | | >100.0 | |
| * 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 | | | >100.0 | |

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 | | | >100.0 | |
| fe80::280:42ff:fe11:226f/128 | | | | | | |
| * | D | 0 | | | >100.0 | |

__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 | | | >100.16385 | |

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>) |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display the routing information as it was received through a particular neighbor using a particular dynamic routing protocol. |
| Options | <p><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.</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 | 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 | <p>show route receive-protocol bgp on page 452</p> <p>show route receive-protocol bgp extensive on page 452</p> <p>show route receive-protocol bgp extensive on page 453</p> <p>show route receive-protocol bgp detail (Layer 2 VPN) on page 454</p> <p>show route receive-protocol bgp extensive (Layer 2 VPN) on page 454</p> <p>show route receive-protocol bgp (Layer 3 VPN) on page 455</p> <p>show route receive-protocol bgp detail (Layer 3 VPN) on page 455</p> <p>show route receive-protocol bgp extensive (Layer 3 VPN) on page 456</p> |
| Output Fields | Table 115 on page 450 describes the output fields for the show route receive-protocol command. Output fields are listed in the approximate order in which they appear. |

Table 115: 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 |
| <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 ■ holddown (routes in that are pending state before being declared inactive) ■ hidden (the routes are not used because of a routing policy) | All levels |

Table 115: show route receive-protocol Output Fields (continued)

| Field Name | Field Description | Level of Output |
|-------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| Prefix | Destination prefix. | none brief |
| MED | Multiple exit discriminator value included in the route. | none brief |
| <i>destination-prefix</i> (entry, announced) | Destination prefix. The entry value is the number of routes for this destination, and the announced value is the number of routes being announced for this destination. | detail extensive |
| Route Distinguisher | 64-bit prefix added to IP subnets to make them unique. | detail extensive |
| Label-Base, range | First label in a block of labels and label block size. A remote PE 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 a Resource Reservation Protocol (RSVP) or a Label Distribution Protocol (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 Lc1pref | Local preference value included in the route. | All levels |
| 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 was originated:</p> <ul style="list-style-type: none"> ■ I—IIGP. ■ 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. | All levels |
| Cluster list | (For router reflected output only) Cluster ID sent by the route reflector. | detail extensive |
| Originator ID | (For router reflected output only) Address of router that originally sent the route to the route reflector. | detail extensive |
| Communities | Community path attribute for the route. See Table 107 on page 367 for all possible values for this field. | detail extensive |

Table 115: show route receive-protocol Output Fields (continued)

| Field Name | Field Description | Level of Output |
|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| 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 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 |

```

show route      user@host> show route receive-protocol bgp 10.255.245.215
receive-protocol bgp
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      user@host> show route receive-protocol bgp 10.255.245.63 extensive
receive-protocol bgp
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      user@host> show route receive-protocol bgp 207.17.136.192 table inet.0
receive-protocol bgp 66.117.68.0/24 extensive
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
detail (Layer 2 VPN)**

```
user@host> show route receive-protocol bgp 10.255.14.171 detail
inet.0: 68 destinations, 68 routes (67 active, 0 holddown, 1 hidden)
Prefix          Nexthop          MED      Lclpref AS path
inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
frame-vpn.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0
hidden)
Prefix          Nexthop          MED      Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 1 announced)
  Route Distinguisher: 10.255.245.35:1
  Label-base : 800000, range : 4, status-vector : 0x0
  Nexthop: 10.255.245.35
  Localpref: 100
  AS path: I
  Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
control flags: 0, mtu: 0
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 0 announced)
  Route Distinguisher: 10.255.245.35:1
  Label-base : 800000, range : 4, status-vector : 0x0
  Nexthop: 10.255.245.35
  Localpref: 100
  AS path: I
  Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
control flags:0, mtu: 0
```

**show route
receive-protocol bgp
extensive (Layer 2 VPN)**

```
user@host> show route receive-protocol bgp 10.255.14.171 extensive
inet.0: 68 destinations, 68 routes (67 active, 0 holddown, 1 hidden)
Prefix          Nexthop          MED      Lclpref AS path
inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
frame-vpn.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 1 announced)
  Route Distinguisher: 10.255.245.35:1
  Label-base : 800000, range : 4, status-vector : 0x0
  Nexthop: 10.255.245.35
  Localpref: 100
  AS path: I
  Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
control flags:0, mtu: 0
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 0 announced)
  Route Distinguisher: 10.255.245.35:1
  Label-base : 800000, range : 4, status-vector : 0x0
  Nexthop: 10.255.245.35
  Localpref: 100
  AS path: I
  Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
control flags:0, mtu: 0
```

**show route
receive-protocol bgp
(Layer 3 VPN)**

```
user@host> show route receive-protocol bgp 10.255.14.171
inet.0: 33 destinations, 33 routes (32 active, 0 holddown, 1 hidden)
Prefix          Nexthop          MED    Lclpref AS path
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
VPN-A.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
10.255.14.175/32 10.255.14.171          100 2 I
10.255.14.179/32 10.255.14.171          2    100 I
VPN-B.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
10.255.14.175/32 10.255.14.171          100 2 I
10.255.14.177/32 10.255.14.171          100 I
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
mpls.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
bgp.l3vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
10.255.14.171:300:10.255.14.177/32
                  10.255.14.171          100 I
10.255.14.171:100:10.255.14.179/32
                  10.255.14.171          2    100 I
10.255.14.171:200:10.255.14.175/32
                  10.255.14.171          100 2 I
```

**show route
receive-protocol bgp
detail (Layer 3 VPN)**

```
user@host> show route receive-protocol bgp 10.255.14.174 detail
inet.0: 16 destinations, 17 routes (15 active, 0 holddown, 1 hidden)
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
vpna.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
* 10.49.0.0/30 (1 entry, 1 announced)
  Route Distinguisher: 10.255.14.176:2
  VPN Label: 101264
  Nexthop: 10.255.14.174
  Localpref: 100
  AS path: I
  Communities: target:200:100
  AttrSet AS: 100
    Localpref: 100
    AS path: I
* 10.255.14.172/32 (1 entry, 1 announced)
  Route Distinguisher: 10.255.14.176:2
  VPN Label: 101280
  Nexthop: 10.255.14.174
  Localpref: 100
  AS path: I
  Communities: target:200:100
  AttrSet AS: 100
    Localpref: 100
    AS path: I
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
* 10.255.14.174:2:10.49.0.0/30 (1 entry, 0 announced)
  Route Distinguisher: 10.255.14.174:2
  VPN Label: 101264
  Nexthop: 10.255.14.174
  Localpref: 100
  AS path: I
  Communities: target:200:100
  AttrSet AS: 100
```

```

        Localpref: 100
        AS path: I
* 10.255.14.174:2:10.255.14.172/32 (1 entry, 0 announced)
    Route Distinguisher: 10.255.14.174:2
    VPN Label: 101280
    Nexthop: 10.255.14.174
    Localpref: 100
    AS path: I
    Communities: target:200:100
    AttrSet AS: 100
        Localpref: 100
        AS path: I
inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

```

**show route
receive-protocol bgp
extensive (Layer 3 VPN)**

```

user@host> show route receive-protocol bgp 10.255.245.63 extensive
inet.0: 244 destinations, 244 routes (243 active, 0 holddown, 1 hidden)
  Prefix          Nexthop          MED      Lclpref AS path
  1.1.1.0/24 (1 entry, 1 announced)
    Nexthop: 10.0.50.3
    Localpref: 100
    AS path: I <Originator>
    Cluster list: 10.2.3.1
    Originator ID: 10.255.245.45
  165.3.0.0/16 (1 entry, 1 announced)
    Nexthop: 111.222.5.254
    Localpref: 100
    AS path: I <Originator>
    Cluster list: 10.2.3.1
    Originator ID: 10.255.245.68
  165.4.0.0/16 (1 entry, 1 announced)
    Nexthop: 111.222.5.254
    Localpref: 100
    AS path: I <Originator>
    Cluster list: 10.2.3.1
    Originator ID: 10.255.245.45
  195.1.2.0/24 (1 entry, 1 announced)
    Nexthop: 111.222.5.254
    Localpref: 100
    AS path: I <Originator>
    Cluster list: 10.2.3.1
    Originator ID: 10.255.245.68
inet.2: 63 destinations, 63 routes (63 active, 0 holddown, 0 hidden)
  Prefix          Nexthop          MED      Lclpref AS path
inet.3: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
  Prefix          Nexthop          MED      Lclpref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
  Prefix          Nexthop          MED      Lclpref AS path
mpls.0: 48 destinations, 48 routes (48 active, 0 holddown, 0 hidden)

```

show route resolution

Syntax show route resolution
 <brief | detail | extensive | summary>
 <index *index*>
 <logical-system (all | *logical-system-name*)>
 <*prefix*>
 <table *routing-table-name*>
 <unresolved>

Release Information Command introduced before JUNOS Release 7.4.

Description Display the entries in the next-hop resolution database. This database provides for recursive resolution of next hops through other prefixes in the routing table.

Options none—Display standard information about all entries in the next-hop resolution database on all logical systems.

brief | detail | extensive | summary—(Optional) Display the specified level of output.

index *index*—(Optional) Show the index of the resolution tree.

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

prefix network/destination-prefix—(Optional) Display database entries for the specified address.

table *routing-table-name*—(Optional) Display information about a particular routing table (for example, inet.0) where policy-based export is currently enabled. (For information about the different types of routing tables, see the *JUNOS Routing Protocols Configuration Guide*.)

unresolved—(Optional) Display routes that could not be resolved.

Required Privilege Level view

List of Sample Output show route resolution detail on page 458
 show route resolution summary on page 459
 show route resolution unresolved on page 459

Output Fields Table 116 on page 458 describes the output fields for the **show route resolution** command. Output fields are listed in the approximate order in which they appear.

Table 116: 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 via 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. |

```

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 user@host> show route resolution summary
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 user@host> show route resolution unresolved
unresolved Tree Index 1
vt-3/2/0.32769.0 /16
Protocol Nexthop: 10.255.71.238 Push 800000
Indirect nexthop: 0 -
vt-3/2/0.32772.0 /16
Protocol Nexthop: 10.255.70.103 Push 800008
Indirect nexthop: 0 -
Tree Index 2

```

show route snooping

| | |
|-----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | <pre>show route snooping <brief detail extensive terse> <all> <best address/prefix> <exact address> <range prefix-range> <summary> <table table-name></pre> |
| Release Information | Command introduced in JUNOS Release 8.5. |
| Description | Display the entries in the routing table that were learned from snooping. |
| 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>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 460 |
| Output Fields | For information about output fields, see the show route command (Table 101 on page 331), the show route detail command (Table 104 on page 360), the show route extensive command (Table 110 on page 380), or the show route terse command (Table 118 on page 483). |
| show route snooping detail | <pre>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</pre> |

```

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)

224.0.0.0.0.0.0.0/24 (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 | <pre>show route source-gateway address <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)></pre> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | <p>Display the entries in the routing table that were learned from a particular address. The Source field in the show route detail command output lists the source for each route, if known.</p> |
| Options | <p><i>address</i>—IP address of the system.</p> <p><i>brief detail extensive terse</i>—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to <i>brief</i>.</p> <p><i>logical-system (all <i>logical-system-name</i>)</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 source-gateway on page 468</p> <p>show route source-gateway detail on page 469</p> <p>show route source-gateway extensive on page 471</p> |
| Output Fields | For information about output fields, see the show route command (Table 101 on page 331), the show route detail command (Table 104 on page 360), the show route extensive command (Table 110 on page 380), or the show route terse command (Table 118 on page 483). |
| show route source-gateway | <pre>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.12vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden) Restart Complete + = Active Route, - = Last Active, * = Both</pre> |

```

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      user@host> show route source-gateway 10.255.70.103 extensive
source-gateway inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
extensive       Restart Complete

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
10.255.70.103:1:3:1/96 (1 entry, 1 announced)
  *BGP      Preference: 170/-101
            Route Distinguisher: 10.255.70.103:1
            Next-hop reference count: 7
            Source: 10.255.70.103
            Protocol next hop: 10.255.70.103
            Indirect next hop: 2 no-forward
            State: <Secondary Active Int Ext>
            Local AS: 69 Peer AS: 69
            Age: 12:15:24 Metric2: 1
            Task: BGP_69.10.255.70.103+179
            Announcement bits (1): 0-green-l2vpn
            AS path: I
            Communities: target:11111:1 Layer2-info: encaps:VPLS,
            control flags:, mtu: 0
            Label-base: 800008, range: 8
            Localpref: 100
            Router ID: 10.255.70.103
            Primary Routing Table bgp.l2vpn.0

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:2:3:1/96 (1 entry, 1 announced)
  *BGP      Preference: 170/-1
            Route Distinguisher: 10.255.70.103:2
            Next-hop reference count: 7
            Source: 10.255.70.103
            Protocol next hop: 10.255.70.103
            Indirect next hop: 2 no-forward
            State: <Secondary Active Int Ext>

```

```

Local AS:    69 Peer AS:    69
Age: 12:15:24 Metric2: 1
Task: BGP_69.10.255.70.103+179
Announcement bits (1): 0-red-l2vpn
AS path: I
Communities: target:11111:2 Layer2-info: encaps:VPLS,
control flags:Site-Down, mtu: 0
Label-base: 800016, range: 8
Localpref: 0
Router ID: 10.255.70.103
Primary Routing Table bgp.l2vpn.0

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:1:3:1/96 (1 entry, 0 announced)
  *BGP Preference: 170/-101
    Route Distinguisher: 10.255.70.103:1
    Next-hop reference count: 7
    Source: 10.255.70.103
    Protocol next hop: 10.255.70.103
    Indirect next hop: 2 no-forward
    State: <Active Int Ext>
    Local AS:    69 Peer AS:    69
    Age: 12:15:24 Metric2: 1
    Task: BGP_69.10.255.70.103+179
    AS path: I
    Communities: target:11111:1 Layer2-info: encaps:VPLS,
    control flags:, mtu: 0
    Label-base: 800008, range: 8
    Localpref: 100
    Router ID: 10.255.70.103
    Secondary Tables: green.l2vpn.0
    Indirect next hops: 1
      Protocol next hop: 10.255.70.103 Metric: 2
      Indirect next hop: 2 no-forward
      Indirect path forwarding next hops: 1
    Next hop:      via so-0/3/0.0 weight 0x1
      10.255.70.103/32 Originating RIB: inet.3
      Metric: 2 Node path count: 1
      Forwarding nexthops: 1
      Nexthop: via so-0/3/0.0

10.255.70.103:2:3:1/96 (1 entry, 0 announced)
  *BGP Preference: 170/-1
    Route Distinguisher: 10.255.70.103:2
    Next-hop reference count: 7
    Source: 10.255.70.103
    Protocol next hop: 10.255.70.103
    Indirect next hop: 2 no-forward
    State: <Active Int Ext>
    Local AS:    69 Peer AS:    69
    Age: 12:15:24 Metric2: 1
    Task: BGP_69.10.255.70.103+179
    AS path: I
    Communities: target:11111:2 Layer2-info: encaps:VPLS,
    control flags:Site-Down,
    mtu: 0
    Label-base: 800016, range: 8
    Localpref: 0
    Router ID: 10.255.70.103

```

```
Secondary Tables: red.12vpn.0
Indirect next hops: 1
  Protocol next hop: 10.255.70.103 Metric: 2
  Indirect next hop: 2 no-forward
  Indirect path forwarding next hops: 1
Next hop:          via so-0/3/0.0 weight 0x1
10.255.70.103/32 Originating RIB: inet.3
Metric: 2          Node path count: 1
Forwarding nexthops: 1
  Nexthop: via so-0/3/0.0
```

show route summary

| | |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show route summary <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display summary statistics about the entries in the routing table. |
| Options | <p>none—Display summary statistics about the entries in the routing table on all logical systems.</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 summary on page 475 |
| Output Fields | Table 117 on page 474 lists the output fields for the show route summary command. Output fields are listed in the approximate order in which they appear. |

Table 117: show route summary Output Fields

| Field Name | Field Description |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>routing-table-name</i> | Name of the routing table (for example, inet.0). |
| 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 not used because of routing policy. |
| Direct | Routes on the directly connected network. |
| Local | Local routes. |
| <i>protocol-name</i> | Name of the protocol from which the route was learned. For example, OSPF, RSVP, and Static. |


```

show route summary  user@host> show route summary
Autonomous system number: 69
Router ID: 10.255.71.52

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 table

| | |
|----------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show route table <i>routing-table-name</i> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display the route entries in a particular routing table. |
| Options | <p><i>routing-table-name</i>—Display information about a particular routing table (for example, inet.0) where policy-based export is currently enabled. (For information about the different types of routing tables, see the <i>JUNOS Routing Protocols Configuration Guide</i>.)</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 Topics | show route summary |
| List of Sample Output | <p>show route table bgp.l2.vpn on page 476</p> <p>show route table bgp.l3vpn.0 on page 477</p> <p>show route table bgp.l3vpn.0 detail on page 477</p> <p>show route table inet.0 on page 478</p> <p>show route table inet6.0 on page 479</p> <p>show route table inet6.3 on page 479</p> <p>show route table l2circuit.0 on page 479</p> <p>show route table mpls on page 479</p> <p>show route table mpls extensive on page 480</p> <p>show route table mpls.0 on page 480</p> <p>show route table vpls_1 detail on page 480</p> <p>show route table vpn-a on page 481</p> <p>show route table vpn-a.mdt.0 on page 481</p> <p>show route table VPN-AB.inet.0 on page 481</p> |
| Output Fields | For information about output fields, see the show route command (Table 101 on page 331), the show route detail command (Table 104 on page 360), the show route extensive command (Table 110 on page 380), or the show route terse command (Table 118 on page 483). |
| show route table bgp.l2.vpn | <pre> user@host> show route table bgp.l2.vpn bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden) + = Active Route, - = Last Active, * = Both 192.168.24.1:1:4:1/96 *[BGP/170] 01:08:58, localpref 100, from 192.168.24.1 AS path: I > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am </pre> |

```

show route table      user@host> show route table bgp.l3vpn.0
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      user@host> show route table bgp.l3vpn.0 detail
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 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 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 l2circuit.0

```

user@host> show route table l2circuit.0
l2circuit.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.1.1.195:NoCtrlWord:1:1:Local/96
                  *[L2CKT/7] 00:50:47
                  > via so-0/1/2.0, Push 100049
                  via so-0/1/3.0, Push 100049
10.1.1.195:NoCtrlWord:1:1:Remote/96
                  *[LDP/9] 00:50:14
                  Discard
10.1.1.195:CtrlWord:1:2:Local/96
                  *[L2CKT/7] 00:50:47
                  > via so-0/1/2.0, Push 100049
                  via so-0/1/3.0, Push 100049
10.1.1.195:CtrlWord:1:2:Remote/96
                  *[LDP/9] 00:50:14
                  Discard

```

show route table mpls

```

user@host> show route table mpls
mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                  *[MPLS/0] 00:13:55, metric 1
                  Receive
1                  *[MPLS/0] 00:13:55, metric 1
                  Receive
2                  *[MPLS/0] 00:13:55, metric 1
                  Receive

```

```

1024                *[VPN/0] 00:04:18
                    to table red.inet.0, Pop

show route table mpls extensive
user@host> show route table mpls extensive
100000 (1 entry, 1 announced)
TSI:
KRT in-kernel 100000 /36 -> {so-1/0/0.0}
    *LDP      Preference: 9
              Next hop: via so-1/0/0.0, selected
              Pop
              State: <Active Int>
              Age: 29:50      Metric: 1
              Task: LDP
              Announcement bits (1): 0-KRT
              AS path: I
              Prefixes bound to route: 10.0.0.194/32

show route table mpls.0
user@host> show route table mpls.0
mpls.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                *[MPLS/0] 00:45:09, metric 1
                  Receive
1                *[MPLS/0] 00:45:09, metric 1
                  Receive
2                *[MPLS/0] 00:45:09, metric 1
                  Receive
100000           *[L2VPN/7] 00:43:04
                  > via so-0/1/0.1, Pop
100001           *[L2VPN/7] 00:43:03
                  > via so-0/1/0.2, Pop      Offset: 4
100002           *[LDP/9] 00:43:22, metric 1
                  via so-0/1/2.0, Pop
                  > via so-0/1/3.0, Pop
100002(S=0)      *[LDP/9] 00:43:22, metric 1
                  via so-0/1/2.0, Pop
                  > via so-0/1/3.0, Pop
100003           *[LDP/9] 00:43:22, metric 1
                  > via so-0/1/2.0, Swap 100002
                  via so-0/1/3.0, Swap 100002
100004           *[LDP/9] 00:43:16, metric 1
                  via so-0/1/2.0, Swap 100049
                  > via so-0/1/3.0, Swap 100049
so-0/1/0.1       *[L2VPN/7] 00:43:04
                  > via so-0/1/2.0, Push 100001, Push 100049(top)
                  via so-0/1/3.0, Push 100001, Push 100049(top)
so-0/1/0.2       *[L2VPN/7] 00:43:03
                  via so-0/1/2.0, Push 100000, Push 100049(top) Offset: -4
                  > via so-0/1/3.0, Push 100000, Push 100049(top) Offset: -4

show route table 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+179
AS path: I
Communities: Layer2-info: encaps:VPLS, control flags:Site-Down
Label-base: 800000, range: 8, status-vector: 0xFF
```

```
show route table vpn-a user@host> show route table vpn-a
vpn-a.12vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

+ = Active Route, - = Last Active, * = Both
192.168.16.1:1:1:1/96
    *[VPN/7] 05:48:27
    Discard
192.168.24.1:1:2:1/96
    *[BGP/170] 00:02:53, localpref 100, from 192.168.24.1
    AS path: I
    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
192.168.24.1:1:3:1/96
    *[BGP/170] 00:02:53, localpref 100, from 192.168.24.1
    AS path: I
    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
```

```
show route table vpn-a.mdt.0 user@host> show route table vpn-a.mdt.0
vpn-a.mdt.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:1:0:10.255.14.216:232.1.1.1/144
    *[MVPN/70] 01:23:05, metric2 1
    Indirect
1:1:1:10.255.14.218:232.1.1.1/144
    *[BGP/170] 00:57:49, localpref 100, from 10.255.14.218
    AS path: I
    > via so-0/0/0.0, label-switched-path r0e-to-r1
1:1:2:10.255.14.217:232.1.1.1/144
    *[BGP/170] 00:57:49, localpref 100, from 10.255.14.217
    AS path: I
    > via so-0/0/1.0, label-switched-path r0-to-r2
```

```
show route table VPN-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 terse

| | |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show route terse <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display a high-level summary of the routes in the routing table. |
| Options | <p>none—Display a high-level summary of the routes in the routing table on all logical systems.</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 485 |
| Output Fields | Table 118 on page 483 describes the output fields for the show route terse command. Output fields are listed in the approximate order in which they appear. |

Table 118: 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> | <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) ■ hidden (routes that are not used because of a routing policy) |
| <i>route key</i> | <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. |
| Destination | Destination of the route. |

Table 118: show route terse Output Fields (*continued*)

| Field Name | Field Description |
|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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 Software 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 | <p>First metric value in the route. For routes learned from BGP, this is the MED metric.</p> |
| Metric 2 | <p>Second metric value in the route. For routes learned from BGP, this is the IGP metric.</p> |
| Next hop | <p>Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.</p> |
| 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 was originated:</p> <ul style="list-style-type: none"> ■ I—IGP. ■ E—EGP. ■ ?—Incomplete; typically, the AS path was aggregated. |

```

show route terse user@host> show route terse
inet.0: 12 destinations, 12 routes (11 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf Metric 1    Metric 2    Next hop      AS path
* 0.0.0.0/0        S   5                      >111.222.5.254
* 1.0.0.1/32       D   0                      >at-5/3/0.0
* 1.0.0.2/32       L   0                      Local
* 12.12.12.21/32   L   0                      Reject
* 13.13.13.13/32   D   0                      >t3-5/2/1.0
* 13.13.13.14/32   L   0                      Local
* 13.13.13.21/32   L   0                      Local
* 13.13.13.22/32   D   0                      >t3-5/2/0.0
  127.0.0.1/32     D   0                      >lo0.0
* 111.222.5.0/24   D   0                      >fxp0.0
* 111.222.5.81/32  L   0                      Local
* 224.0.0.5/32     O  10                      1      MultiRecv

```


Chapter 11

RIP Operational Mode Commands

Table 119 on page 487 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 119: 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” on page 321.

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>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Clear Routing Information Protocol (RIP) general statistics. |
| 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 | clear |
| Related Topics | show rip general-statistics |
| List of Sample Output | clear rip general-statistics on page 488 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| 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>)> < <i>neighbor</i> > |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Clear Routing Information Protocol (RIP) statistics. |
| Options | <p>none—Reset RIP counters for all neighbors for all routing instances on all logical systems.</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><i>neighbor</i>—(Optional) Clear RIP statistics for the specified neighbor only.</p> |
| Required Privilege Level | clear |
| Related Topics | show rip statistics |
| List of Sample Output | clear rip statistics on page 489 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| 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>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display brief Routing Information Protocol (RIP) statistics. |
| Options | none—Display brief RIP statistics on all logical systems. 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 Topics | clear rip general-statistics |
| List of Sample Output | show rip general-statistics on page 490 |
| Output Fields | Table 120 on page 490 lists the output fields for the show rip general-statistics command. Output fields are listed in the approximate order in which they appear. |

Table 120: 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. |

```

show rip      user@host> show rip general-statistics
general-statistics
RIPv2 I/O info:
  bad msgs      :      0
  no rcv intf   :      0
  curr memory   :      0
  max memory    :      0

```


show rip neighbor

| | |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show rip neighbor <instance (all <i>instance-name</i>)> <logical-system (all <i>logical-system-name</i>)> <name> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display information about Routing Information Protocol (RIP) neighbors. |
| Options | <p>none—Display information about all RIP neighbors for all instances on all logical systems.</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 492 |
| Output Fields | Table 121 on page 491 lists the output fields for the show rip neighbor command. Output fields are listed in the approximate order in which they appear. |

Table 121: show rip neighbor Output Fields

| Field Name | Field Description |
|---------------------|----------------------------------------------------------------------------------------------------------|
| Neighbor | Name of RIP neighbor. |
| State | State of the connection: Up or Dn (Down). |
| Source Address | Source address. |
| Destination Address | Destination address. |
| 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. |

show rip neighbor user@host> **show rip neighbor**

| Neighbor | 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 statistics

| | |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show rip statistics <instance (all <i>instance-name</i>)> <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display Routing Information Protocol (RIP) statistics about messages sent and received on an interface, as well as information received from advertisements from other routers. |
| Options | <p>none—Display RIP statistics for all routing instances on all logical systems.</p> <p>instance (all <i>instance-name</i>)—(Optional) Display RIP statistics 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> |
| Required Privilege Level | view |
| Related Topics | clear rip statistics |
| List of Sample Output | show rip statistics on page 494 |
| Output Fields | Table 122 on page 493 lists the output fields for the show rip statistics command. Output fields are listed in the approximate order in which they appear. |

Table 122: 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. ■ holddown—Hold-down interval, 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. ■ restart—Graceful restart status. Displayed when RIP is or has been in the process of graceful restart. |
| <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. ■ timeout—Timeout interval, in seconds. ■ update interval—Number of seconds since last update. |

Table 122: show rip statistics Output Fields (continued)

| Field Name | Field Description |
|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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 that were 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. |

```

show rip statistics  user@host> show rip statistics so-0/0/0.0
RIP info: port 520; update interval: 30s; holddown 180s; timeout 120s
restart in progress: restart time 60s; restart will complete in 55s
    rts learned  rts held down  rqsts dropped  resps dropped
              0              0              0              0
so-0/0/0.0: 0 routes learned; 501 routes advertised
Counter              Total    Last 5 min  Last minute
-----
Updates Sent          0         0         0
Triggered Updates Sent 0         0         0
Responses Sent        0         0         0
Bad Messages          0         0         0
RIPv1 Updates Received 0         0         0
RIPv1 Bad Route Entries 0         0         0
RIPv1 Updates Ignored  0         0         0
RIPv2 Updates Received 0         0         0
RIPv2 Bad Route Entries 0         0         0
RIPv2 Updates Ignored  0         0         0
Authentication Failures 0         0         0
RIP Requests Received  0         0         0
RIP Requests Ignored   0         0         0

```

Chapter 12

RIPng Operational Mode Commands

Table 123 on page 495 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 123: RIPng Operational Mode Commands

| Task | Command |
|-----------------------------|--------------------------------|
| Clear general statistics. | clear ripng general-statistics |
| Clear statistics. | clear ripng statistics |
| Display general statistics. | show ripng general-statistics |
| Display RIPng neighbors. | show ripng neighbor |
| Display statistics. | show ripng statistics |



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” on page 321.

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>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Clear Routing Information Protocol next generation (RIPng) general statistics. |
| 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 | clear |
| Related Topics | clear ripng general-statistics |
| List of Sample Output | clear ripng general-statistics on page 496 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| 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)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Clear Routing Information Protocol next-generation (RIPng) statistics. |
| Options | <p>none—Reset RIPng counters for all neighbors for all routing instances on all logical systems.</p> <p>instance—(Optional) Reset RIPng counters for the specified instance.</p> <p>name—(Optional) Reset RIPng counters for the specified neighbor.</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 | clear |
| Related Topics | clear ripng statistics |
| List of Sample Output | clear ripng statistics on page 497 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| 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>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display general Routing Information Protocol next-generation (RIPng) statistics. |
| Options | none—Display general RIPng statistics on all logical systems. 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 Topics | clear ripng general-statistics |
| List of Sample Output | show ripng general-statistics on page 498 |
| Output Fields | Table 124 on page 498 lists the output fields for the show ripng general-statistics command. Output fields are listed in the approximate order in which they appear. |

Table 124: 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. |

```

show ripng      user@host> show ripng general-statistics
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> > |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display information about Routing Information Protocol next-generation (RIPng) neighbors. |
| Options | none—Display information about all RIPng neighbors on all logical systems. logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system. <i>name</i> —(Optional) Display detailed information about a specific RIPng neighbor. |
| Required Privilege Level | view |
| List of Sample Output | show ripng neighbor on page 499 |
| Output Fields | Table 125 on page 499 lists the output fields for the show ripng neighbor command. Output fields are listed in the approximate order in which they appear. |

Table 125: 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 Mode | Send options: broadcast, multicast, none, version 1, or yes. |
| Receive Mode | 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. |

| | | | | | | |
|----------------------------|-------|--------------------------------|--------|---------|------|----------|
| show ripng neighbor | | user@host> show ripng neighbor | | | | |
| | | | Source | Dest | | In |
| Neighbor | State | Address | | Address | Send | Recv 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>)> <name> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display Routing Information Protocol next generation (RIPng) statistics about messages sent and received on an interface, as well as information received from advertisements from other routers. |
| Options | <p>none—Display RIPng statistics for all neighbors on all logical systems.</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 Topics | clear ripng statistics |
| List of Sample Output | show ripng statistics on page 501 |
| Output Fields | Table 126 on page 500 lists the output fields for the show ripng statistics command. Output fields are listed in the approximate order in which they appear. |

Table 126: 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 RIP. ■ holddown—Hold-down interval, 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. ■ restart—Graceful restart status. Displayed when RIPng is or has been in the process of graceful restart. |
| <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. ■ timeout—Timeout interval, in seconds. ■ update interval—Number of seconds since last update. |

Table 126: show ripng statistics Output Fields *(continued)*

| Field Name | Field Description |
|-------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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. |

```

show ripng statistics  user@host> show ripng statistics
RIPng info: port 521; holddown 120s;
      rts learned  rts held down  rqsts dropped  resps dropped
                0             0             0             0

so-0/1/3.0: 0 routes learned; 1 routes advertised; timeout 180s; update interval
20s
Counter                Total    Last 5 min  Last minute
-----
Updates Sent            934         16         4
Triggered Updates Sent    1          0          0
Responses Sent           0          0          0
Bad Messages             0          0          0
Updates Received          0          0          0
Bad Route Entries         0          0          0
Updates Ignored           0          0          0
RIPng Requests Received   0          0          0
RIPng Requests Ignored    0          0          0

```


Part 2

Policy Framework

- Firewall Filter Operational Mode Commands on page 505
- Forwarding Operational Mode Commands on page 515
- Routing Policy Operational Mode Commands on page 531

Chapter 13

Firewall Filter Operational Mode Commands

Table 127 on page 505 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot firewall filters. Commands are listed in alphabetical order.

Table 127: Firewall Filter Operational Mode Commands

| Task | Command |
|------------------------------------------------|-----------------------------------|
| Clear firewall filter counters. | clear firewall |
| Operational statistics for firewall filters. | show firewall |
| Firewall filter log information. | show firewall log |
| Prefix-action statistics for firewall filters. | show firewall prefix-action-stats |
| Counters for policers. | show policer |



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 *counter-name* | filter *filter-name* | logical-system *logical-system-name*)

Release Information Command introduced before JUNOS Release 7.4.
The logical-system option introduced in JUNOS Release 9.3.

Description Clear statistics about configured firewall filters.



NOTE: The clear firewall command cannot be used to clear the Routing Engine filter counters on a backup Routing Engine that is enabled for GRES.

Options all—Clear the packet and byte counts for all filters.

counter *counter-name*—Clear the packet and byte counts for a filter counter that has been configured with the counter firewall filter action.

filter *filter-name*—Clear the packet and byte counts for the specified firewall filter.

logical-system *logical-system-name*—Clear the packet and byte counts for the specified logical system.

Required Privilege Level clear

Related Topics show firewall

List of Sample Output clear firewall all on page 506

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

clear firewall all user@host> clear firewall all

show firewall

| | |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show firewall <filter <i>filter-name</i> > <counter <i>counter-name</i> > <logical-system (<i>logical-system-name</i> all)> |
| Release Information | Command introduced before JUNOS Release 7.4. The logical-system option introduced in JUNOS Release 9.3. |
| Description | Display statistics about configured firewall filters. |
| Options | <i>filter-name</i> —(Optional) Name of a configured filter. counter <i>counter-name</i> —(Optional) Name of a filter counter. 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 Topics | clear firewall |
| List of Sample Output | show firewall filter on page 509 show firewall filter (Dynamic Input Filter) on page 509 show firewall (Logical Systems) on page 509 |
| Output Fields | Table 128 on page 507 lists the output fields for the show firewall command. Output fields are listed in the approximate order in which they appear. |

Table 128: 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>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.</p> <p>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).</p> |
| 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 128: show firewall Output Fields (*continued*)

| Field Name | Field Description |
|------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Policers | Display policer information: <ul style="list-style-type: none">■ Name—Name of policer.■ Packets—Number of packets that matched the filter term under which the policer action is specified. This is only the number of out-of-spec packet counts, not all packets policed by the policer. |

```

show firewall filter user@host> show firewall filter test
Filter: test
Counters:
Name                               Bytes          Packets
Counter-1                         0              0
Counter-2                         0              0
Policers:
Name                               Packets
Policer-1                        0

show firewall filter user@host> show firewall filter dfwd-ge-5/0/0.1-in
(Dynamic Input Filter) Filter: dfwd-ge-5/0/0.1-in
Counters:
Name                               Bytes          Packets
cl-ge-5/0/0.1-in                  0              0

show firewall (Logical user@host>show firewall
Systems)
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 log

| | |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show firewall log <detail> <interface <i>interface-name</i> > <logical-system (<i>logical-system-name</i> all)> |
| Release Information | Command introduced before JUNOS Release 7.4. logical-system option introduced in JUNOS Release 9.3. |
| Description | Display log information about firewall filters. |
| Options | detail—(Optional) Display detailed information. interface <i>interface-name</i> —(Optional) Display log information about a specific interface. logical-system (<i>logical-system-name</i> all)—(Optional) Perform this operation on all logical systems or on a particular system. |
| Required Privilege Level | view |
| List of Sample Output | show firewall log detail on page 511 |
| Output Fields | Table 129 on page 510 lists the output fields for the show firewall log command. Output fields are listed in the approximate order in which they appear. |

Table 129: show firewall log Output Fields

| Field Name | Field Description |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Time of Log | Time that the event occurred. |
| Filter | Name of a filter that has been configured with the filter statement at the [edit firewall] hierarchy level. <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. |
| 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 , ipip , ospf , pim , rsvp , tcp , or udp . |
| Packet length | Length of the packet. |

Table 129: show firewall log Output Fields *(continued)*

| Field Name | Field Description |
|---------------------|----------------------------------------|
| Source address | Packet's source address. |
| Destination address | Packet's destination address and port. |

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 Release 7.4. logical-system option introduced in JUNOS Release 9.3. |
| Description | Display prefix action statistics about configured firewall filters. |
| 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 Topics | clear firewall |
| List of Sample Output | show firewall prefix-action-stats on page 512 |
| Output Fields | Table 130 on page 512 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 130: 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). |

show firewall prefix-action-stats user@host> show firewall prefix-action-stats filter test prefix-action act1
Filter: __ls2/test

show policer

| | |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show policer < <i>policer-name</i> > |
| Release Information | Command introduced before JUNOS Release 7.4. |
| 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. <i>policer-name</i> —(Optional) Display the number of policed packets for the specified policer. |
| Required Privilege Level | view |
| List of Sample Output | show policer on page 513 show policer (Aggregate Policar) on page 513 |
| Output Fields | Table 131 on page 513 lists the output fields for the show policer command. Output fields are listed in the approximate order in which they appear. |

Table 131: show policer Output Fields

| Field Name | Field Description |
|------------|-----------------------------------------------------------|
| Name | Name of the policer. |
| Packets | Total number of packets policed by the specified policer. |

| | |
|-----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| show policer | <pre> user@host> show policer Policers: Name Packets __default_arp_policer__ 0 fe-1/2/0.0-out-policer 3496661237 fe-1/2/1.0-out-policer 3432710964 </pre> |
| show policer (Aggregate Policar) | <pre> user@host> show policer Policers: Name 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 </pre> |

| | |
|----------------------|---|
| __policer_tmpl__-fc1 | 0 |
| __policer_tmpl__-fc2 | 0 |
| __policer_tmpl__-fc3 | 0 |

Chapter 14

Forwarding Operational Mode Commands

Table 132 on page 515 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot forwarding options. Commands are listed in alphabetical order.

Table 132: 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 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 statistics collected by the UDP forwarding process. | <code>show helper statistics</code> |

clear dhcp relay binding

| | |
|-------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | <pre>clear dhcp relay binding <(all <i>ip-address</i> <i>mac-address</i>)> <interface <i>interface-name</i>> <logical-system <i>logical-system-name</i>> <routing-instance <i>routing-instance-name</i>></pre> |
| Release Information | <p>Command introduced in JUNOS Release 8.3.</p> <p>all and interface options added in JUNOS Release 8.4.</p> |
| Description | Clear the binding state of a Dynamic Host Configuration Protocol (DHCP) client from the client table. |
| Options | <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><i>ip-address</i>—(Optional) IP address of the DHCP client.</p> <p><i>mac-address</i>—(Optional) MAC address of the DHCP client.</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 |
| List of Sample Output | <p>clear dhcp relay binding on page 516</p> <p>clear dhcp relay binding all on page 516</p> <p>clear dhcp relay binding interface on page 517</p> |
| Output Fields | For information about output fields, see the show dhcp relay binding command (Table 133 on page 521). |
| clear dhcp relay binding | <p>The following sample output displays the address bindings in the DHCP client table before and after the clear dhcp relay binding command is issued.</p> <pre>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 user@host> clear dhcp relay binding 100.20.32.1 user@host> show dhcp relay binding</pre> |
| clear dhcp relay binding all | <pre>user@host> clear dhcp relay binding all</pre> |

```
clear dhcp relay binding user@host> clear dhcp relay binding interface fe-0/0/2  
interface
```

clear dhcp relay statistics

| | |
|------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | clear dhcp relay statistics <logical-system <i>logical-system-name</i> > <routing-instance <i>routing-instance-name</i> > |
| Release Information | Command introduced in JUNOS Release 8.3. |
| Description | Clear all Dynamic Host Configuration Protocol (DHCP) 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 cleared 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 cleared for the default routing instance.</p> |
| Required Privilege Level | view |
| List of Sample Output | clear dhcp relay statistics on page 518 |
| Output Fields | For information about output fields, see the show dhcp relay statistics command (Table 134 on page 526). |
| clear dhcp relay statistics | <p>The following sample output displays the DHCP relay statistics before and after the clear dhcp relay statistics command is issued.</p> <pre> 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 </pre> |

| | |
|----------------|---|
| DHCPRELEASE | 0 |
| DHCPREQUEST | 0 |
| Messages sent: | |
| BOOTREPLY | 0 |
| DHCPOFFER | 0 |
| DHCPACK | 0 |
| DHCPNAK | 0 |

clear helper statistics

| | |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | clear helper statistics |
| Release Information | Command introduced before JUNOS 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 Topics | show helper statistics |
| List of Sample Output | clear helper statistics on page 520 |
| Output Fields | See show helper statistics for an explanation of output fields. |
| clear helper statistics | <p>The following sample output displays statistics counters before and after the clear helper statistics command is issued:</p> <pre> 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 </pre> |

show dhcp relay binding

| | |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show dhcp relay binding <detail> <interface <i>interface-name</i> > < <i>ip-address</i> <i>mac-address</i> > <logical-system <i>logical-system-name</i> > <routing-instance <i>routing-instance-name</i> > |
| Release Information | Command introduced in JUNOS Release 8.3. interface and <i>mac-address</i> options added in JUNOS Release 8.4. |
| Description | Display the address bindings in the Dynamic Host Configuration Protocol (DHCP) client table. |
| Options | detail—(Optional) Display detailed client binding information. interface <i>interface-name</i> —(Optional) Perform this operation on the specified interface. <i>ip-address</i> —(Optional) IP address of the DHCP client. <i>mac-address</i> —(Optional) MAC address of the DHCP client. logical-system <i>logical-system-name</i> —(Optional) Perform this operation on the specified logical system. routing-instance <i>routing-instance-name</i> —(Optional) Perform this operation on the specified routing instance. |
| Required Privilege Level | view |
| Related Topics | clear dhcp relay binding |
| List of Sample Output | show dhcp relay binding on page 522 show dhcp relay binding detail on page 523 show dhcp relay binding interface on page 523 show dhcp relay binding ip-address on page 523 show dhcp relay binding ip-address detail on page 523 show dhcp relay binding mac-address on page 523 |
| Output Fields | Table 133 on page 521 lists the output fields for the show dhcp relay binding command. Output fields are listed in the approximate order in which they appear. |

Table 133: show dhcp relay binding Output Fields

| Field Name | Field Description | Level of Output |
|-------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------|-----------------|
| <i>number</i> clients, (<i>number</i> bound, <i>number</i> selecting, <i>number</i> renewing, <i>number</i> rebinding) | Summary counts of the total number of DHCP clients and the number of DHCP clients in each state | detail none |

Table 133: show dhcp relay binding Output Fields (continued)

| Field Name | Field Description | Level of Output |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------|
| IP address | IP address of the DHCP client. | All levels |
| Hardware address | Hardware address of the DHCP client. | All levels |
| 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 |
| State | State of the DHCP relay address binding table on the DHCP client: <ul style="list-style-type: none"> ■ init—Initial state. ■ reboot—Client sends DHCP DISCOVER request. ■ select—Client receives offers from DHCP servers. ■ request—Client requests a DHCP server. ■ add—Client is in process of being added. ■ delete—Client is in process of being deleted. ■ bound—Client has active IP address lease. ■ renew—Client sends request to renew IP address lease. ■ rebind—Client broadcasts request to renew IP address lease. | detail |
| Active binding information | Information about active IP address binding: <ul style="list-style-type: none"> ■ IP address—IP address of the DHCP client. ■ Hardware address—Hardware address of the DHCP client. ■ Request received on—(detail level only) Interface on which the client request was received. ■ relayed by—(detail level only) IP address on which the client request was relayed. | All levels (unless specified otherwise) when command includes <i>ip-address</i> or <i>mac-address</i> value |
| Lease information | Information about the client's IP address lease: <ul style="list-style-type: none"> ■ Type—Type of IP address lease; always DHCP. ■ Obtained at—Date and time at which the client's IP address lease was obtained. ■ Expires at—Date and time at which the client's IP address lease expires. ■ State—(detail level only) State of the DHCP relay address binding table on the DHCP client. | All levels (unless specified otherwise) when command includes <i>ip-address</i> or <i>mac-address</i> value |

show dhcp relay binding user@host> **show dhcp relay binding**

10 clients, (10 bound, 0 selecting, 0 renewing, 0 rebinding)

| | | | |
|--------------|-------------------|--------|-------------------------|
| IP address | Hardware address | Type | Lease expires at |
| 100.20.0.152 | 00:10:95:01:00:01 | active | 2007-02-15 05:25:47 PST |
| 100.20.0.153 | 00:10:95:01:00:02 | active | 2007-02-15 05:25:48 PST |
| 100.20.0.154 | 00:10:95:01:00:03 | active | 2007-02-15 05:25:49 PST |

| | | | |
|--------------|-------------------|--------|-------------------------|
| 100.20.0.155 | 00:10:95:01:00:04 | active | 2007-02-15 05:25:50 PST |
| 100.20.0.156 | 00:10:95:01:00:05 | active | 2007-02-15 05:25:51 PST |
| 100.20.0.157 | 00:10:95:01:00:06 | active | 2007-02-15 05:25:52 PST |
| 100.20.0.158 | 00:10:95:01:00:07 | active | 2007-02-15 05:25:53 PST |
| 100.20.0.159 | 00:10:95:01:00:08 | active | 2007-02-15 05:25:54 PST |
| 100.20.0.160 | 00:10:95:01:00:09 | active | 2007-02-15 05:25:55 PST |
| 100.20.0.161 | 00:10:95:01:00:0a | active | 2007-02-15 05:25:56 PST |

show dhcp relay binding detail user@host> **show dhcp relay binding detail**

10 clients, (10 bound, 0 selecting, 0 renewing, 0 rebinding)

| IP address | Hardware address | Type | Lease expires | State |
|--------------|-------------------|--------|-------------------------|-------|
| 100.20.0.152 | 00:10:95:01:00:01 | active | 2007-02-15 05:25:47 PST | bound |
| 100.20.0.153 | 00:10:95:01:00:02 | active | 2007-02-15 05:25:48 PST | bound |
| 100.20.0.154 | 00:10:95:01:00:03 | active | 2007-02-15 05:25:49 PST | bound |
| 100.20.0.155 | 00:10:95:01:00:04 | active | 2007-02-15 05:25:50 PST | bound |
| 100.20.0.156 | 00:10:95:01:00:05 | active | 2007-02-15 05:25:51 PST | bound |
| 100.20.0.157 | 00:10:95:01:00:06 | active | 2007-02-15 05:25:52 PST | bound |
| 100.20.0.158 | 00:10:95:01:00:07 | active | 2007-02-15 05:25:53 PST | bound |
| 100.20.0.159 | 00:10:95:01:00:08 | active | 2007-02-15 05:25:54 PST | bound |
| 100.20.0.160 | 00:10:95:01:00:09 | active | 2007-02-15 05:25:55 PST | bound |
| 100.20.0.161 | 00:10:95:01:00:0a | active | 2007-02-15 05:25:56 PST | bound |

show dhcp relay binding interface user@host> **show dhcp relay binding interface fe-0/0/2**

1 clients, (1 bound, 0 selecting, 0 renewing, 0 rebinding)

| 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 ip-address user@host> **show dhcp relay binding 100.20.4.138**

Active binding information:

| | |
|------------------|-------------------|
| IP address | 100.20.4.138 |
| Hardware address | 00:10:95:01:00:01 |

Lease information:

| | |
|-------------|-------------------------|
| Type | DHCP |
| Obtained at | 2007-02-14 06:37:33 PST |
| Expires at | 2007-02-15 06:37:33 PST |

show dhcp relay binding ip-address detail user@host> **show dhcp relay binding 100.20.32.1 detail**

Active binding information:

| | |
|---------------------|------------------------------------|
| IP address | 100.20.32.1 |
| Hardware address | 90:00:00:01:00:01 |
| Request received on | fe-0/0/2.0, relayed by 100.20.32.2 |

Lease information:

| | |
|-------------|-------------------------|
| Type | DHCP |
| Obtained at | 2007-01-29 15:43:27 EST |
| Expires at | 2007-01-29 15:53:27 EST |
| State | rebind |

show dhcp relay binding mac-address user@host> **show dhcp relay binding 90:00:00:01:00:01**

Active binding information:

| | |
|------------------|-------------------|
| IP address | 100.20.32.1 |
| Hardware address | 90:00:00:01:00:01 |

Lease information:

| | |
|------|------|
| Type | DHCP |
|------|------|

| | |
|-------------|-------------------------|
| Obtained at | 2007-01-17 11:28:47 PST |
| Expires at | 2007-01-17 11:38:47 PST |

show dhcp relay statistics

| | |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show dhcp relay statistics <logical-system <i>logical-system-name</i> > <routing-instance <i>routing-instance-name</i> > |
| Release Information | Command introduced in JUNOS Release 8.3. |
| Description | Display Dynamic Host Configuration Protocol (DHCP) 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 Topics | clear dhcp relay statistics |
| List of Sample Output | show dhcp relay statistics on page 526 |
| Output Fields | Table 134 on page 526 lists the output fields for the show dhcp relay statistics command. Output fields are listed in the approximate order in which they appear. |

Table 134: 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 |

```

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 | 44 |
| DHCPOFFER | 11 |
| DHCPACK | 11 |
| DHCPNAK | 11 |

show helper statistics

| | |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show helper statistics |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Show statistics collected by the UDP forwarding process. |
| Options | This command has no options. |
| Required Privilege Level | view |
| Related Topics | clear helper statistics |
| List of Sample Output | show helper statistics on page 529 |
| Output Fields | Table 135 on page 528 lists the output fields for the <code>show helper statistics</code> command. Output fields are listed in the approximate order in which they appear. |

Table 135: show helper statistics Output Fields

| Field Name | Field Description |
|------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| domain | Statistics for Domain Name System (DNS) forwarding: <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 135: 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. |

```

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
tftp: Due to no interface in fud database: 0
tftp: Due to no matching routing instance: 0
tftp: Due to an error during packet read: 0
tftp: Due to an error during packet send: 0
tftp: Due to invalid server address: 0
tftp: Due to no valid local address: 0
tftp: Due to no route to server/client: 0

```


Chapter 15

Routing Policy Operational Mode Commands

Table 136 on page 531 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot routing policy filters. Commands are listed in alphabetical order.

Table 136: 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> > |
| Release Information | Command introduced before JUNOS Release 7.4. |
| 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 Topics | show policy damping |
| List of Sample Output | <p>show policy on page 532</p> <p>show policy policy-name on page 533</p> <p>show policy (Multicast Scoping) on page 533</p> |
| Output Fields | Table 137 on page 532 lists the output fields for the show policy command. Output fields are listed in the approximate order in which they appear. |

Table 137: 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. |

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 | show policy conditions <condition-name> <detail> |
| Release Information | Command introduced in JUNOS Release 9.0. |
| Description | 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. |
| 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> |
| Required Privilege Level | view |
| List of Sample Output | show policy conditions detail on page 534 |
| Output Fields | Table 138 on page 534 lists the output fields for the show policy conditions command. Output fields are listed in the approximate order in which they appear. |

Table 138: 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 |
| If-route-exists conditions | List of conditions configured to look for a route in the specified table. | All levels |

```

show policy conditions    user@host> show policy conditions detail
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
cond2

test policy

| | |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | <code>test policy <i>policy-name</i> <i>prefix</i></code> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Test a policy configuration to determine which prefixes match routes in the routing table. |
| Options | <p><i>policy-name</i>—Name of a policy.</p> <p><i>prefix</i>—Destination prefix to match.</p> |
| 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 Border Gateway Protocol (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 Topics | show policy damping |
| List of Sample Output | test policy on page 536 |
| Output Fields | For information about output fields, see the show route command (Table 101 on page 331), the show route detail command (Table 104 on page 360), the show route extensive command (Table 110 on page 380), or the show route terse command (Table 118 on page 483). The last line of output is unique to the test policy command. It provides the policy name and number of prefixes accepted and rejected. |

```

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 539
- MPLS Operational Mode Commands on page 565
- RSVP Operational Mode Commands on page 609

Chapter 16

LDP Operational Mode Commands

Table 139 on page 539 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 139: 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” on page 321.

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>)> <neighbor> |
| Description | Tear down Label Distribution Protocol (LDP) neighbor connections. |
| Options | <p>none—Tear down connections with all LDP neighbors for all routing instances on all logical systems.</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 Topics | show ldp neighbor |
| List of Sample Output | clear ldp neighbor on page 540 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| clear ldp neighbor | user@host> clear ldp neighbor |

clear ldp session

| | |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | clear ldp session <destination> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Clear Label Distribution Protocol (LDP) sessions. |
| Options | <p>none—Clear LDP sessions for all destinations for all routing instances on all logical systems.</p> <p><i>destination</i>—(Optional) Clear an LDP session for the specified destination (IP address).</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 Topics | show ldp session |
| List of Sample Output | clear ldp session on page 541 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| 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 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 on all logical systems.</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 Topics | show ldp statistics show ldp traffic-statistics |
| List of Sample Output | clear ldp statistics on page 542 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| clear ldp statistics | user@host> clear ldp statistics |

show ldp database

| | |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | 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> > |
| Release Information | Command introduced before JUNOS 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 on all logical systems.</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 | <p>show ldp database on page 545</p> <p>show ldp database l2circuit detail on page 545</p> <p>show ldp database session on page 545</p> |
| Output Fields | Table 140 on page 543 describes the output fields for the show ldp database command. Output fields are listed in the approximate order in which they appear. |

Table 140: 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 140: 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 <i>L2CKT control word status encapsulation-type vc-number</i>, for example, <i>L2CKT CtlfWord FRAME RELAY VC 2</i></p> <ul style="list-style-type: none"> ■ <i>control-word-status</i>—Displays whether the use of the control word has been negotiated for this virtual circuit: <ul style="list-style-type: none"> ■ NoCtrlWord ■ CtrlWord ■ <i>encapsulation-type</i>—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 |
| 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 |
| State | <p>State of the label binding:</p> <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 |

```

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 Release 7.4. |
| Description | Display information about configured Label Distribution Protocol (LDP) forwarding equivalence class (FEC) filters. |
| Options | <p><i>fec</i>—(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 546 |
| Output Fields | Table 141 on page 546 lists the output fields for the show ldp fec-filters command. Output fields are listed in the approximate order in which they appear. |

Table 141: 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. |

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> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS 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 on all logical systems.</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> |
| Required Privilege Level | view |
| List of Sample Output | show ldp interface extensive on page 548 |
| Output Fields | Table 142 on page 547 describes the output fields for the show ldp interface command. Output fields are listed in the approximate order in which they appear. |

Table 142: 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 |
| 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 |

```
show ldp interface   user@host> show ldp interface extensive  
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 Release 7.4. neighbor-address option added in JUNOS Release 8.5. |
| Description | Display Label Distribution Protocol (LDP) neighbor information. |
| Options | <p>none—Display standard information about LDP neighbors for all routing instances on all logical systems.</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 Topics | clear ldp neighbor |
| List of Sample Output | show ldp neighbor extensive on page 550 |
| Output Fields | Table 143 on page 549 describes the output fields for the show ldp neighbor command. Output fields are listed in the approximate order in which they appear. |

Table 143: 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 143: 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 |

```

show ldp neighbor      user@host> show ldp neighbor extensive
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 show ldp path
 <brief | detail | extensive>
 <destination>
 <instance *instance-name*>
 <logical-system (all | *logical-system-name*)>

Release Information Command introduced before JUNOS Release 7.4.

Description Display Label Distribution Protocol (LDP) label-switched paths (LSPs).

Options none—Display standard information about all LDP LSPs for all routing instances on all logical systems.

brief | detail | extensive—(Optional) Display the specified level of output.

destination—(Optional) Restrict the output to entries that match the specified destination prefix.

instance *instance-name*—(Optional) Display information for the specified routing instance only.

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

Required Privilege Level view

List of Sample Output show ldp path extensive on page 552

Output Fields Table 144 on page 551 describes the output fields for the **show ldp path** command. Output fields are listed in the approximate order in which they appear.

Table 144: 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. |
| <i>route</i> | 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 144: show ldp path Output Fields *(continued)*

| Field Name | Field Description |
|--------------|-----------------------------------|
| Global label | MPLS label that is used globally. |

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 | show ldp route <brief detail extensive> <destination> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS 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 on all logical systems.</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 <i>instance-name</i>—(Optional) Display entries 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 |
| List of Sample Output | show ldp route detail on page 554 |
| Output Fields | Table 145 on page 553 describes the output fields for the show ldp route command. Output fields are listed in the approximate order in which they appear. |

Table 145: show ldp route Output Fields

| Field Name | Field Description |
|-------------------------|---------------------------------------------------------------------------------|
| Destination | Destination prefix. |
| Next-hop intf | Interface that is the next hop to the destination prefix. |
| Next-hop address | IP address of the next hop. |
| Bound to outgoing label | The route has been bound to LSPs with the label being distributed for that LSP. |

```

show ldp route detail  user@host> show ldp route detail
Destination                Next-hop intf  Next-hop address
10.10.255.1/32             so-2/3/0
*10.10.255.3/32           so-1/0/0      10.10.1.3
    Bound to outgoing label 100001
*10.10.255.1/32           so-2/3/0
10.10.255.4/32            so-0/0/0      192.168.1.213
*10.10.255.4/32           so-0/0/0      192.168.1.213
    Bound to outgoing label 100002
10.10.255.6/32            so-0/0/0      192.168.1.215
*10.10.255.6/32           so-0/0/0      192.168.1.215
    Bound to outgoing label 100000
*10.10.255.2/32
    Bound to outgoing label 3
0.0.0.0/0                 so-0/0/0      192.168.1.254
10.10.255.3/32           so-1/0/0      10.10.1.3

```


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 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 on all logical systems.</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 Topics | clear ldp session |
| List of Sample Output | show ldp session extensive on page 557 |
| Output Fields | Table 146 on page 555 describes the output fields for the show ldp session command. Output fields are listed in the approximate order in which they appear. |

Table 146: show ldp session Output Fields

| Field Name | Field Description | Level of Output |
|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| Address | Transport address of the session. | to be provided |
| 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. | to be provided |
| Connection | TCP connection state: Closed , Opening , or Open . | to be provided |
| Hold time | Time remaining until the session will be closed, in seconds. | to be provided |
| Session ID | LDP identifiers of the peers of this session. | to be provided |
| Next keepalive | Time until next keepalive is sent, in seconds. | detail extensive |

Table 146: show ldp session Output Fields (continued)

| Field Name | Field Description | Level of Output |
|-------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| Active | Whether the local router is playing the active role in the session and during session establishment. | detail extensive |
| Maximum PDU | Maximum PDU 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 <code>keepalive-timeout</code> statement configured at the <code>[edit protocols ldp]</code> 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 |
| 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 |
| restart complete in <i>nnn</i> msec | Amount of time (in milliseconds) remaining until graceful restart is declared complete. | detail extensive |

Table 146: show ldp session Output Fields (continued)

| Field Name | Field Description | Level of Output |
|-----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| 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 |
| Message type | <p>Type of message being sent.</p> <ul style="list-style-type: none"> ■ Total—Messages sent and received during the lifetime of the session. ■ Last 5 seconds—Messages sent and received during the current session. | extensive |

```

show ldp session extensive  user@host> show ldp session extensive
                             Address: 10.255.70.103, State: Operational, Connection: Open, Hold time: 20
                             Session ID: 10.255.71.52:0--10.255.70.103:0
                             Next keepalive in 0 seconds
                             Active, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
                             Keepalive interval: 10, Connect retry interval: 1
                             Local address: 10.255.71.52, Remote address: 10.255.70.103
                             Up for 00:01:40
                             Last down 00:01:41 ago; Reason: received notification from peer
                             Number of session flaps: 3
                             Local - Restart: disabled, Helper mode: enabled
                             Remote - Restart: disabled, Helper mode: enabled

```

Local maximum recovery time: 240000 msec

Next-hop addresses received:

so-2/0/0.0

11.2.3.1

11.2.4.1

11.2.5.1

11.1.2.2

192.168.70.103

10.255.70.103

Queue depth: 0

| Message type | Total | | Last 5 seconds | |
|------------------|-------|----------|----------------|----------|
| | Sent | Received | Sent | Received |
| Initialization | 4 | 4 | 0 | 0 |
| Keepalive | 114 | 114 | 0 | 0 |
| Notification | 0 | 3 | 0 | 0 |
| Address | 4 | 4 | 0 | 0 |
| Address withdraw | 0 | 0 | 0 | 0 |
| Label mapping | 8 | 8 | 0 | 0 |
| Label request | 0 | 0 | 0 | 0 |
| Label withdraw | 0 | 0 | 0 | 0 |
| Label release | 0 | 0 | 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 Release 7.4. |
| Description | Display Label Distribution Protocol (LDP) statistics. |
| Options | none—Display LDP statistics for all routing instances on all logical systems. instance <i>instance-name</i> —(Optional) Display information for the specified routing instance only. logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | view |
| Related Topics | clear ldp statistics |
| List of Sample Output | show ldp statistics on page 560 |
| Output Fields | Table 147 on page 559 lists the output fields for the show ldp statistics command. Output fields are listed in the approximate order in which they appear. |

Table 147: 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 147: 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. |
| Event type | LDP events and errors. |
| Total | Total number of each event or error. |
| Last 5 seconds | Number of each event or error in the last 5 seconds. |

```

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>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display Label Distribution Protocol (LDP) traffic statistics. |
| Options | <p>none—Display LDP traffic statistics for all routing instances on all logical systems.</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> |
| 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 Topics | clear ldp statistics |
| List of Sample Output | show ldp traffic-statistics on page 563 |
| Output Fields | Table 148 on page 562 lists the output fields for the show ldp traffic-statistics command. Output fields are listed in the approximate order in which they appear. |

Table 148: 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. |
| 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. |
| 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. |

| | | | | | |
|----------------------------------------|----------------------------------------|---------|---------|-------|--------|
| <div>show ldp traffic-statistics</div> | 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 |

Chapter 17

MPLS Operational Mode Commands

Table 149 on page 565 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 149: 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> |
| Display configured LSPs on this router, as well as all ingress, transit, and egress LSPs. | <code>show mpls lsp</code> |
| Display configured named paths that are used in dynamic MPLS. | <code>show mpls path</code> |

Table 149: MPLS Operational Mode Commands (*continued*)

| Task | Command |
|---------------------------------------------------------------------|--------------------------------|
| Display entries in the traffic engineering database. | <code>show ted database</code> |
| Display current traffic engineering database links. | <code>show ted link</code> |
| Display protocols contributing to the traffic engineering database. | <code>show ted protocol</code> |



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” on page 321.

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 | *logical-system-name*)>
 <name *name*>
 <optimize | optimize-aggressive>
 <path *regular-expression*>
 <statistics>

Release Information Command introduced before JUNOS Release 7.4.

Description Release the routes and states associated with Multiprotocol Label Switching (MPLS) label-switched paths (LSPs), and start new LSPs.



CAUTION: This command disconnects existing Resource Reservation Protocol (RSVP) sessions on the ingress router. 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 none—Reset and restart all LSPs that originated from this router; that is, all LSPs for which this router is the ingress router. Depending on the number of LSPs involved, it might take a while to restart all the LSPs.

autobandwidth—(Optional) Clear LSP autobandwidth counters.

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

name *name*—(Optional) Reset and restart the specified LSP or group of LSPs. You can include wildcard characters in the interface name, as described in the *JUNOS Network Interfaces Configuration Guide*.

optimize | optimize-aggressive—(Optional) Run nonpreemptive optimization or aggressive optimization computation now.

path *regular-expression*—(Optional) Clear the specific LSP path matching the specified regular expression.

statistics—(Optional) Clear LSP statistics.

Required Privilege Level clear

Related Topics show mpls lsp
 show rsvp session

List of Sample Output clear mpls lsp on page 568

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

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>)> < <i>lsp-name</i> > |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Manually trigger a bandwidth allocation adjustment for active label-switched paths (LSPs). |
| Options | <p>none—Manually trigger a bandwidth allocation adjustment for all active LSP paths on all logical systems.</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) 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 |
| List of Sample Output | request mpls lsp adjust-auto-bandwidth on page 569 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| request mpls lsp adjust-auto-bandwidth | user@host> request mpls lsp adjust-auto-bandwidth |

show connections

Syntax 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 | *logical-system-name*)>
 <name>
 <status>

Release Information Command introduced before JUNOS Release 7.4.

Description Display information about the configured circuit cross-connect (CCC) connections.

Options none—Display the standard level of output for all configured CCC connections on all logical systems.

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.

history—(Optional) Display information about connection history.

labels—(Optional) Display labels used for transmit and receive.

status—(Optional) Display information about the connection and interface status.

all—(Optional) Display all connections.

interface-switch—(Optional) Display interface switch connections only.

lsp-switch—(Optional) Display LSP switch connections only.

p2mp-receive-switch—(Optional) Display point-to-multipoint LSP to local interfaces switch connections only.

p2mp-transmit-switch—(Optional) Display local interface to point-to-multipoint LSP switch connections only.

remote-interface-switch—(Optional) Display remote interface switch connections only.

down | up | up-down—(Optional) Display nonoperational, operational, or both kinds of connections.

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.

Required Privilege Level view

Output Fields Table 150 on page 571 describes the output fields for the `show connections` command. Output fields are listed in the approximate order in which they appear.

Table 150: 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. |

show connections

```

user@switch> show connections
CCC and TCC connections [Link Monitoring On]
  Legend for status (St)      Legend for connection types
  UN -- uninitialized         if-sw: interface switching
  NP -- not present           rmt-if: remote interface switching
  WE -- wrong encapsulation    lsp-sw: LSP switching
  DS -- disabled
  Dn -- down
  -> -- only outbound conn is up
  <- -- only inbound conn is up
  Up -- operational
  RmtDn -- remote CCC down
  Restart -- restarting
  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 | if-sw | Up | Aug 5 15:39:15 | 1 |
| so-1/0/2.0 | intf | Up | | |
| t1-0/1/2.0 | intf | Up | | |
| SW-db | rmt-if | Restart | | 0 |
| 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 Release 7.4. |
| Description | Display Multiprotocol Label Switching (MPLS) peer and traffic engineering link information. |
| Options | This command has no options. |
| Required Privilege Level | view |
| Related Topics | show link-management peer show link-management routing show link-management statistics show link-management te-link |
| List of Sample Output | show link-management on page 575 |
| Output Fields | Table 151 on page 573 describes the output fields for the show link-management command. Output fields are listed in the approximate order in which they appear. |

Table 151: 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. |
| 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. |

Table 151: show link-management Output Fields (continued)

| Field Name | Field Description |
|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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 . |
| LSP-name | LSP name. |

```

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 Release 7.4. |
| 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 Topics | show link-management show link-management routing show link-management statistics show link-management te-link |
| List of Sample Output | show link-management peer on page 577 |
| Output Fields | Table 152 on page 576 describes the output fields for the show link-management peer command. Output fields are listed in the approximate order in which they appear. |

Table 152: 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 router 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 router 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. |
| 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. |

Table 152: show link-management peer Output Fields *(continued)*

| Field Name | Field Description |
|------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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. |

```

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> >> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| 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>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 Topics | show link-management show link-management peer show link-management statistics show link-management te-link |
| List of Sample Output | show link-management routing on page 580 |
| Output Fields | Table 153 on page 578 describes the output fields for the show link-management routing command. Output fields are listed in the approximate order in which they appear. |

Table 153: 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. |
| 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. |

Table 153: show link-management routing Output Fields *(continued)*

| Field Name | Field Description |
|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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. |

```

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              Active

Peer name: __rpd:so-0/2/1.0, System identifier: 2147483652
  State: Down, Control address: (null)
    Control-channel          State
    so-0/2/1.0              Active

...

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: 0bps, 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: 0bps, 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: 0bps, Maximum bandwidth: 155.52Mbps,
  Total bandwidth: 155.52Mbps, Available bandwidth: 155.52Mbps

...

Resource: falsp-bd, Type: LSP, State: Dn System identifier: 2147483652,
Total bandwidth: 0bps, 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 Release 8.0. |
| 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 Topics | show link-management show link-management peer show link-management routing show link-management te-link |
| List of Sample Output | show link-management statistics on page 582 |
| Output Fields | Table 154 on page 581 describes the output fields for the show link-management statistics command. Output fields are listed in the approximate order in which they appear. |

Table 154: 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. |
| 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. |

Table 154: show link-management statistics Output Fields (*continued*)

| Field Name | Field Description |
|-----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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. |

```
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 <name <i>name</i> > |
| Release Information | Command introduced before JUNOS Release 7.4. |
| 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. name <i>name</i> —(Optional) Display information for the specified traffic-engineered link only. |
| Required Privilege Level | view |
| Related Topics | show link-management show link-management peer show link-management routing show link-management statistics |
| List of Sample Output | show link-management te-link on page 584 |
| Output Fields | Table 155 on page 583 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 155: 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. |
| 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. |

Table 155: show link-management te-link Output Fields (continued)

| Field Name | Field Description |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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. |

```

show link-management user@host> show link-management te-link
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  e2e1sp-bf

```

show mpls admin-groups

| | |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show mpls admin-groups <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display information about configured Multiprotocol Label Switching (MPLS) administrative groups. |
| Options | <p>none—Display information about the configured MPLS administrative groups on all logical systems.</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 mpls admin-groups on page 585 |
| Output Fields | Table 156 on page 585 describes the output fields for the show mpls admin-groups command. Output fields are listed in the approximate order in which they appear. |

Table 156: show mpls admin-groups Output Fields

| Field Name | Field Description |
|------------|---------------------------------------------|
| Group | Name of the administrative group. |
| Bit index | Value assigned to the administrative group. |

```

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>)> <lsp-name> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display Multiprotocol Label Switching (MPLS) label-switched path (LSP) call admission control (CAC) information. |
| Options | <p>none—Display CAC information for all LSPs on all logical systems.</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 587 |
| Output Fields | Table 157 on page 586 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 157: 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 path name. |
| Neighbor address | Neighbor address from which CAC and bandwidth booking are configured for Layer 2 circuits. |
| Circuit | Interface name and circuit information. |
| Primary | LSP's primary standby path. |
| Standby | LSP's secondary standby path. |
| VC bandwidth | Bandwidth constraints associated with a Layer 2 circuit route. |


```

show mpls      user@host# show mpls call-admission-control
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>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display Multiprotocol Label Switching (MPLS) Constrained Shortest Path First (CSPF) statistics. |
| Options | none—Display MPLS CSFP statistics on all logical systems. 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 589 |
| Output Fields | Table 158 on page 588 describes the output fields for the show mpls cspf command. Output fields are listed in the approximate order in which they appear. |

Table 158: 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. |
| 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. |

Table 158: show mpls cspf Output Fields *(continued)*

| Field Name | Field Description |
|--------------|-----------------------------------------------------------------|
| 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. |

```

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>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| 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 on all logical systems. 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 590 |
| Output Fields | Table 159 on page 590 describes the output fields for the show mpls diffserv-te command. Output fields are listed in the approximate order in which they appear. |

Table 159: 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. |

```

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>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display information about Multiprotocol Label Switching (MPLS)-enabled interfaces. |
| Options | none—Display information about MPLS-enabled interfaces on all logical systems. 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 <code>set protocol mpls interface <i>interface-name</i></code> and <code>set interface <i>interface-name</i> unit 0 family mpls</code> statements. |
| Required Privilege Level | view |
| List of Sample Output | show mpls interface on page 591 |
| Output Fields | Table 160 on page 591 describes the output fields for the <code>show mpls interface</code> command. Output fields are listed in the approximate order in which they appear. |

Table 160: 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. |

```

show mpls interface  user@host> show mpls interface
Interface  State      Administrative groups
so-1/0/0.0 Up         Blue Yellow Red

```

show mpls lsp

Syntax show mpls lsp
 <brief | detail | extensive | terse>
 <bidirectional | unidirectional>
 <bypass>
 <descriptions>
 <down | up>
 <logical-system (all | *logical-system-name*)>
 <lsp-type>
 <name *name*>
 <p2mp>
 <statistics>
 <transit>

Release Information Command introduced before JUNOS Release 7.4.

Description Display information about configured and active dynamic Multiprotocol Label Switching (MPLS) label-switched paths (LSPs).

Options none—Display standard information about all configured and active dynamic MPLS LSPs on all logical systems.

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.

bidirectional | unidirectional—(Optional) Display bidirectional or unidirectional LSP information, respectively.

bypass—(Optional) Display LSPs used for protecting other LSPs.

descriptions—(Optional) Display the MPLS label-switched path (LSP) descriptions. To view this information, you must configure the description statement at the [edit protocol mpls lsp] hierarchy level. Only LSPs with a description are displayed. This command is only valid for the ingress router 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 router.
- **ingress**—Sessions that originate from this router.
- **transit**—Sessions that pass through this router.

name *name*—(Optional) Display information about the specified LSP or group of LSPs.

p2mp—(Optional) Display information about point-to-multipoint LSPs.

statistics—(Optional) (Egress and transit routers only.) Display accounting information about LSPs. Statistics are not available for LSPs on the egress router, because the penultimate router in the LSP sets the label to 0. Also, as the packet arrives at the egress router, 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.

transit—(Optional) Display LSPs transiting this router.

Required Privilege Level view

Related Topics clear mpls lsp

List of Sample Output show mpls lsp descriptions on page 597
 show mpls lsp detail on page 597
 show mpls lsp extensive on page 598
 show mpls lsp p2mp on page 598
 show mpls lsp p2mp detail on page 599

Output Fields Table 161 on page 593 describes the output fields for the **show mpls lsp** command. Output fields are listed in the approximate order in which they appear.

Table 161: show mpls lsp Output Fields

| Field Name | Field Description | Level of Output |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| Ingress LSP | Information about LSPs on the ingress router. Each session has one line of output. | All levels |
| Egress LSP | Information about the LSPs on the egress router. 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 routers 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> automatically generated by JUNOS. | 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 router) of the LSP. | detail extensive |
| To | Destination (egress router) of the session. | brief |
| From | Source (ingress router) of the session. | brief detail |

Table 161: show mpls lsp Output Fields (continued)

| Field Name | Field Description | Level of Output |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| 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 |
| P | Path. An asterisk (*) underneath this column indicates that the LSP is a primary path. | brief |
| LSPname | Name of the LSP. | brief detail |
| DiffServInfo | Type of LSP: multiclass LSP (multiclass diffServ-TE LSP) or Differentiated-Services-aware traffic engineering LSP (diffServ-TE LSP). | detail |
| Bypass | (Bypass LSP) Destination address (egress router) 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 routers. | 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 |
| 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 |
| ActivePath | (Ingress LSP) Name of the active path: Primary or Secondary . | detail extensive |
| FastReroute desired | Fast reroute has been requested by the ingress router. | detail |
| Link protection desired | Link protection has been requested by the ingress router. | 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 |
| 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 |

Table 161: show mpls lsp Output Fields (continued)

| Field Name | Field Description | Level of Output |
|----------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| 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 |
| 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 |
| 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 |
| OptimizeTimer | (Ingress LSP) Configured value of the optimize timer, indicating the total amount of time allowed before path reoptimization, in seconds. | detail extensive |

Table 161: show mpls lsp Output Fields (continued)

| Field Name | Field Description | Level of Output |
|--------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| 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 |
| 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 router has a backup path providing the same bandwidth guarantee as the protected LSP for the protected section. ■ 0x08—Node protection. The downstream router has a backup path providing protection against link and node failure on the corresponding path section. If the downstream router 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 |

Table 161: show mpls lsp Output Fields (continued)

| Field Name | Field Description | Level of Output |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------|
| 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 |
| 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) router 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) router or client, interface used to reach this neighbor, and number of packets sent to the downstream router. | detail |
| RESV rcvfrom | Address of the previous-hop (upstream) router or client, interface the neighbor used to reach this router, 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 |

show mpls lsp descriptions

```

user@host> show mpls lsp descriptions
Ingress LSP: 3 sessions
To          LSP name
10.0.0.195   to-sanjose
10.0.0.195   to-sanjose-other-desc
Total 2 displayed, Up 2, Down 0

```

```

Description
to-sanjose-desc
other-desc

```

show mpls lsp detail

```

user@host> show mpls lsp detail
Ingress LSP: 1 sessions

10.255.245.3
From: 10.255.245.5, State: Up, ActiveRoute: 1, LSPname: lsp-ec
ActivePath: long-path (primary)
LoadBalance: Random

```

```

Autobandwidth
MaxBW: 5Mbps
AdjustTimer: 4800 secs AdjustThreshold: 1%
Max AvgBW util: 0bps, Bandwidth Adjustment in 3383 second(s).
Overflow limit: 5, Overflow sample count: 0
Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary long-path State: Up
  SmartOptimizeTimer: 180
  Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 5)
192.168.37.89 S 192.168.37.87 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):
    192.168.37.89 192.168.37.87
Total 1 displayed, Up 1, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

show mpls lsp extensive

```

user@host> show mpls lsp extensive
Ingress LSP: 5 sessions

10.255.71.242
  From: 10.255.71.238, State: Up, ActiveRoute: 1009, LSPname: sample-ccc
  ActivePath: path3 (primary)
  Link protection desired
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary path3 State: Up
  OptimizeTimer: 30
  SmartOptimizeTimer: 180
  Reoptimization in 26 second(s).
  Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 1)
10.35.1.41 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):
    10.35.1.41(Label=3)
10 Dec 8 13:51:58.986 CSPF: computation result ignored
9 Dec 8 13:51:30.547 Record Route: 10.35.1.41(Label=3)
8 Dec 8 13:51:30.547 Up
7 Dec 8 13:51:30.397 Originate make-before-break call
6 Dec 8 13:51:30.397 CSPF: computation result accepted 10.35.1.41
5 Dec 8 13:50:41.467 Selected as active path
4 Dec 8 13:50:41.467 Record Route: 10.35.1.41(Label=3)
3 Dec 8 13:50:41.466 Up
2 Dec 8 13:50:41.371 Originate Call
1 Dec 8 13:50:41.371 CSPF: computation result accepted 10.35.1.41
Created: Fri Dec 8 13:50:40 2006
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 p2mp

```

user@host> show mpls lsp p2mp
Ingress LSP: 2 sessions
P2MP name: p2mp-lsp1, P2MP branch count: 1
To          From          State Rt ActivePath    P    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 | ActivePath | P | 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 path

| | |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show mpls path <path-name> <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display dynamic Multiprotocol Label Switching (MPLS) label-switched paths (LSPs). |
| Options | none—Display standard information about all MPLS LSPs on all logical systems. <i>path-name</i> —(Optional) Display information about the specified LSP only. logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level | view |
| List of Sample Output | show mpls path on page 600 |
| Output Fields | Table 162 on page 600 describes the output fields for the show mpls path command. Output fields are listed in the approximate order in which they appear. |

Table 162: 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 routers that form the LSP. |
| Strict/loose address | Whether the address is configured as a strict or loose address. |

```

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 ted database

| | |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show ted database <brief detail extensive> <logical-system (all <i>logical-system-name</i>)> < <i>system-name</i> > |
| Release Information | Command introduced before JUNOS Release 7.4. |
| 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 on all logical systems.</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 | <p>show ted database brief on page 603</p> <p>show ted database detail <i>system-name</i> on page 604</p> <p>show ted database extensive on page 604</p> |
| Output Fields | Table 163 on page 601 describes the output fields for the show ted database command. Output fields are listed in the approximate order in which they appear. |

Table 163: 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 router 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 router itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode. | extensive |
| 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 |

Table 163: show ted database Output Fields (continued)

| Field Name | Field Description | Level of Output |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| 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 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 163: 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 |

```

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      user@host> show ted database detail merino
detail system-name    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      user@host> show ted database extensive
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>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display Multiprotocol Label Switching (MPLS) traffic engineering database link information. |
| Options | <p>none—Display standard information about traffic engineering database link information on all logical systems.</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 | <p>show ted link brief on page 606</p> <p>show ted link detail on page 606</p> |
| Output Fields | Table 164 on page 605 describes the output fields for the show ted link command. Output fields are listed in the approximate order in which they appear. |

Table 164: 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 router 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 router itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode. | brief |
| hostname | Hostname and address of the node that the link is coming from. An address of .00 indicates that the node is the router itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode. | detail |
| hostname | Hostname and address of the node that the link is going to. An address of .00 indicates that the node is the router 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 router has placed on the link. | All levels |
| Local BW | Amount of bandwidth the local router has placed on the link. | All levels |

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>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| 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 on all logical systems.</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 607 |
| Output Fields | Table 165 on page 607 describes the output fields for the show ted protocol command. Output fields are listed in the approximate order in which they appear. |

Table 165: 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. |

```

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)

```


Chapter 18

RSVP Operational Mode Commands

Table 166 on page 609 summarizes the command-line interface (CLI) commands you can use to monitor Resource Reservation Protocol (RSVP) sessions. Commands are listed in alphabetical order.

Table 166: 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” on page 321.

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 clear rsvp session
 <connection-source *address*>
 <connection-destination *address*>
 <gracefully>
 <logical-system (all | *logical-system-name*)>
 <lsp-id *identifier*>
 <name *name*>
 <optimize-fast-reroute>
 <tunnel-id *identifier*>

Release Information Command introduced before JUNOS Release 7.4.

Description Reset and restart Resource Reservation Protocol (RSVP) sessions.

Options none—Reset and restart all RSVP sessions for which this router is the ingress, transit, or egress router.

connection-source *address*—(Optional) Source address for GMPLS and MPLS LSPs from the RSVP sender template.

connection-destination *address*—(Optional) Destination address for GMPLS and MPLS LSPs from the RSVP sender template.

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 router of the RSVP session and is only valid for nonpacket LSPs.

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

lsp-id *identifier*—(Optional) LSP identifier (source port) for the RSVP sender template.

name *name*—(Optional) Reset and restart the specified RSVP session.

optimize-fast-reroute—(Optional) Begin fast reroute optimization.

tunnel-id *identifier*—(Optional) Tunnel identifier (destination port) for the RSVP session.

Required Privilege Level clear

Related Topics clear mpls lsp
 show rsvp session

List of Sample Output clear rsvp session on page 611

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

clear rsvp session user@host> **clear rsvp session**

clear rsvp statistics

| | |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | clear rsvp statistics <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Clear Resource Reservation Protocol (RSVP) packet and error statistics. |
| Options | none—Clear RSVP packet and error statistics on all logical systems. 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 Topics | show rsvp statistics |
| List of Sample Output | clear rsvp statistics on page 612 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| 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>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| 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 on all logical systems.</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 | <p>show rsvp interface brief on page 615</p> <p>show rsvp interface detail on page 615</p> <p>show rsvp interface extensive on page 616</p> <p>show rsvp interface link-management on page 616</p> |
| Output Fields | Table 167 on page 613 lists the output fields for the show rsvp interface command. Output fields are listed in the approximate order in which they appear. |

Table 167: 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 |
| 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 |

Table 167: show rsvp interface Output Fields (continued)

| Field Name | Field Description | Level of Output |
|------------------------|-----------------------------------------------------------------------------------------------------------------------|-----------------|
| 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 |
| 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 |

Table 167: show rsvp interface Output Fields (continued)

| Field Name | Field Description | Level of Output |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| 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 |
| queue-priority-value | Low, High, None, or Exact. None indicates no rate limiting. Exact indicates the queue transmits at the configured rate only. | extensive |

```

show rsvp interface brief    user@host> show rsvp interface brief
                               RSVP interface: 1 active
                               Active Subscr- Static
                               Interface State resv  option BW      Available  Reserved  Highwater
                               de0.0    Up        1    23%   10Mbps   989.992kbps 1.31Mbps  mark
                               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>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| 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 on all logical 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 rsvp neighbor on page 621 show rsvp neighbor detail on page 621 |
| Output Fields | Table 168 on page 618 lists the output fields for the show rsvp neighbor command. Output fields are listed in the approximate order in which they appear. |

Table 168: show rsvp neighbor Output Fields

| Field Name | Field Description | Level of Output |
|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| RSVP neighbor | Number of neighbors about which the router has learned. 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 Software Release 3.2 or earlier, are not reported as up or down. | All levels |
| 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 Software Release 3.2 or earlier, are not reported as up or down. | detail |

Table 168: show rsvp neighbor Output Fields (continued)

| Field Name | Field Description | Level of Output |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| status | <p>State of the RSVP neighbor:</p> <ul style="list-style-type: none"> ■ Up—Router can detect RSVP Hello messages from the neighbor. ■ Down—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. ■ 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—Router 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 router has received from the neighbor. | detail |
| Remote Instance | Identification provided by the remote router during Hello message exchange. | detail |
| Local Instance | Identification sent to the remote router during Hello message exchange. | detail |
| Refresh reduction | <p>Measure of processing overhead requests of refresh messages. Refresh reduction extensions improve router performance by reducing the process overhead, thus increasing the number of LSPs a router 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 routers. For a detailed explanation of these extensions, see RFC 2961. ■ incomplete—Some RSVP refresh reduction extensions are functional between the two neighboring routers. ■ no operational—Either the refresh reduction feature has been turned off, or the remote router cannot support the refresh reduction extensions. | detail |

Table 168: show rsvp neighbor Output Fields (continued)

| Field Name | Field Description | Level of Output |
|-----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Remote end | Neighboring router's status in regard to refresh reduction: <ul style="list-style-type: none"> ■ enabled—Remote router has requested refresh reduction during RSVP message exchanges. ■ disabled—Remote router does not require refresh reduction. | detail |
| Ack-extension | An RSVP refresh reduction extension: <ul style="list-style-type: none"> ■ enabled—Both local and remote routers support the ack-extension (RFC 2961). ■ disabled—Remote router does not support the ack-extension. | detail |
| Link protection | Status of the MPLS fast reroute mechanism that protects traffic from link failure: <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 | Status of the bypass LSP. It can have the following values: <ul style="list-style-type: none"> ■ does not exist—Bypass LSP is not available. ■ connecting—Router 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 |
| 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 |

```

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 user@host> show rsvp neighbor detail
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 show rsvp session
 <brief | detail | extensive | terse>
 <bidirectional | unidirectional>
 <bypass>
 <down | up>
 <interface *interface-name*>
 <logical-system (all | *logical-system-name*)>
 <lsp-type>
 <name *session-name*>
 <p2mp>
 <session-type>
 <statistics>
 <te-link *te-link*>

Release Information Command introduced before JUNOS Release 7.4.

Description Display information about Resource Reservation Protocol (RSVP) sessions.

Options none—Display standard information about all RSVP sessions on all logical systems.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

bidirectional | unidirectional—(Optional) Display information about bidirectional or unidirectional RSVP sessions only, respectively.

bypass—(Optional) Display RSVP sessions for bypass LSPs.

down | up—(Optional) Display only LSPs that are inactive or active, respectively.

interface *interface-name*—(Optional) Display RSVP sessions for the specified interface only.

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 RSVP sessions with regard to LSPs:

- 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 router.
- *ingress*—Sessions that originate from this router.
- *transit*—Sessions that transit through this router.

statistics—(Optional) Display packet statistics.

te-link te-link—(Optional) Display sessions with reservations on the specified TE link.

Required Privilege Level view

Related Topics clear rsvp session

List of Sample Output show rsvp session on page 626
 show rsvp session statistics on page 626
 show rsvp session detail on page 627
 show rsvp session detail (Path MTU Output Field) on page 627
 show rsvp session detail (GMPLS) on page 627
 show rsvp session extensive on page 628
 show rsvp session p2mp on page 628

Output Fields Table 169 on page 623 describes the output fields for the `show rsvp session` command. Output fields are listed in the approximate order in which they appear.

Table 169: 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 <code>p2mp</code> option is specified). Name of the point-to-multipoint LSP path. | All levels |
| P2MP branch count | (Appears only when the <code>p2mp</code> option is specified). Number of LSPs receiving packets from the point-to-multipoint LSP. | All levels |
| To | Destination (egress router) of the session. | All levels |
| From | Source (ingress router) 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 router) of the LSP. | detail |
| From | Source (ingress router) of the session. | detail |

Table 169: show rsvp session Output Fields (continued)

| Field Name | Field Description | Level of Output |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| 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 routers. LSPpath can also indicate when a graceful LSP deletion has been triggered. | detail |
| Bypass | (Egress router) 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 |
| 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 |

Table 169: show rsvp session Output Fields (continued)

| Field Name | Field Description | Level of Output |
|----------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| 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 |
| FastReroute desired | Fast reroute has been requested by the ingress router. | detail |
| Soft preemption desired | Soft preemption has been requested by the ingress router. | 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 router. | 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 router. | detail extensive |
| 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 router. | detail extensive |
| Type | LSP type: <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 |

Table 169: show rsvp session Output Fields (continued)

| Field Name | Field Description | Level of Output |
|---------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| 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 routers with the <code>allow-fragmentation</code> statement configured at the <code>[edit protocols mpls path-mtu]</code> 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) router or client, interface the neighbor used to reach this router, and number of packets received from the upstream neighbor. | detail |
| Adspec | MTU signaled from the ingress router to the egress router by means of the adspec object. | detail |
| PATH sentto | Address of the next-hop (downstream) router or client, interface used to reach this neighbor (or peer-name in the GMPLS LSP case), and number of packets sent to the downstream router. | detail |
| Explot 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 explot route. Differences indicate that path rerouting has occurred, typically during fast-reroute. | detail |

```

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      user@host> show rsvp session statistics
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: gmpls-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**
10.255.245.13

```

From: 10.255.245.48, LSPstate: Up, ActiveRoute: 0
....
Link protection desired
Type: Link protected LSP, using p2
11 Feb 6 15:24:16 Backup LSP: Call was cleared by RSVP
10 Feb 6 15:24:16 Backup LSP: Session preempted
9 Feb 6 15:24:16 Deleting backup LSP, protected LSP restored
8 Feb 6 15:23:22 Backup LSP: Up 192.168.208.117(Label=3)
7 Feb 6 15:23:22 Backup LSP: Record Route: 192.168.208.117(Label=3)
6 Feb 6 15:23:19 Backup LSP: Explicit Route: wrong delivery
5 Feb 6 15:23:19 Creating backup LSP, link down
4 Feb 6 12:36:03 Link protection up, using p2
3 Feb 6 12:35:56 New bypass p2
2 Feb 6 12:35:47 Bypass state down, p1[2 times]
1 Feb 6 12:35:39 New bypass p1

```

show rsvp session p2mp user@host> **show rsvp session p2mp**

```

Ingress RSVP: 3 sessions
P2MP name: p2mp-lsp1, P2MP branch count: 1
To          From          State Rt Style Labelin Labelout LSPname
10.255.245.34 10.255.245.25 Up      0 1 FF      -      100128 p2mp-branch-1
P2MP name: p2mp-lsp2, P2MP branch count: 1
To          From          State Rt Style Labelin Labelout LSPname
10.255.245.34 10.255.245.25 Up      0 1 FF      -      3 p2mp-st-br1
P2MP name: lsp-a_b, P2MP branch count: 1
Total 2 displayed, Up 2, 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 statistics

| | |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show rsvp statistics <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| Description | Display Resource Reservation Protocol (RSVP) packet and error statistics. |
| Options | none—Display RSVP packet and error statistics on all logical systems. 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 Topics | clear rsvp statistics |
| List of Sample Output | show rsvp statistics on page 631 |
| Output Fields | Table 170 on page 629 describes the output fields for the show rsvp statistics command. Output fields are listed in the approximate order in which they appear. |

Table 170: 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 routers along a path. |
| 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. |

Table 170: show rsvp statistics Output Fields (continued)

| Field Name | Field Description |
|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|
| 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 cksum | The RSVP checksum check failed. |
| Rcv pkt bad format | General packet processing failed because the packet was badly formed. |
| Memory alloc fail | An internal resource failure occurred. |
| No path info | 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. |
| 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 router replies with an ACK object that contains an unknown message ID. This can indicate a message ID handshake problem. |

Table 170: show rsvp statistics Output Fields (continued)

| Field Name | Field Description |
|------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Recv nack | A neighboring router explicitly rejects a message ID in a summary refresh message. This can happen if that neighbor has been rebooted. In this case, the router sends a regular RSVP refresh message to recover the state, and starts 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 router restarts. |
| No TE-link to recv Hop | Counter of packets discarded because a TE link was not found. |

```

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

```

show rsvp version

| | |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show rsvp version <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS Release 7.4. |
| 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 router. |
| Options | none—Display RSVP protocol settings on all logical systems. 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 (Router in Steady State) on page 633 show rsvp version (Router Restarting) on page 633 |
| Output Fields | Table 171 on page 632 describes the output fields for the show rsvp version command. Output fields are listed in the approximate order in which they appear. |

Table 171: 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 . |
| Graceful restart | Status of the graceful restart feature for RSVP on the restarting router: Enabled or Disabled . |
| Restart helper mode | Status of the helper mode feature: Enabled or Disabled . When this feature is enabled, the restarting router 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 router 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 router 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. |

Table 171: show rsvp version Output Fields *(continued)*

| Field Name | Field Description |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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. |
| 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. |

show rsvp version
(Router in Steady State)

```

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  60 seconds
  Graceful restart         Enabled
  Restart helper mode      Enabled
  Restart time             60000 msec

```

show rsvp version
(Router Restarting)

```

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
  Graceful restart:         Disabled
  Restart helper mode:      Enabled
  Maximum helper restart time: 20000 msec
  Maximum helper recovery time: 180000 msec
  Restart time:             0 msec

```


Part 4

Layer 2 Bridging and Switching Operational Mode Commands

- Layer 2 Bridging and Switching Operational Mode Commands on page 637
- Spanning Tree Operational Mode Commands on page 663

Chapter 19

Layer 2 Bridging and Switching Operational Mode Commands

Table 172 on page 637 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 172: 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 | 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> > |
| Release Information | Command introduced in JUNOS 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 638 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| clear bridge mac-table | user@host> clear bridge mac-table |

clear error bpdu

| | |
|-----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | clear error bpdu <interface <i>interface-name</i> > |
| Release Information | Command introduced in JUNOS 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 | interface <i>interface-name</i> —(Optional) Clear the BPDU error condition for the specified interface. |
| Required Privilege Level | clear |
| List of Sample Output | clear error bpdu interface on page 639 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| clear error bpdu interface | user@host> clear error bpdu interface ge-1/1/1 |

clear error mac-rewrite

| | |
|------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | clear error mac-rewrite <interface <i>interface-name</i> > |
| Release Information | Command introduced in JUNOS 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 | interface <i>interface-name</i> —(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 640 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| clear error mac-rewrite interface | user@host> clear error mac-rewrite interface ge-1/0/1 |

show bridge domain

| | | | | |
|----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--|--|--|
| Syntax | show bridge domain <brief detail extensive> <bridge-domain (all <i>domain-name</i>)> <instance <i>instance-name</i> > <operational> | | | |
| Release Information | Command introduced in JUNOS Release 8.4. | | | |
| Description | (MX Series routers only) Display bridge domain information. | | | |
| Options | none—Display information for all bridge domains. brief detail extensive—(Optional) Display the specified level of output. bridge-domain (all <i>domain-name</i>)— (Optional) Display information about all bridge domains or the specified bridge domain. instance <i>instance-name</i> —(Optional) Display information for the specified routing instance. operational—(Optional) Display information for the operational routing instances. | | | |
| Required Privilege Level | view | | | |
| List of Sample Output | show bridge domain on page 641 show bridge domain brief on page 641 show bridge domain detail on page 641 | | | |
| show bridge domain | <pre> user@host> show bridge domain Instance Bridging Domain Type Active Primary Table vs1 vlan100 bridge 2 vs1 bridge.0 bridge 0 vs1 vlan200 bridge vs1 bridge.0 </pre> | | | |
| show bridge domain brief | <pre> user@host> show bridge domain brief Instance Bridging Domain Type Active Primary Table vs1 vlan100 bridge 2 vs1 bridge.0 bridge 0 vs1 vlan200 bridge vs1 bridge.0 </pre> | | | |
| show bridge domain detail | <pre> 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 </pre> | | | |

```
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 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 | <pre>show bridge flood on page 644 show bridge flood brief on page 644 show bridge flood detail on page 644 show bridge flood extensive on page 645</pre> |
| Output Fields | to be provided |

```

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 show bridge mac-table
 <brief | count | detail | extensive>
 <bridge-domain (all | *bridge-domain-name*)>
 <global-count>
 <interface *interface-name*>
 <mac-address>
 <vlan-id (all-vlan | *vlan-id*)>

Release Information Command introduced in JUNOS Release 8.4.

Description (MX Series routers only) Display Layer 2 MAC address information.

Options none—Display all learned Layer 2 MAC address information.

brief | count | detail | extensive—(Optional) Display the specified level of output.

bridge-domain (all | *bridge-domain-name*)—(Optional) Display learned Layer 2 MAC addresses for all bridging domains or for the specified bridging domain.

global-count—(Optional) Display the total number of learned Layer 2 MAC addresses on the system.

instance *instance-name*—(Optional) Display learned Layer 2 MAC addresses for the specified routing instance.

interface *interface-name*—(Optional) Display learned Layer 2 MAC addresses for the specified interface.

mac-address—(Optional) Display the specified learned Layer 2 MAC address information.

vlan-id (all-vlan | *vlan-id*)—(Optional) Display learned Layer 2 MAC addresses for all VLANs or for the specified VLAN.

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 651
 show bridge mac-table brief on page 651
 show brief mac-table count on page 652
 show bridge mac-table detail on page 652

Output Fields Table 173 on page 651 describes the output fields for the show bridge mac-table command. Output fields are listed in the approximate order in which they appear.

Table 173: 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 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. |

show bridge mac-table 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
```

show bridge mac-table brief 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

```

show brief mac-table count user@host> **show bridge mac-table count**
 2 MAC address learned in routing instance vs1 bridge domain vlan100

MAC address count per interface within routing instance:

| Logical interface | MAC count |
|-------------------|-----------|
| ge-11/0/3.0 | 1 |
| ge-11/1/4.100 | 0 |
| ge-11/1/1.100 | 0 |
| ge-11/1/0.100 | 0 |
| xe-10/2/0.100 | 1 |
| xe-10/0/0.100 | 0 |

MAC address count per learn VLAN within routing instance:

| Learn VLAN ID | MAC count |
|---------------|-----------|
| 0 | 2 |

0 MAC address learned in routing instance vs1 bridge domain vlan200

MAC address count per interface within routing instance:

| Logical interface | MAC count |
|-------------------|-----------|
| ge-11/1/0.200 | 0 |
| ge-11/1/1.200 | 0 |
| ge-11/1/4.200 | 0 |
| xe-10/0/0.200 | 0 |
| xe-10/2/0.200 | 0 |

MAC address count per learn VLAN within routing instance:

| Learn VLAN ID | MAC count |
|---------------|-----------|
| 0 | 0 |

show bridge mac-table detail user@host> **show bridge mac-table detail**
 MAC address: 00:00:00:19:1c:db

```

Routing instance: vs1
Bridging domain: vlan100
Learning interface: ge-11/0/3.0   Learning VLAN: 0
Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
Epoch: 4                         Sequence number: 0
Learning mask: 0x800              IPC generation: 0

```

MAC address: 00:00:00:59:3a:2f

```

Routing instance: vs1
Bridging domain: vlan100
Learning interface: xe-10/2/0.100 Learning VLAN: 0
Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
Epoch: 7                         Sequence number: 0
Learning mask: 0x400              IPC generation: 0

```

show bridge statistics

| | |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show bridge statistics <bridge-domain <i>domain-name</i> > <instance <i>instance-name</i> > |
| Release Information | Command introduced in JUNOS Release 8.4. |
| Description | (MX Series routers only) Display bridge statistics. |
| Options | <p>none—Display bridge statistics for all bridge domains in all routing instances.</p> <p>bridge-domain <i>domain-name</i>—(Optional) Display statistics for the specified bridge domain.</p> <p>instance <i>instance-name</i>—(Optional) Display statistics for the specified routing instance.</p> |
| Required Privilege Level | view |

List of Sample Output show bridge statistics on page 653

```

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 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 656 |
| Output Fields | Table 174 on page 656 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 174: 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). |

```

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 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 657 |
| Output Fields | Displays the total number of dynamic and static MAC addresses learned for the entire router. |
| show l2-learning global-mac-count | <pre>user@host> show l2-learning global-mac-count 100 dynamic and static MAC addresses learned globally</pre> |

show l2-learning instance

| | |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show l2-learning instance |
| Release Information | (MX Series routers only) Command introduced in JUNOS 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 658 |
| Output Fields | Table 175 on page 658 describes the output fields for the show l2-learning instance command. Output fields are listed in the approximate order in which they appear. |

Table 175: show l2-learning instance Output Fields

| Field Name | Field Description |
|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Routing Instance | Name of routing instance. |
| bridging domain | Name of bridging domain. |
| 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. |

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 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 659 |
| Output Fields | Table 176 on page 659 describes the output fields for the show l2-learning interface command. Output fields are listed in the approximate order in which they appear. |

Table 176: show l2-learning interface Output Fields

| Field Name | Field Description |
|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Logical interface | Name of the logical interface. |
| Index | Index of the interface. |
| Routing Instance | Number of the routing instance to which the interface belongs. |
| Interface device | Value of the order in which the JUNOS Software 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. |

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 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 661 |
| Output Fields | Table 177 on page 661 lists the output fields for the show mac-rewrite interface command. Output fields are listed in the approximate order in which they appear. |

Table 177: 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 |

```

show mac-rewrite  user@host> show mac-rewrite interface
interface         Interface      Protocols

                    ge-1/0/1      STP VTP CDP

```


Chapter 20

Spanning Tree Operational Mode Commands

Table 178 on page 663 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 178: STP Operational Mode Commands

| Task | Command |
|-----------------------------------------------------|----------------------------------------|
| Clear STP protocol. | clear spanning-tree protocol-migration |
| Clear STP statistics. | clear spanning-tree statistics |
| Display STP bridge domain configuration and status. | show spanning-tree bridge |
| Display STP interface configuration and status. | show spanning-tree interface |
| Display MSTP configuration and status. | show spanning-tree mstp configuration |
| Display STP statistics. | show spanning-tree statistics |



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” on page 637.

For information about how to configure STP, see the *JUNOS Routing Protocols Configuration Guide*.

clear spanning-tree protocol-migration

| | |
|-----------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | clear spanning-tree protocol-migration <interface <i>interface-name</i> > <routing-instance <i>routing-instance-name</i> > |
| Release Information | Command introduced in JUNOS 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 | <p>none—Reset the STP protocol for all interfaces and all routing instances.</p> <p>interface <i>interface-name</i>—(Optional) Reset the STP protocol for the specified interface only.</p> <p>routing-instance <i>routing-instance-name</i>—(Optional) Reset the STP protocol for a particular routing instance.</p> |
| Additional Information | For information about the force-version statement, see the <i>JUNOS Routing Protocols Configuration Guide</i> . |
| Required Privilege Level | clear |
| clear spanning-tree protocol-migration | user@host> clear spanning-tree protocol-migration |

clear spanning-tree statistics

| | |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | clear spanning-tree statistics <interface <i>interface-name</i> > <logical-system <i>logical-system-name</i> > |
| Release Information | Command introduced in JUNOS Release 8.4. |
| Description | Clear Spanning Tree Protocol statistics. |
| Options | <p>none—Reset STP counters for all interfaces for all routing instances on all logical systems.</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> |
| Required Privilege Level | clear |
| Related Topics | show spanning-tree statistics |
| List of Sample Output | clear stp statistics on page 665 |
| 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> > |
| Release Information | Command introduced in JUNOS Release 8.4. |
| 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 | <p>show spanning-tree bridge routing-instance on page 667</p> <p>show spanning-tree bridge msti on page 668</p> <p>show spanning-tree bridge vlan-id (MSTP) on page 668</p> <p>show spanning-tree bridge (VSTP) on page 669</p> <p>show spanning-tree bridge vlan-id (VSTP) on page 669</p> |
| Output Fields | Table 179 on page 666 lists the output fields for the show spanning-tree bridge command. Output fields are listed in the approximate order in which they appear. |

Table 179: show spanning-tree bridge Output Fields

| Field Name | Field Description |
|-----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| Routing instance name | Name of the routing instance under which the bridging domain is configured. |
| Context ID | Internally generated identifier. |
| 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. |
| Root cost | Calculated cost to reach the root bridge from the bridge where the command is entered. |

Table 179: show spanning-tree bridge Output Fields (continued)

| Field Name | Field Description |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| 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 | Configured time 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 router 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 | Internally generated system identifier. |
| MSTI regional root | Bridge ID of the elected MSTP regional root bridge. |

**show spanning-tree
bridge routing-instance**

```
user@host> show spanning-tree bridge routing-instance vs1 detail
```

```
STP bridge parameters
```

```
Routing instance name      : vs1
Context ID                 : 1
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 : 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 : 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

STP bridge parameters for MSTI 2
  MSTI regional root : 32770.00:13: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 : 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**

```

STP bridge parameters
Routing instance name : vs1
Context ID : 1
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 1 101 routing-instance vs1 detail**

```

STP bridge parameters
Routing instance name : vs1
Context ID : 1
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 (VSTP) user@host> **show spanning-tree bridge**

```

STP bridge parameters
Routing instance name : GLOBAL
Context ID : 0
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 (VSTP) 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> > |
| Release Information | Command introduced in JUNOS Release 8.4. |
| 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 | <p>show spanning-tree interface on page 672</p> <p>show spanning-tree interface detail on page 672</p> <p>show spanning-tree interface msti on page 674</p> <p>show spanning-tree interface vlan-id 101 on page 674</p> <p>show spanning-tree interface (VSTP) on page 675</p> <p>show spanning-tree interface vlan-id (VSTP) on page 675</p> |
| Output Fields | Table 180 on page 671 lists the output fields for the show spanning-tree Interface command. Output fields are listed in the approximate order in which they appear. |

Table 180: show spanning-tree Interface Output Fields

| Field Name | Field Description |
|----------------------|---------------------------------------------------------------------------------------------|
| Interface name | Interface configured to participate in the STP, RSTP, VSTP, or MSTP instance. |
| 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. |

Table 180: show spanning-tree Interface Output Fields (continued)

| Field Name | Field Description |
|---------------|------------------------------------------------------------------------------------------------------------------------|
| 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). |

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 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 101 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> > |
| Release Information | Command introduced in JUNOS Release 8.4. |
| 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 on page 676 |
| Output Fields | Table 181 on page 676 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 181: 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 ID | MSTI instance identifier. |
| Member VLANs | VLAN identifiers associated with the MSTI. |

```

show spanning-tree mstp configuration
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 statistics

| | |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show spanning-tree statistics <brief detail > <routing-instance <i>routing-instance-name</i> > |
| Release Information | Command introduced in JUNOS Release 8.4. |
| Description | Display STP statistics. |
| Options | none—Display brief STP statistics. brief detail—(Optional) Display the specified level of output. routing-instance <i>routing-instance-name</i> —(Optional) Display STP statistics for the specified routing instance. |
| Required Privilege Level | view |
| List of Sample Output | show spanning-tree statistics routing-instance on page 678 show spanning-tree statistics interface on page 679 |
| Output Fields | Table 182 on page 678 lists the output fields for the show spanning-tree statistics command. Output fields are listed in the approximate order in which they appear. |

Table 182: 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. |
| Next BPDU transmission | Number of seconds until the next BPDU is scheduled to be sent. |

```

show spanning-tree user@host> show spanning-tree statistics routing-instance vs1 detail
statistics Routing instance level STP statistics
routing-instance 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

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 683

Chapter 21

VPN Operational Mode Commands

Table 183 on page 683 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 183: 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 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> |
| Display VPLS information related to the level 2 address learning process for the specified routing instance. | <code>show vpls flood instance</code> |
| Display VPLS route information related to the level 2 address learning process. | <code>show vpls flood route</code> |
| Display learned VPLS MAC address information. | <code>show vpls mac-table</code> |
| Display VPLS statistics. | <code>show vpls statistics</code> |



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>)> <mac-address> |
| Release Information | Command introduced before JUNOS 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 on all logical systems.</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>mac-address—(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 685 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| clear vpls mac-address | user@host> clear vpls mac-address |

clear vpls mac-table

| | |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | clear vpls mac-table <instance <i>instance-name</i> > <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)> <mac-address> <vlan-id> |
| Release Information | Command introduced before JUNOS 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 on all logical systems.</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 686 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| 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 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) Triggers 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) Triggers 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 687 |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |
| 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 Release 7.4. |
| Description | Display dynamic tunnel database information. |
| Options | <p>none—Display dynamic tunnel database information for all destinations and routing tables on all logical systems.</p> <p><i>destination</i>—(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 689 show dynamic-tunnels database (No Tunnel PIC) on page 689 show dynamic-tunnels database (Tunnel Is Expiring) on page 689 show dynamic-tunnels database (Destination Specified) on page 689 |
| Output Fields | Table 184 on page 688 lists the output fields for the show dynamic-tunnels database command. Output fields are listed in the approximate order in which they appear. |

Table 184: 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 184: 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). |

**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 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 Release 7.4. |
| Description | Display status information about Layer 2 virtual circuits (VCs) 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 on all logical systems.</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 VCs 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 692</p> <p>show l2circuit connections extensive on page 693</p> |
| Output Fields | Table 185 on page 690 lists the output fields for the <code>show l2circuits connections</code> command. Output fields are listed in the approximate order in which they appear. |

Table 185: show l2circuits connections Output Fields

| Field Name | Field Description |
|-----------------------------|-----------------------------------------------------------|
| Layer-2 Circuit Connections | Displays the legends for connection and interface status. |

Table 185: show l2circuits connections Output Fields (continued)

| Field Name | Field Description |
|-----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Neighbor | Remote PE neighbor. |
| Interface | Logical PE-to-CE interface on which the VC is configured. |
| Type | VC type: rmt (remote) or loc (local). |
| Legend for connection status (St) | <p>Status of the VC connection:</p> <ul style="list-style-type: none"> ■ EI—The local VC 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 VC from the neighbor does not match the local VC 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 VC from the neighbor. There is no outgoing label available for use by this VC. ■ NC—The interface is not configured as a CCC or TCC interface. ■ CB—The remote PE router is advertising a different cell bundle from that configured on the local PE router. ■ 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 VC is down because the local VC interface is down. ■ VC-Dn—The VC is down because there is no tunnel LSP from the local PE router to the neighbor. ■ UP—The VC is operational. ■ CF—The router cannot find enough bandwidth to the remote router to satisfy the Layer 2 circuit bandwidth requirement. ■ XX—The VC is down for an unknown reason. This is a programming error. |
| Time last up | Date and time the VC was last operational. |
| # Up trans | Number of times the VC 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. |

Table 185: show l2circuits connections Output Fields (continued)

| Field Name | Field Description |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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 VC: Yes (Null) or No. |
| Incoming label | Label used by the remote side of the VC to send packets destined to the local side. This label is routed to the local VC interface. |
| Outgoing label | Label used by the local side of the VC to send packets to the remote side of the VC. Packets originated on the local VC interface are encapsulated with this label before being placed on the tunnel LSP to the neighbor for this VC. This label is allocated by the neighbor and is used in demultiplexing incoming packets destined for this VC. |
| VC bandwidth | Bandwidth requirement of the Layer 2 circuit. |
| Time | Time at which the event occurred. |
| Event | <p>Event types logged in history.</p> <ul style="list-style-type: none"> ■ loc intf up—Local VC interface went up. ■ loc intf down—Local VC 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 VC 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 VC, and determines whether it should be advertised or withdrawn to or from the remote side. |
| Interface/Lbl/PE | Name associated with an event. It can be the name of the interface, the name of the new label, or if the route to a PE router changed, the name of the PE router that went down. |
| (vc-number) | Unique number (between 1 and $2^{32} - 1$) that identifies the VC configured between two PE routers, but not across the entire network. A VC identifier in conjunction with a neighbor address uniquely identifies the VC. |

show l2circuit connections

user@host> show l2circuit 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         XX -- unknown
NC -- intf encaps not CCC/TCC
CB -- rcvd cell-bundle size bad
+ IB -- incompatible bit rate
+ TM -- TDM misconfiguration

```

Legend for interface status

Up -- operational

Dn -- down

Neighbor: 10.255.245.51

| Interface | Type | St | Time last up | # Up trans |
|------------------|------|----|----------------------|------------|
| t1-0/0/2.0(vc 1) | rmt | Up | Oct 24 11:09:50 2007 | 1 |

Local interface: t1-0/0/2.0, Status: Up, Encapsulation: SATOP-T1
 Remote PE: 10.255.245.51, Negotiated control-word: Yes
 Incoming label: 100048, Outgoing label: 100128

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         XX -- unknown
NC -- intf encaps not CCC/TCC
CB -- rcvd cell-bundle size bad

```

Neighbor: 10.1.1.195

| Interface | Type | St | Time last up | # Up trans |
|-----------------------------|------|----|---------------------|------------|
| so-0/1/0.1 (vc 1 primary-W) | rmt | Up | Oct 8 15:44:19 2001 | 1 |

Local interface: so-0/1/0.1, Status: Up
 Remote PE: 10.1.1.195, , Negotiated control-word: Yes (Null)
 Incoming label: 100002, Outgoing label: 138264604

| Time | Event | Interface/Lbl/PE |
|---------------------|---------------------|------------------|
| Oct 8 15:44:19 2001 | status update timer | |
| Oct 8 15:44:18 2001 | PE route changed | |
| Oct 8 15:44:18 2001 | Out lbl Update | 138264604 |
| Oct 8 15:44:18 2001 | In lbl Update | 100002 |
| Oct 8 15:44:18 2001 | loc intf up | so-0/1/0.1 |

| Interface | Type | St | Time last up | # Up trans |
|-----------------------------|------|----|---------------------|------------|
| so-0/1/0.2 (vc 1 primary-W) | rmt | Up | Oct 8 15:44:19 2001 | 1 |

Local interface: so-0/1/0.2, Status: Up
 Remote PE: 10.1.1.195, , Negotiated control-word: Yes (Null)
 Incoming label: 100003, Outgoing label: 138264604

| Time | Event | Interface/Lbl/PE |
|---------------------|---------------------|------------------|
| Oct 8 15:44:19 2001 | status update timer | |
| Oct 8 15:44:17 2001 | PE route changed | |
| Oct 8 15:44:17 2001 | Out lbl Update | 138264604 |
| Oct 8 15:44:17 2001 | In lbl Update | 100003 |
| Oct 8 15:44:17 2001 | loc intf up | so-0/1/0.2 |

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 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 on all logical systems.</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 696</p> <p>show l2vpn connections extensive on page 696</p> |
| Output Fields | Table 186 on page 695 lists the output fields for the show l2vpn connections command. Output fields are listed in the approximate order in which they appear. |

Table 186: 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. |
| Status | Status of local or remote circuit: <ul style="list-style-type: none"> ■ Up—Operational ■ Dn—Down ■ NP—Not present ■ DS—Disabled ■ WE—Wrong encapsulation ■ UN—Uninitialized |
| Remote PE | Address of the remote provider edge router. |
| Incoming label | Name of the incoming label. |
| Outgoing label | Name of the outgoing label. |
| Time | Date and time of Layer 2 VPN connection event. |
| Event | Type of event. |

Table 186: show l2vpn connections Output Fields (continued)

| Field Name | Field Description |
|------------------|---------------------------------|
| Interface/Lbl/PE | Interface, label, or PE router. |

```

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/Lb1/PE |
|---------------------|--|----------------|------------------|
| Aug 3 00:08:14 2001 | | PE route up | |
| Aug 3 00:08:14 2001 | | Out lb1 Update | 32768 |
| Aug 3 00:08:14 2001 | | In lb1 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 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 698 show mvpn c-multicast summary on page 699 show mvpn c-multicast extensive on page 699 |
| Output Fields | Table 187 on page 698 lists the output fields for the show mvpn c-multicast command. Output fields are listed in the approximate order in which they appear. |

Table 187: 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) | Cutomer 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 |
| MVPN instance | Name of the multicast VPN routing instance | extensive none |
| C-multicast IPv4 route count | Number of c-multicast IPv4 routes associated with the multicast VPN routing instance. | summary |

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 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 701 show mvpn instance summary on page 701 show mvpn instance extensive on page 701 |
| Output Fields | Table 188 on page 700 lists the output fields for the show mvpn instance command. Output fields are listed in the approximate order in which they appear. |

Table 188: 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) | Cutomer 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 |
| Neighbor count | Number of neighbors associated with the multicast VPN routing instance. | summary |
| C-multicast IPv4 route count | Number of c-multicast IPv4 routes associated with the multicast VPN routing instance. | summary |

```

show mvpn instance user@host> show mvpn instance
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
  Provider tunnel: I-P-tnl:PIM-SM:10.255.14.144, 239.1.1.1
  Neighbor
    10.255.14.160          I-P-tnl
    10.255.70.17          PIM-SM:10.255.14.160, 239.1.1.1
    10.255.70.17          PIM-SM:10.255.70.17, 239.1.1.1
  C-mcast IPv4 (S:G)      Ptnl          St
    192.168.195.78/32:225.5.5.5/32 PIM-SM:10.255.14.144, 239.1.1.1      RM
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-B
  Provider tunnel: I-P-tnl:PIM-SM:10.255.14.144, 239.2.0.0
  Neighbor
    10.255.14.160          I-P-tnl
    10.255.70.17          PIM-SM:10.255.14.160, 239.2.0.0
    10.255.70.17          PIM-SM:10.255.70.17, 239.2.0.0
  C-mcast IPv4 (S:G)      Ptnl          St
    192.168.195.94/32:226.6.6.6/32 PIM-SM:10.255.14.144, 239.2.0.0      RM

show mvpn instance summary user@host> show mvpn instance summary
MVPN Summary:
Instance: VPN-A
  Neighbor count: 2
  C-multicast IPv4 route count: 1
Instance: VPN-B
  Neighbor count: 4
  C-multicast IPv4 route count: 2

show mvpn instance extensive user@host> show mvpn instance extensive
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
  Provider tunnel: I-P-tnl:PIM-SM:10.255.14.144, 239.1.1.1
  Neighbor
    10.255.14.160          I-P-tnl
    10.255.70.17          PIM-SM:10.255.14.160, 239.1.1.1
    10.255.70.17          PIM-SM:10.255.70.17, 239.1.1.1
  C-mcast IPv4 (S:G)      Ptnl          St
    192.168.195.78/32:225.5.5.5/32 PIM-SM:10.255.14.144, 239.1.1.1      RM
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)

```

```

DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-B
  Provider tunnel: I-P-tnl:PIM-SM:10.255.14.144, 239.2.0.0
  Neighbor                               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 neighbor

| | |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show mvpn neighbor <extensive summary> <instance <i>instance-name</i> neighbor-address <i>address</i> > |
| Release Information | Command introduced in JUNOS Release 8.4. |
| Description | Display the multicast VPN neighbor information. |
| Options | extensive summary—(Optional) Display the specified level of output. instance <i>instance-name</i> neighbor-address <i>address</i> —(Optional) Display multicast VPN neighbor information for the specified instance or the specified neighbor. |
| Required Privilege Level | view |
| List of Sample Output | show mvpn neighbor on page 703 show mvpn neighbor extensive on page 704 show mvpn neighbor instance-name on page 704 show mvpn neighbor neighbor-address on page 704 show mvpn neighbor neighbor-address summary on page 705 show mvpn neighbor neighbor-address extensive on page 705 show mvpn neighbor neighbor-address instance-name on page 705 |
| Output Fields | Table 189 on page 703 lists the output fields for the show mvpn neighbor command. Output fields are listed in the approximate order in which they appear. |

Table 189: 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 |

```

show mvpn neighbor  user@host> show mvpn neighbor
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
Neighbor                             I-P-tnl

```

```

10.255.14.160          PIM-SM:10.255.14.160, 239.1.1.1
10.255.70.17          PIM-SM:10.255.70.17, 239.1.1.1
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-B
Neighbor
10.255.14.160          I-P-tnl
                       PIM-SM:10.255.14.160, 239.2.0.0
10.255.70.17          PIM-SM:10.255.70.17, 239.2.0.0

```

show mvpn neighbor extensive user@host> **show mvpn neighbor extensive**
MVPN instance:

```

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
Neighbor
10.255.14.160          I-P-tnl
                       PIM-SM:10.255.14.160, 239.1.1.1
10.255.70.17          PIM-SM:10.255.70.17, 239.1.1.1
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
10.255.14.160          I-P-tnl
                       PIM-SM:10.255.14.160, 239.2.0.0
10.255.70.17          PIM-SM:10.255.70.17, 239.2.0.0

```

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
10.255.14.160          I-P-tnl
                       PIM-SM:10.255.14.160, 239.1.1.1
10.255.70.17          PIM-SM:10.255.70.17, 239.1.1.1

```

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

```

```

show mvpn neighbor      user@host> show mvpn neighbor neighbor-address 10.255.70.17 summary
neighbor-address      MVPN Summary:
summary                Instance: VPN-A
                          Instance: VPN-B

```

```

show mvpn neighbor      user@host> show mvpn neighbor neighbor-address 10.255.70.17 extensive
neighbor-address      MVPN instance:
extensive

```

```

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

```

```

show mvpn neighbor      user@host> show mvpn neighbor neighbor-address 10.255.70.17 instance-name VPN-A
neighbor-address      MVPN instance:
instance-name

```

```

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 show vpls connections
 <brief | extensive>
 <down | up | up-down>
 <history>
 <instance *instance-name* local-site *local-site-name* remote-site *remote-site-name*>
 <logical-system (all | *logical-system-name*)>
 <status>
 <summary>

Release Information Command introduced before JUNOS Release 7.4.

Description (T Series and M Series routers, except for the M160 router) Display virtual private LAN service (VPLS) connection information.

Options none—Display information about all VPLS connections for all routing instances on all logical systems.

brief | extensive—(Optional) Display the specified level of output.

down | up | up-down—(Optional) Display nonoperational, operational, or both types of connections.

history—(Optional) Display information about connection history.

instance *instance-name*—(Optional) Display the VPLS connections for the specified routing instance only.

local-site *local-site-name*—(Optional) Display the VPLS connections for the specified local site name or ID only.

remote-site *remote-site-name*—(Optional) Display the VPLS connections for the specified remote site name or ID only.

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

status—(Optional) Display information about the connection and interface status.

summary—(Optional) Display summary of all VPLS connections information.

Required Privilege Level view

List of Sample Output show vpls connections on page 708
 show vpls connections extensive on page 708

Output Fields Table 190 on page 707 lists the output fields for the **show vpls connections** command. Output fields are listed in the approximate order in which they appear.

Table 190: show vpls connections Output Fields

| Field Name | Field Description |
|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Instance | Name of the VPLS instance. |
| Local site | Name of the local site. |
| Local interface | Name of the local VPLS virtual loopback tunnel interface or label switched interface. Virtual loopback tunnel interfaces are displayed using the <i>vt-fpc/pic/port.nnnnn</i> format. Label switched interfaces are displayed using the <i>lsi.nnnnn</i> format. In both cases, <i>nnnnn</i> is a dynamically generated virtual port used to transport and receive packets from other provider edge (PE) routers in the VPLS domain. |
| 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. |
| Range | 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: <i>loc</i> (local) or <i>rmt</i> (remote). |
| St | Status of the connection. (For a list of possible values, see the Legend for connection status (St) field.) |
| Time last up | Time connection was last in the Up condition. |
| # Up trans | Number of transitions from Down to Up condition. |
| Status | Status of the (local or remote circuit) local interface: <ul style="list-style-type: none"> ■ Up—Operational ■ Dn—Down ■ NP—Not present ■ DS—Disabled ■ WE—Wrong encapsulation ■ UN—Uninitialized |
| Encapsulation | Type of encapsulation: VPLS . |
| Remote PE | Address of the remote provider edge router. |
| Negotiated control-word | Whether a control word has been negotiated: Yes or No . |
| Incoming label | Name of the incoming label. |
| Outgoing label | Name of the outgoing label. |
| Time | Date and time of VPLS connection event. |

Table 190: show vpls connections Output Fields (continued)

| Field Name | Field Description |
|------------------|---------------------------------|
| Event | Type of event. |
| Interface/Lbl/PE | Interface, label, or PE router. |

show vpls connectionsuser@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 -- unknown connection status | |

Legend for interface status

Up -- operational
Dn -- down

Instance: green

Local site: greenR1 (1)

| connection-site | Type | St | Time last up | # Up trans |
|-----------------|------|----|----------------------|------------|
| 2 | rmt | RD | | |
| 4 | rmt | Up | Oct 31 12:46:23 2005 | 1 |

Local interface: vt-3/3/0.3, Status: Up, Encapsulation: VPLS
Remote PE: 10.255.14.222, Negotiated control-word: No
Incoming label: 800011, Outgoing label: 800000

Local site: SW0 (4)

| connection-site | Type | St | Time last up | # Up trans |
|-----------------|------|----|--------------|------------|
| 2 | rmt | LN | | |
| 4 | rmt | SC | | |
| 5 | rmt | LN | | |

Local site: SW1 (5)

| connection-site | Type | St | Time last up | # Up trans |
|-----------------|------|----|--------------|------------|
| 2 | rmt | LM | | |
| 4 | rmt | LM | | |
| 5 | rmt | SC | | |

Instance: red

Local site: redR1 (1)

| connection-site | Type | St | Time last up | # Up trans |
|-----------------|------|----|--------------|------------|
| 2 | rmt | LD | | |

show vpls connections extensiveuser@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 -- unknown connection status

```

Legend for interface status

```

Up -- operational
Dn -- down

```

Instance: red

Local site: redR1 (1)

```

Number of local interfaces: 1
Number of local interfaces up: 1

```

```
lt-2/1/0.4
```

```
vt-2/1/0.32780 2 Intf - vpls red local site 1 remote site 2
```

```
Interface flags: VC-Down
```

```
vt-2/1/0.32782 3 Intf - vpls red local site 1 remote site 3
```

```
vt-2/1/0.32773 4 Intf - vpls red local site 1 remote site 4
```

```
vt-2/1/0.32779 5 Intf - vpls red local site 1 remote site 5
```

```
Interface flags: VC-Down
```

```
800024 1 8 100
```

status-vector: 87

| connection-site | Type | St | Time last up | # Up trans |
|-----------------|------|----|----------------------|------------|
| 2 | rmt | RD | | |
| 3 | rmt | Up | Oct 28 07:13:18 2005 | 1 |

```
Local interface: vt-2/1/0.32782, Status: Up, Encapsulation: VPLS
```

```
Remote PE: 10.255.14.218, Negotiated control-word: No
```

```
Incoming label: 800026, Outgoing label: 800024
```

| Time | Event | Interface/Lbl/PE |
|----------------------|---------------------|------------------|
| Oct 28 07:13:18 2005 | status update timer | |
| Oct 28 07:13:18 2005 | PE route changed | |
| Oct 28 07:13:18 2005 | Out lbl Update | 800024 |
| Oct 28 07:13:18 2005 | In lbl Update | 800026 |
| Oct 28 07:13:18 2005 | loc intf up | vt-2/1/0.32782 |

| | | | | |
|---|-----|----|----------------------|---|
| 4 | rmt | Up | Oct 28 07:13:12 2005 | 1 |
|---|-----|----|----------------------|---|

```
Local interface: vt-2/1/0.32773, Status: Up, Encapsulation: VPLS
```

```
Remote PE: 10.255.14.216, Negotiated control-word: No
```

```
Incoming label: 800027, Outgoing label: 800008
```

| Time | Event | Interface/Lbl/PE |
|----------------------|---------------------|------------------|
| Oct 28 07:13:12 2005 | status update timer | |
| Oct 28 07:13:12 2005 | PE route changed | |
| Oct 28 07:13:12 2005 | Out lbl Update | 800008 |
| Oct 28 07:13:12 2005 | In lbl Update | 800027 |
| Oct 28 07:13:12 2005 | loc intf up | vt-2/1/0.32773 |

| | | | | |
|---|-----|----|--|--|
| 5 | rmt | RM | | |
|---|-----|----|--|--|

Local site: SW0 (4)

```

Number of local interfaces: 1
Number of local interfaces up: 1

```

```
fe-0/1/2.1
```

```
Interface flags: VC-Down
```

```
vt-2/1/0.32771 2 Intf - vpls red local site 4 remote site 2
```

```
Interface flags: VC-Down
```

```
vt-2/1/0.32783 3 Intf - vpls red local site 4 remote site 3
```

```

      Interface flags: VC-Down
vt-2/1/0.32772  5      Intf - vpls red local site 4 remote site 5
      Interface flags: VC-Down
      800016          1          8          100
status-vector: 97
  connection-site      Type  St      Time last up      # Up trans
    2                  rmt   LN
    3                  rmt   LN
    4                  rmt   SC
    5                  rmt   LN
Local site: SW1 (5)
Number of local interfaces: 1
Number of local interfaces up: 1
fe-0/1/3.1
  Interface flags: VC-Down
vt-2/1/0.32778  2      Intf - vpls red local site 5 remote site 2
  Interface flags: VC-Down

```

show vpls flood event-queue

| | |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show vpls flood event-queue |
| Release Information | Command introduced in JUNOS 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 711 |
| Output Fields | Table 191 on page 711 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 191: 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. |

```

show vpls flood      user@host> show vpls flood event-queue
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 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 on all logical systems.</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 | <p>show vpls flood instance on page 713</p> <p>show vpls flood instance logical-system-name on page 713</p> <p>show vpls flood instance detail on page 713</p> |
| Output Fields | Table 192 on page 712 lists the output fields for the show vpls flood instance command. Output fields are listed in the approximate order in which they appear. |

Table 192: show vpls flood instance Output Fields

| Field Name | Field Description |
|----------------|---------------------------------------------------------------|
| Logical system | Name of the logical system. |
| Name | Name of the VPLS routing instance. |
| CEs | Number of CE routers connected to the VPLS instance. |
| VEs | Number of VE routers connected to the VPLS instance. |
| Flood routes | List of all flood routes associated with the VPLS instance. |
| Prefix | Prefix for the route. |
| Type | Type of route. |
| Owner | VPLS routing instance or interface associated with the route. |
| Nhype | Next-hop type. For example, flood for a flood route. |

Table 192: show vpls flood instance Output Fields *(continued)*

| Field Name | Field Description |
|------------|--------------------------------------|
| Nhindex | Next-hop index number for the route. |

```

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    user@host:__juniper_ls1__> show vpls flood instance juniper_ls1
logical-system-name

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    user@host:__juniper_ls1__> show vpls flood instance detail
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 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 <code>__juniper_ls1__</code> (the name must be entered in the command with the underscore characters).</p> |
| Required Privilege Level | view |
| List of Sample Output | <p>show vpls flood route all-ce-flood on page 715</p> <p>show vpls flood route ce-flood on page 715</p> |
| Output Fields | Table 193 on page 714 lists the output for the <code>show vpls flood route</code> command. Output fields are listed in the approximate order in which they appear. |

Table 193: 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 <code>CE_FLOOD</code> or <code>ALL_CE_FLOOD</code>). |
| Flood route owner | VPLS routing instance or interface associated with the flood route. |
| Nexthop type | Next-hop type. For example, <code>flood</code> 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. |
| Type | Type of VPLS router (CE or VE). |
| Nh type | Next-hop type. |

Table 193: show vpls flood route Output Fields (continued)

| Field Name | Field Description |
|------------|-----------------------------------|
| Index | Index number for the flood route. |

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 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 | <pre>show vpls mac-table on page 717 show vpls mac-table count on page 717 show vpls mac-table detail on page 718 show vpls mac-table extensive on page 719</pre> |
| Output Fields | Table 194 on page 716 describes the output fields for the show bridge mac-table command. Output fields are listed in the approximate order in which they appear. |

Table 194: show vpls mac-table Output fields

| Field Name | Field Description |
|------------------|-------------------------------|
| Routing instance | Name of the routing instance. |

Table 194: 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. |

```

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 |
|-------------------|-----------|
| lc-0/0/0.32769 | 0 |
| lc-0/1/0.32769 | 0 |
| lc-0/2/0.32769 | 0 |
| lc-2/0/0.32769 | 0 |
| lc-0/3/0.32769 | 0 |
| lc-2/1/0.32769 | 0 |
| lc-9/0/0.32769 | 0 |
| lc-11/0/0.32769 | 0 |
| lc-2/2/0.32769 | 0 |
| lc-9/1/0.32769 | 0 |
| lc-11/1/0.32769 | 0 |
| lc-2/3/0.32769 | 0 |
| lc-9/2/0.32769 | 0 |
| lc-11/2/0.32769 | 0 |
| lc-11/3/0.32769 | 0 |
| lc-9/3/0.32769 | 0 |

MAC address count per learn VLAN within routing instance:

| Learn VLAN ID | MAC count |
|---------------|-----------|
| 0 | 0 |

1 MAC address learned in routing instance vpls_ldp1

MAC address count per interface within routing instance:

| Logical interface | MAC count |
|-------------------|-----------|
| lsi.1051137 | 0 |
| ge-0/2/5.400 | 1 |

MAC address count per learn VLAN within routing instance:

| Learn VLAN ID | MAC count |
|---------------|-----------|
| 0 | 1 |

1 MAC address learned in routing instance vpls_red

MAC address count per interface within routing instance:

| Logical interface | MAC count |
|-------------------|-----------|
| ge-0/2/5.300 | 1 |

MAC address count per learn VLAN within routing instance:

| Learn VLAN ID | MAC count |
|---------------|-----------|
| 0 | 1 |

show vpls mac-table detail

user@host> show vpls mac-table detail

MAC address: 00:90:69:9c:1c:5d

Routing instance: vpls_ldp1

Learning interface: ge-0/2/5.400

Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel

Epoch: 0

Sequence number: 1

Learning mask: 0x1

IPC generation: 0

MAC address: 00:90:69:9c:1c:5d

Routing instance: vpls_red

Learning interface: ge-0/2/5.300

Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel

Epoch: 0

Sequence number: 1

Learning mask: 0x1

IPC generation: 0

```

show vpls mac-table extensive user@host> show vpls mac-table extensive
MAC address: 00:00:aa:12:12:12
  Routing instance: vpls_ldp1
  Learning interface: lsi.1051137
  Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
  Epoch: 0                               Sequence number: 1
  Learning mask: 0x1                       IPC generation: 0

MAC address: 00:05:85:74:9f:f0
  Routing instance: vpls_ldp1
  Learning interface: lsi.1051137
  Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
  Epoch: 0                               Sequence number: 1
  Learning mask: 0x1                       IPC generation: 0

MAC address: 00:90:69:9c:1c:5d
  Routing instance: vpls_ldp1
  Learning interface: ge-0/2/5.400
  Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
  Epoch: 0                               Sequence number: 1
  Learning mask: 0x1                       IPC generation: 0

MAC address: 00:00:aa:12:12:12
  Routing instance: vpls_red
  Learning interface: lsi.1051138
  Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
  Epoch: 0                               Sequence number: 0
  Learning mask: 0x1                       IPC generation: 0

MAC address: 00:05:85:74:9f:f0
  Routing instance: vpls_red
  Learning interface: lsi.1051138
  Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
  Epoch: 0                               Sequence number: 0
  Learning mask: 0x1                       IPC generation: 0

```

show vpls statistics

| | |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax | show vpls statistics <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> |
| Release Information | Command introduced before JUNOS 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 on all logical systems.</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 | <p>show vpls statistics on page 721</p> <p>show vpls statistics instance on page 721</p> |
| Output Fields | Table 195 on page 720 lists the output fields for the show vpls statistics command. Output fields are listed in the approximate order in which they appear. |

Table 195: show vpls statistics Output Fields

| Field Name | Field Description |
|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Instance | Name of the VPLS instance. |
| Local interface | Name of the local VPLS virtual loopback tunnel interface, <i>vt-fpc/pic/port.nnnnn</i> , where <i>nnnnn</i> is a dynamically generated virtual port used to transport and receive packets from other provider edge (PE) routers in the VPLS domain. |
| Index | Number associated with the next hop. |
| Remote provider edge router | Address of the remote provider edge router. |
| Multicast packets | Number of multicast packets received. |
| Multicast bytes | Number of multicast bytes received. |
| Flood packets | Number of VPLS flood packets received. |
| Flood bytes | Number of VPLS flood bytes received. |
| 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. |

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 user@host> **show vpls statistics instance red**
instance

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 725
- Index of Statements and Commands on page 735

Index

Symbols

| | |
|----------------------------------------------|-----|
| #, comments in configuration statements..... | xx |
| (), in syntax descriptions..... | xx |
| < >, in syntax descriptions..... | xx |
| [], in configuration statements..... | xx |
| { }, in configuration statements..... | xxi |
| (pipe), in syntax descriptions..... | xx |

A

administrative groups *See* groups
 advertisements *See* router advertisements
 ANCP

| | |
|----------------------|----|
| CoS state | |
| displaying ANCP..... | 5 |
| local access loop | |
| displaying..... | 10 |
| neighbor | |
| clearing..... | 4 |
| neighbors | |
| displaying..... | 8 |
| subscribers | |
| displaying..... | 10 |

AS paths

| | |
|-----------------------------------------------|-----|
| distribution of, displaying..... | 324 |
| domain information, displaying..... | 328 |
| matching regular expressions, displaying..... | 344 |
| summary of, displaying..... | 330 |

ASN

| | |
|---------------------------------------|-----|
| BGP community routes, displaying..... | 352 |
|---------------------------------------|-----|

automatic bandwidth allocation

| | |
|-----------|-----|
| LSPs..... | 569 |
|-----------|-----|

autonomous system number *See* ASN

autonomous system paths *See* AS paths

B

bandwidth, allocating for LSPs.....569

best routes, displaying.....346

BFD

| | |
|-----------------|----|
| sessions | |
| clearing..... | 17 |
| displaying..... | 18 |

BGP

| | |
|----------------------------------------|-----|
| community ASN, displaying routes..... | 352 |
| community name, displaying routes..... | 353 |
| damping parameters | |
| clearing..... | 26 |
| displaying..... | 54 |
| damping routes, displaying..... | 355 |
| groups | |
| general information, displaying..... | 31 |
| traffic statistics, displaying..... | 37 |
| neighbors | |
| clearing connections..... | 27 |
| displaying..... | 39 |
| summary information, displaying..... | 50 |
| table | |
| clearing..... | 29 |

BGP Monitoring Protocol

| | |
|-----------------|----|
| displaying | |
| statistics..... | 30 |

Bidirectional Forwarding Detection *See* BFD

binding state of DHCP client

| | |
|-----------------|-----|
| clearing..... | 516 |
| displaying..... | 521 |

bootstrap routers, displaying.....174

Border Gateway Protocol *See* BGP

braces, in configuration statements.....xxi

brackets

| | |
|------------------------------------------|----|
| angle, in syntax descriptions..... | xx |
| square, in configuration statements..... | xx |

bridge domain

| | |
|--------------------------|-----|
| Layer 2, displaying..... | 641 |
|--------------------------|-----|

C

CAC

| | |
|--------------------------|-----|
| displaying for LSPs..... | 586 |
|--------------------------|-----|

call admission control *See* CAC

CCC

| | |
|------------------------------|-----|
| connections, displaying..... | 570 |
|------------------------------|-----|

circuit cross-connect *See* CCC

clear (ospf | ospf3) database command.....271

clear (ospf | ospf3) io-statistics command.....274

clear (ospf | ospf3) neighbor command.....275

clear (ospf | ospf3) statistics command.....277

clear ancip neighbor command.....4

clear bfd adaptation command.....16

| | |
|--------------------------------------------------------|-----|
| clear bfd session command..... | 17 |
| clear bgp command..... | 29 |
| clear bgp damping command..... | 26 |
| clear bgp neighbor command..... | 27 |
| clear bridge mac-table command..... | 638 |
| clear dhcp relay binding command..... | 516 |
| clear dhcp relay statistics command..... | 518 |
| clear error bpdu command..... | 639 |
| clear error mac-rewrite command..... | 640 |
| clear esis adjacency command..... | 58 |
| clear esis statistics command..... | 59 |
| clear firewall command..... | 506 |
| clear helper statistics command..... | 520 |
| clear igmp membership command..... | 69 |
| clear igmp snooping membership command..... | 72 |
| clear igmp snooping statistics command..... | 73 |
| clear igmp statistics command..... | 74 |
| clear ipv6 neighbors command..... | 204 |
| clear ipv6 router-advertisement command..... | 205 |
| clear isis adjacency command..... | 213 |
| clear isis database command..... | 214 |
| clear isis overload command..... | 216 |
| clear isis statistics command..... | 217 |
| clear ldp neighbor command..... | 540 |
| clear ldp session command..... | 541 |
| clear ldp statistics command..... | 542 |
| clear lldp neighbor command..... | 256 |
| clear lldp statistics command..... | 257 |
| clear mld membership command..... | 76 |
| clear mld statistics command..... | 77 |
| clear msdp cache command..... | 78 |
| clear msdp statistics command..... | 79 |
| clear multicast bandwidth-admission command..... | 82 |
| clear multicast scope command..... | 80 |
| clear multicast sessions command..... | 84 |
| clear multicast snooping statistics command..... | 81 |
| clear multicast statistics command..... | 85 |
| clear ospf overload command..... | 276 |
| clear pgm negative-acknowledgments command..... | 86 |
| clear pgm source-path-messages command..... | 87 |
| clear pgm statistics command..... | 88 |
| clear pim join command..... | 89 |
| clear pim register command..... | 90 |
| clear pim statistics command..... | 91 |
| clear rip general-statistics command..... | 488 |
| clear rip statistics command..... | 489 |
| clear ripng general-statistics command..... | 496 |
| clear ripng statistics command..... | 497 |
| clear rsvp session command..... | 610 |
| clear rsvp statistics command..... | 612 |
| clear spanning-tree protocol-migration command..... | 664 |
| clear spanning-tree statistics command..... | 665 |
| clear vpls mac-address command..... | 685 |
| clear vpls mac-table command..... | 686 |
| comments, in configuration statements..... | xx |

| | |
|-------------------------------------------------|-----|
| community ASN, displaying routes..... | 352 |
| community name, displaying routes..... | 353 |
| conditions routing policy..... | 534 |
| Constrained Shortest Path First <i>See</i> CSPF | |
| conventions text and syntax..... | xx |
| CSPF statistics, displaying..... | 588 |
| curly braces, in configuration statements..... | xxi |
| customer support..... | xxi |
| contacting JTAC..... | xxi |

D

| | |
|-------------------------------------------------------------|-----|
| damping parameters, BGP clearing..... | 26 |
| displaying..... | 54 |
| damping routes, BGP displaying..... | 355 |
| DHCP relay binding clearing..... | 516 |
| relay binding state displaying..... | 521 |
| relay statistics clearing..... | 518 |
| displaying..... | 525 |
| DHCP client binding state clearing..... | 516 |
| displaying..... | 521 |
| statistics clearing..... | 518 |
| DiffServ classes, displaying for MPLS..... | 590 |
| Distance Vector Multicast Routing Protocol <i>See</i> DVMRP | |
| documentation set comments on..... | xxi |
| DVMRP groups, displaying..... | 166 |
| interfaces, displaying..... | 93 |
| neighbors, displaying..... | 95 |
| prefixes, displaying..... | 97 |
| prunes, displaying active..... | 99 |
| dynamic overload bit, resetting for IS-IS..... | 216 |
| dynamic tunnel database, displaying..... | 688 |

E

| | |
|----------------------------------------------------|--|
| End System-to-Intermediate System <i>See</i> ES-IS | |
|----------------------------------------------------|--|

| | |
|-------------------------------------------|-----|
| ES-IS | |
| adjacencies | |
| clearing..... | 58 |
| displaying..... | 60 |
| interfaces, displaying..... | 62 |
| statistics | |
| clearing..... | 59 |
| displaying..... | 64 |
| export route information, displaying..... | 376 |

F

| | |
|-------------------------------------------------|--------------------|
| FEC filters | |
| displaying for LDP..... | 546 |
| firewall | |
| statistics | |
| displaying..... | 507 |
| firewall filters | |
| log information, displaying..... | 510 |
| policed packets, displaying..... | 513 |
| statistics | |
| clearing..... | 506 |
| displaying..... | 512 |
| flooding | |
| Layer 2 bridging, displaying..... | 643 |
| font conventions..... | xx |
| forwarding equivalence class <i>See</i> FEC | |
| forwarding options | |
| DHCP relay agent..... | 516, 518, 521, 525 |
| UDP statistics | |
| clearing..... | 520 |
| displaying..... | 528 |
| forwarding table | |
| multicast information, displaying..... | 146 |
| multicast snooping information, displaying..... | 158 |
| route entries, displaying..... | 394 |

G

| | |
|-----------------------------------------|-----|
| generalized MPLS <i>See</i> GMPLS | |
| GMPLS | |
| link-management information, displaying | |
| all..... | 573 |
| peers..... | 576 |
| routing process..... | 578 |
| statistics..... | 581 |
| traffic-engineered links..... | 583 |
| groups | |
| BGP | |
| general information, displaying..... | 31 |
| traffic statistics, displaying..... | 37 |
| DVMRP, displaying..... | 166 |
| IGMP membership, displaying..... | 101 |
| MLD | |
| clearing..... | 76 |
| displaying..... | 118 |

| | |
|--------------------------------------|-----|
| MPLS, displaying administrative..... | 585 |
| PIM | |
| general information, displaying..... | 179 |
| usage information, displaying..... | 166 |

H

| | |
|--------------------------------|-----|
| hidden routes, displaying..... | 405 |
| hostnames | |
| IS-IS, displaying..... | 237 |

I

| | |
|-------------------------------------------------------------|-----|
| icons defined, notice..... | xix |
| IGMP | |
| group membership, displaying..... | 101 |
| interfaces, displaying..... | 104 |
| PIM-to-IGMP message translation information, | |
| displaying..... | 144 |
| snooping (interface)..... | 107 |
| snooping (membership)..... | 109 |
| snooping (statistics)..... | 112 |
| statistics, displaying..... | 115 |
| interface learning information | |
| Layer 2, displaying..... | 659 |
| Intermediate System-to-Intermediate System <i>See</i> IS-IS | |
| Internet Group Management Protocol <i>See</i> IGMP | |
| invalid routes, displaying..... | 423 |
| IP IGMP snooping | |
| membership | |
| clearing..... | 72 |
| statistics | |
| clearing..... | 73 |
| IP multicast..... | 65 |
| announced sessions, displaying..... | 156 |
| bandwidth admission | |
| clearing..... | 82 |
| flow map information, displaying..... | 136 |
| forwarding table, displaying..... | 146 |
| interface information, displaying..... | 138 |
| network information, displaying..... | 140 |
| next-hop table, displaying..... | 142 |
| PIM-to-IGMP message translation information, | |
| displaying..... | 144 |
| PIM-to-MLD message translation information, | |
| displaying..... | 145 |
| RPF calculations, displaying..... | 150 |
| SAP announcements, displaying..... | 201 |
| scope, clearing..... | 80 |
| scoped information, displaying..... | 154 |
| sessions, clearing..... | 84 |
| statistics | |
| clearing..... | 85 |
| displaying..... | 163 |
| <i>See also</i> DVMRP, MDT, MLD, MSDP, PGM, PIM | |

| | |
|-------------------------------------------|-----|
| IP multicast snooping | |
| forwarding table, displaying..... | 158 |
| statistics | |
| clearing..... | 81 |
| displaying..... | 160 |
| IPv6 | |
| neighbor cache information | |
| clearing..... | 204 |
| displaying..... | 206 |
| router advertisements | |
| clearing..... | 205 |
| displaying..... | 207 |
| IS-IS | |
| adjacency database entries, clearing..... | 213 |
| authentication, displaying..... | 223 |
| backup coverage | |
| displaying..... | 225 |
| backup MPLS LSPs..... | 227 |
| backup paths | |
| SPF calculations..... | 228 |
| dynamic overload bit, resetting..... | 216 |
| hostname database, displaying..... | 237 |
| interfaces, displaying..... | 238 |
| link-state database entries | |
| clearing..... | 214 |
| displaying..... | 231 |
| neighbors, displaying..... | 219 |
| routes, displaying..... | 244 |
| SPF calculations, displaying..... | 247 |
| traffic statistics | |
| clearing..... | 217 |
| displaying..... | 252 |

J

| | |
|--------------------------------|----|
| join states, clearing PIM..... | 89 |
|--------------------------------|----|

L

Label Distribution Protocol *See* LDP

| | |
|-------------------------------------------------|----------|
| Layer 2 BPDU loop prevention | |
| clearing errors..... | 639 |
| Layer 2 bridging..... | 637 |
| bridge domain, displaying..... | 641 |
| flooding, displaying..... | 643 |
| interface learning information, displaying..... | 659 |
| learning information, displaying..... | 656, 657 |
| learning properties, displaying..... | 658 |
| MAC address table, clearing..... | 638 |
| MAC address table, displaying..... | 650 |
| statistics, displaying..... | 653 |
| Layer 2 circuits | |
| connections, displaying..... | 690 |
| redundant pseudowires..... | 687 |

| | |
|-------------------------------------------|----------|
| Layer 2 protocol tunneling | |
| clearing errors..... | 640 |
| interface information, displaying..... | 661 |
| Layer 2 switching..... | 637 |
| Layer 2 VPNs | |
| connections, displaying..... | 694 |
| VPLS connections, displaying..... | 706 |
| VPLS statistics, displaying..... | 720 |
| Layer 3 VPNs | |
| dynamic tunnel database, displaying..... | 688 |
| PIM MDTs, displaying..... | 184 |
| PIM, displaying..... | 187 |
| LDP | |
| database entries, displaying..... | 543 |
| FEC filters, displaying..... | 546 |
| interfaces, displaying..... | 547 |
| LSPs, displaying..... | 551 |
| neighbors | |
| clearing connections..... | 540 |
| displaying..... | 549 |
| routes, displaying..... | 553 |
| sessions | |
| clearing..... | 541 |
| displaying..... | 555 |
| statistics | |
| clearing..... | 542 |
| displaying..... | 559 |
| traffic statistics, displaying..... | 562 |
| learning information | |
| Layer 2, displaying..... | 656, 657 |
| learning properties | |
| Layer 2, displaying..... | 658 |
| link-state paths <i>See</i> LSPs | |
| LLDP | |
| displaying..... | 258 |
| local information, displaying..... | 260 |
| neighbor | |
| clearing..... | 256 |
| neighbors, displaying..... | 262 |
| remote global statistics, displaying..... | 265 |
| statistics | |
| clearing..... | 257 |
| statistics, displaying..... | 267 |
| LSPs | |
| bandwidth allocation, adjusting..... | 569 |
| CAC information, displaying..... | 586 |
| clearing..... | 567 |
| LDP, displaying..... | 551 |
| MPLS, displaying..... | 592 |

M

| | |
|--------------------------------------------|-----|
| MAC address table | |
| Layer 2, clearing address information..... | 638 |
| Layer 2, displaying..... | 650 |

manuals
 comments on.....xxi
 martians, displaying.....423
 MDT
 displaying information.....184
 MLD
 group membership
 clearing.....76
 displaying.....118
 interfaces, displaying.....122
 PIM-to-MLD message translation information,
 displaying.....145
 statistics
 clearing.....77
 displaying.....125
 MPLS.....539
 administrative groups, displaying.....585
 CCC connections, displaying.....570
 CSPF statistics, displaying.....588
 DiffServ classes, displaying.....590
 interfaces, displaying.....591
 labels, displaying routes.....420
 link-management information, displaying
 all.....573
 peers.....576
 routing process.....578
 statistics.....581
 traffic-engineered links.....583
 LSPs, displaying.....600
 See also LDP, LSPs, RSVP, traffic engineering
 database
 MSDP
 cache entries, clearing.....78
 general information, displaying.....128
 message source information, displaying.....130
 peer statistics
 clearing.....79
 displaying.....134
 source-active cache, displaying.....132
 MSTP
 configuration
 displaying.....676
 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
 customer multicast routes, displaying.....698
 neighbors, displaying.....703
 routing instances, displaying.....700
 Multiprotocol Label Switching *See* MPLS
 MVPN
 displaying information.....187

N

neighbor
 ANCP.....4
 LLDP.....256
 next hops
 multicast entries, displaying.....142
 resolution database, displaying.....457
 routes sent to, displaying.....425
 notice icons defined.....xix

O

Open Shortest Path First *See* OSPF
 OSPF
 interfaces, displaying.....296
 link-state database entries, displaying
 version 2.....279
 version 3.....286
 neighbors
 clearing connections.....275
 displaying.....304
 overview
 displaying.....309
 routing table entries, displaying.....313
 SPF calculations, displaying.....302
 statistics, general
 clearing.....277
 displaying.....318
 statistics, I/O
 clearing.....274
 displaying.....301
 overload bit, resetting for IS-IS.....216

P

parentheses, in syntax descriptions.....xx
 PGM
 negative acknowledgments
 clearing.....86
 displaying.....168
 source path messages
 clearing.....87
 displaying.....170
 statistics
 clearing.....88
 displaying.....171
 PIM
 bootstrap routers, displaying.....174
 groups
 general information, displaying.....179
 usage information, displaying.....166
 interfaces
 displaying.....176
 join states, clearing.....89
 MDTs, displaying.....184
 MVPN, displaying.....187

| | | | |
|-----------------------------------------------------------------|-----|-----------------------------------------|-----|
| neighbors, displaying..... | 188 | RIPng | |
| PIM-to-IGMP message translation information, displaying..... | 144 | general statistics | |
| PIM-to-MLD message translation information, displaying..... | 145 | clearing..... | 496 |
| prune states, clearing..... | 89 | displaying..... | 498 |
| register | | neighbors | |
| clearing..... | 90 | displaying..... | 499 |
| RPF, displaying source state..... | 197 | statistics | |
| RPs | | clearing..... | 497 |
| displaying..... | 192 | displaying..... | 500 |
| statistics | | route advertisements, displaying..... | 339 |
| clearing..... | 91 | route, displaying | |
| displaying..... | 199 | next-hop..... | 425 |
| policers, displaying..... | 513 | router advertisements | |
| Pragmatic General Multicast <i>See</i> PGM | | IPv6 | |
| Protocol Independent Multicast <i>See</i> PIM | | clearing..... | 205 |
| prune states, clearing PIM..... | 89 | displaying..... | 207 |
| prunes, DVMRP, displaying..... | 99 | routes, displaying | |
| pseudowires | | active..... | 331 |
| redundant switchover..... | 687 | active path..... | 335 |
| | | advertising protocol..... | 339 |
| | | all..... | 343 |
| | | AS paths | |
| | | distribution of..... | 324 |
| | | domain information..... | 328 |
| | | regular expressions, matching..... | 344 |
| | | summary of..... | 330 |
| | | best..... | 346 |
| | | brief information..... | 349 |
| | | CCC..... | 351 |
| | | community ASN..... | 352 |
| | | community name..... | 353 |
| | | damping, BGP..... | 355 |
| | | detailed information..... | 360 |
| | | extensive information..... | 380 |
| | | flow validation..... | 392 |
| | | hidden..... | 405 |
| | | in a prefix range..... | 446 |
| | | in a specific routing table..... | 476 |
| | | in the forwarding table..... | 394 |
| | | in the LDP internal topology table..... | 553 |
| | | inactive path..... | 408 |
| | | inactive prefix..... | 411 |
| | | instances..... | 413 |
| | | learned from a specific address..... | 468 |
| | | learned from a specific protocol..... | 438 |
| | | learned from snooping..... | 460 |
| | | LSP..... | 422 |
| | | martian..... | 423 |
| | | matching the specified address..... | 374 |
| | | MPLS labels..... | 420 |
| | | next-hop resolution..... | 457 |
| | | not associated with a community..... | 431 |
| | | policy-based route export..... | 376 |
| | | received through a neighbor..... | 450 |
| | | sent to a specific interface..... | 434 |
| R | | | |
| redundant pseudowires, switchover..... | 687 | | |
| regular expressions | | | |
| AS paths, displaying matching routes..... | 344 | | |
| IP multicast scope | | | |
| clearing..... | 80 | | |
| IP multicast sessions | | | |
| clearing..... | 84 | | |
| displaying..... | 156 | | |
| LSPs, clearing..... | 567 | | |
| target communities, displaying..... | 378 | | |
| rendezvous points <i>See</i> RPs | | | |
| request l2circuit-switchover command..... | 687 | | |
| request mpls lsp adjust-autobandwidth command..... | 569 | | |
| Resource Reservation Protocol <i>See</i> RSVP | | | |
| reverse path forwarding <i>See</i> RPF | | | |
| RIP | | | |
| general statistics | | | |
| clearing..... | 488 | | |
| displaying..... | 490 | | |
| neighbors | | | |
| displaying..... | 491 | | |
| statistics | | | |
| clearing..... | 489 | | |
| displaying..... | 493 | | |

| | |
|---------------------------------------------------------------|-----|
| summary statistics..... | 474 |
| terse..... | 483 |
| Routing Information Protocol <i>See</i> RIP | |
| Routing Information Protocol next generation <i>See</i> RIPng | |
| routing policies | |
| displaying..... | 532 |
| testing the configuration for..... | 536 |
| RPF | |
| calculations, displaying..... | 150 |
| PIM source state, displaying..... | 197 |
| RPs | |
| displaying..... | 192 |
| RSVP | |
| interfaces, displaying..... | 613 |
| neighbors, displaying..... | 618 |
| sessions | |
| clearing..... | 610 |
| displaying..... | 622 |
| statistics | |
| clearing..... | 612 |
| displaying..... | 629 |
| version, displaying..... | 632 |

S

| | |
|------------------------------------------------|-----|
| SAP session announcements, displaying..... | 201 |
| Session Announcement Protocol <i>See</i> SAP | |
| shortest path first <i>See</i> SPF | |
| show (ospf ospf3) interface command..... | 296 |
| show (ospf ospf3) io-statistics command..... | 301 |
| show (ospf ospf3) log command..... | 302 |
| show (ospf ospf3) neighbor command..... | 304 |
| show (ospf ospf3) overview command..... | 309 |
| show (ospf ospf3) route command..... | 313 |
| show (ospf ospf3) statistics command..... | 318 |
| show ancp cos command..... | 5 |
| show ancp neighbor command..... | 8 |
| show ancp subscriber command..... | 10 |
| show as-path command..... | 324 |
| show as-path domain command..... | 328 |
| show as-path summary command..... | 330 |
| show bfd session command..... | 18 |
| show bgp bmp command..... | 30 |
| show bgp group command..... | 31 |
| show bgp group traffic-statistics command..... | 37 |
| show bgp neighbor command..... | 39 |
| show bgp summary command..... | 50 |
| show bridge domain command..... | 641 |
| show bridge flood command..... | 643 |
| show bridge mac-table command..... | 650 |
| show bridge statistics command..... | 653 |
| show connections command..... | 570 |
| show dhcp relay binding command..... | 521 |
| show dhcp relay statistics command..... | 525 |
| show dvmrp interfaces command..... | 93 |

| | |
|---------------------------------------------------|-----|
| show dvmrp neighbors command..... | 95 |
| show dvmrp prefix command..... | 97 |
| show dvmrp prunes command..... | 99 |
| show dynamic-tunnels database command..... | 688 |
| show esis adjacency command..... | 60 |
| show esis interface command..... | 62 |
| show esis statistics command..... | 64 |
| show firewall command..... | 507 |
| show firewall log command..... | 510 |
| show firewall prefix-action-stats command..... | 512 |
| show helper statistics command..... | 528 |
| show igmp group command..... | 101 |
| show igmp interface command..... | 104 |
| show igmp snooping interface command..... | 107 |
| show igmp snooping membership command..... | 109 |
| show igmp snooping statistics command..... | 112 |
| show igmp statistics command..... | 115 |
| show ipv6 neighbors command..... | 206 |
| show ipv6 router-advertisement command..... | 207 |
| show isis adjacency command..... | 219 |
| show isis authentication command..... | 223 |
| show isis backup coverage..... | 225 |
| show isis backup spf results..... | 228 |
| show isis database command..... | 231 |
| show isis hostname command..... | 237 |
| show isis interface command..... | 238 |
| show isis overview command..... | 241 |
| show isis route command..... | 244 |
| show isis spf command..... | 247 |
| show isis statistics command..... | 252 |
| show isis-backup label-switched-path command..... | 227 |
| show l2-learning global-information command..... | 656 |
| show l2-learning global-mac-count command..... | 657 |
| show l2-learning instance command..... | 658 |
| show l2-learning interface command..... | 659 |
| show l2circuit connections command..... | 690 |
| show l2vpn connections command..... | 694 |
| show ldp database command..... | 543 |
| show ldp fec-filters command..... | 546 |
| show ldp interface command..... | 547 |
| show ldp neighbor command..... | 549 |
| show ldp path command..... | 551 |
| show ldp route command..... | 553 |
| show ldp session command..... | 555 |
| show ldp statistics command..... | 559 |
| show ldp traffic-statistics command..... | 562 |
| show link-management command..... | 573 |
| show link-management peer command..... | 576 |
| show link-management routing command..... | 578 |
| show link-management statistics command..... | 581 |
| show link-management te-link command..... | 583 |
| show lldp command..... | 258 |
| show lldp local-information command..... | 260 |
| show lldp neighbors command..... | 262 |
| show lldp remote-global-statistics command..... | 265 |
| show lldp statistics command..... | 267 |

| | |
|-------------------------------------------------------|----------|
| show mac-rewrite interface command..... | 661 |
| show mld group command..... | 118 |
| show mld interface command..... | 122 |
| show mld statistics command..... | 125 |
| show mpls admin-groups command..... | 585 |
| show mpls call-admission-control command..... | 586 |
| show mpls cspf command..... | 588 |
| show mpls diffserv-te command..... | 590 |
| show mpls interface command..... | 591 |
| show mpls lsp command..... | 592 |
| show mpls path command..... | 600 |
| show msdp command..... | 128 |
| show msdp source command..... | 130 |
| show msdp source-active command..... | 132 |
| show msdp statistics command..... | 134 |
| show multicast flow-map command..... | 136 |
| show multicast interface command..... | 138 |
| show multicast mrinfo command..... | 140 |
| show multicast next-hops command..... | 142 |
| show multicast pim-to-igmp-proxy command..... | 144 |
| show multicast pim-to-mld-proxy command..... | 145 |
| show multicast route command..... | 146 |
| show multicast rpf command..... | 150 |
| show multicast scope command..... | 154 |
| show multicast sessions command..... | 156 |
| show multicast snooping route command..... | 158 |
| show multicast snooping statistics command..... | 160 |
| show multicast statistics command..... | 163 |
| show multicast usage command..... | 166 |
| show mvpn c-multicast..... | 698 |
| show mvpn instance command..... | 700 |
| show mvpn neighbor command..... | 703 |
| show ospf database command..... | 279 |
| show ospf3 database command..... | 286 |
| show pgm negative-acknowledgments command..... | 168 |
| show pgm source-path-messages command..... | 170 |
| show pgm statistics command..... | 171 |
| show pim bootstrap command..... | 174 |
| show pim interfaces command..... | 176 |
| show pim join command..... | 179 |
| show pim mdt command..... | 184, 187 |
| show pim neighbors command..... | 188 |
| show pim rps command..... | 192 |
| show pim source command..... | 197 |
| show pim statistics command..... | 199 |
| show policer command..... | 513 |
| show policy command..... | 532 |
| show policy conditions command..... | 534 |
| show policy damping command..... | 54 |
| show rip general-statistics command..... | 490 |
| show rip neighbor command..... | 491 |
| show rip statistics command..... | 493 |
| show ripng general-statistics command..... | 498 |
| show ripng neighbor command..... | 499 |
| show ripng statistics command..... | 500 |
| show route active-path command..... | 335 |
| show route advertising-protocol command..... | 339 |
| show route all command..... | 343 |
| show route aspath-regex command..... | 344 |
| show route best command..... | 346 |
| show route brief command..... | 349 |
| show route ccc command..... | 351 |
| show route command..... | 331 |
| show route community command..... | 352 |
| show route community-name command..... | 353 |
| show route damping command..... | 355 |
| show route detail command..... | 360 |
| show route exact command..... | 374 |
| show route export command..... | 376 |
| show route export vrf-target command..... | 378 |
| show route extensive command..... | 380 |
| show route flow validation command..... | 392 |
| show route forwarding-table command..... | 394 |
| show route hidden command..... | 405 |
| show route inactive-path command..... | 408 |
| show route inactive-prefix command..... | 411 |
| show route instance command..... | 413 |
| show route label command..... | 420 |
| show route label-switched-path command..... | 422 |
| show route martians command..... | 423 |
| show route next-hop command..... | 425 |
| show route no-community command..... | 431 |
| show route output command..... | 434 |
| show route protocol command..... | 438 |
| show route range command..... | 446 |
| show route receive-protocol command..... | 450 |
| show route resolution command..... | 457 |
| show route snooping command..... | 460 |
| show route source-gateway command..... | 468 |
| show route summary command..... | 474 |
| show route table command..... | 476 |
| show route terse command..... | 483 |
| show rsvp interface command..... | 613 |
| show rsvp neighbor command..... | 618 |
| show rsvp session command..... | 622 |
| show rsvp statistics command..... | 629 |
| show rsvp version command..... | 632 |
| show sap listen command..... | 201 |
| show spanning-tree bridge command..... | 666 |
| show spanning-tree interface command..... | 671 |
| show spanning-tree mstp configuration command..... | 676 |
| show spanning-tree statistics command..... | 678 |
| show ted database command..... | 601 |
| show ted link command..... | 605 |
| show ted protocol command..... | 607 |
| show vlpls connections command..... | 706 |
| show vlpls flood event-queue command..... | 711 |
| show vlpls flood instance command..... | 712 |
| show vlpls flood route command..... | 714 |
| show vlpls mac-table command..... | 716 |

show vpls statistics command.....720
 snooping (interface)
 IGMP.....107
 snooping (membership)
 IGMP.....109
 snooping (statistics)
 IGMP.....112
 snooping routes, displaying.....460
 source gateway addresses, displaying.....468
 Spanning Tree Protocol *See* STP
 SPF calculations, displaying.....302
 statistics
 Layer 2 bridging, displaying.....653
 LLDP.....257
 STP
 bridge
 displaying.....666
 interface
 displaying.....671
 protocol
 clearing.....664
 statistics
 clearing.....665
 displaying.....678
 support, technical *See* technical support
 syntax conventions.....xx

T

technical support
 contacting JTAC.....xxi
 TED *See* traffic engineering database
 test msdp command.....202
 test policy command.....536
 traffic engineering database
 database entries, displaying.....601
 link information, displaying.....605
 protocols learned from, displaying.....607

U

UDP, statistics
 clearing.....520
 displaying.....528
 User Datagram Protocol *See* UDP

V

version
 RSVP, displaying.....632
 virtual private LAN service *See* VPLS
 VPLS
 connections, displaying.....706
 statistics, displaying.....720
 VPN routing and forwarding *See* VRF
 VPNs *See* Layer 2 VPNs, Layer 3 VPNs

VRF
 route export, displaying.....378

Index of Statements and Commands

C

| | |
|--------------------------------------------------|-----|
| clear (ospf ospf3) database command..... | 271 |
| clear (ospf ospf3) io-statistics command..... | 274 |
| clear (ospf ospf3) neighbor command..... | 275 |
| clear (ospf ospf3) statistics command..... | 277 |
| clear ancp neighbor command..... | 4 |
| clear bfd adaptation command..... | 16 |
| clear bfd session command..... | 17 |
| clear bgp command..... | 29 |
| clear bgp damping command..... | 26 |
| clear bgp neighbor command..... | 27 |
| clear bridge mac-table command..... | 638 |
| clear dhcp relay binding command..... | 516 |
| clear dhcp relay statistics command..... | 518 |
| clear error bpdu command..... | 639 |
| clear error mac-rewrite command..... | 640 |
| clear esis adjacency command..... | 58 |
| clear esis statistics command..... | 59 |
| clear firewall command..... | 506 |
| clear helper statistics command..... | 520 |
| clear igmp membership command..... | 69 |
| clear igmp snooping membership command..... | 72 |
| clear igmp snooping statistics command..... | 73 |
| clear igmp statistics command..... | 74 |
| clear ipv6 neighbors command..... | 204 |
| clear ipv6 router-advertisement command..... | 205 |
| clear isis adjacency command..... | 213 |
| clear isis database command..... | 214 |
| clear isis overload command..... | 216 |
| clear isis statistics command..... | 217 |
| clear ldp neighbor command..... | 540 |
| clear ldp session command..... | 541 |
| clear ldp statistics command..... | 542 |
| clear lldp neighbor command..... | 256 |
| clear lldp statistics command..... | 257 |
| clear mld membership command..... | 76 |
| clear mld statistics command..... | 77 |
| clear msdp cache command..... | 78 |
| clear msdp statistics command..... | 79 |
| clear multicast bandwidth-admission command..... | 82 |
| clear multicast scope command..... | 80 |
| clear multicast sessions command..... | 84 |
| clear multicast snooping statistics command..... | 81 |
| clear multicast statistics command..... | 85 |

| | |
|--------------------------------------------------------|-----|
| clear ospf overload command..... | 276 |
| clear pgm negative-acknowledgments command..... | 86 |
| clear pgm source-path-messages command..... | 87 |
| clear pgm statistics command..... | 88 |
| clear pim join command..... | 89 |
| clear pim register command..... | 90 |
| clear pim statistics command..... | 91 |
| clear rip general-statistics command..... | 488 |
| clear rip statistics command..... | 489 |
| clear ripng general-statistics command..... | 496 |
| clear ripng statistics command..... | 497 |
| clear rsvp session command..... | 610 |
| clear rsvp statistics command..... | 612 |
| clear spanning-tree protocol-migration command..... | 664 |
| clear spanning-tree statistics command..... | 665 |
| clear vpls mac-address command..... | 685 |
| clear vpls mac-table command..... | 686 |

R

| | |
|-------------------------------------------------------|-----|
| request l2circuit-switchover command..... | 687 |
| request mpls lsp adjust-autobandwidth command..... | 569 |

S

| | |
|------------------------------------------------|-----|
| show (ospf ospf3) interface command..... | 296 |
| show (ospf ospf3) io-statistics command..... | 301 |
| show (ospf ospf3) log command..... | 302 |
| show (ospf ospf3) neighbor command..... | 304 |
| show (ospf ospf3) overview command..... | 309 |
| show (ospf ospf3) route command..... | 313 |
| show (ospf ospf3) statistics command..... | 318 |
| show ancp cos command..... | 5 |
| show ancp neighbor command..... | 8 |
| show ancp subscriber command..... | 10 |
| show as-path command..... | 324 |
| show as-path domain command..... | 328 |
| show as-path summary command..... | 330 |
| show bfd session command..... | 18 |
| show bgp bmp command..... | 30 |
| show bgp group command..... | 31 |
| show bgp group traffic-statistics command..... | 37 |
| show bgp neighbor command..... | 39 |
| show bgp summary command..... | 50 |

| | |
|---------------------------------------------------|----------|
| show bridge domain command..... | 641 |
| show bridge flood command..... | 643 |
| show bridge mac-table command..... | 650 |
| show bridge statistics command..... | 653 |
| show connections command..... | 570 |
| show dhcp relay binding command..... | 521 |
| show dhcp relay statistics command..... | 525 |
| show dvmp interfaces command..... | 93 |
| show dvmp neighbors command..... | 95 |
| show dvmp prefix command..... | 97 |
| show dvmp prunes command..... | 99 |
| show dynamic-tunnels database command..... | 688 |
| show esis adjacency command..... | 60 |
| show esis interface command..... | 62 |
| show esis statistics command..... | 64 |
| show firewall command..... | 507 |
| show firewall log command..... | 510 |
| show firewall prefix-action-stats command..... | 512 |
| show helper statistics command..... | 528 |
| show igmp group command..... | 101 |
| show igmp interface command..... | 104 |
| show igmp snooping interface command..... | 107 |
| show igmp snooping membership command..... | 109 |
| show igmp snooping statistics command..... | 112 |
| show igmp statistics command..... | 115 |
| show ipv6 neighbors command..... | 206 |
| show ipv6 router-advertisement command..... | 207 |
| show isis adjacency command..... | 219 |
| show isis authentication command..... | 223 |
| show isis backup coverage..... | 225 |
| show isis backup spf results..... | 228 |
| show isis database command..... | 231 |
| show isis hostname command..... | 237 |
| show isis interface command..... | 238 |
| show isis overview command..... | 241 |
| show isis route command..... | 244 |
| show isis spf command..... | 247 |
| show isis statistics command..... | 252 |
| show isis-backup label-switched-path command..... | 227 |
| show l2-learning global-information command..... | 656 |
| show l2-learning global-mac-count command..... | 657 |
| show l2-learning instance command..... | 658 |
| show l2-learning interface command..... | 659 |
| show l2circuit connections command..... | 690 |
| show l2vpn connections command..... | 694 |
| show ldp database command..... | 543 |
| show ldp fec-filters command..... | 546 |
| show ldp interface command..... | 547 |
| show ldp neighbor command..... | 549 |
| show ldp path command..... | 551 |
| show ldp route command..... | 553 |
| show ldp session command..... | 555 |
| show ldp statistics command..... | 559 |
| show ldp traffic-statistics command..... | 562 |
| show link-management peer command..... | 573 |
| show link-management routing command..... | 578 |
| show link-management statistics command..... | 581 |
| show link-management te-link command..... | 583 |
| show lldp command..... | 258 |
| show lldp local-information command..... | 260 |
| show lldp neighbors command..... | 262 |
| show lldp remote-global-statistics command..... | 265 |
| show lldp statistics command..... | 267 |
| show mac-rewrite interface command..... | 661 |
| show mld group command..... | 118 |
| show mld interface command..... | 122 |
| show mld statistics command..... | 125 |
| show mpls admin-groups command..... | 585 |
| show mpls call-admission-control command..... | 586 |
| show mpls cspf command..... | 588 |
| show mpls diffserv-te command..... | 590 |
| show mpls interface command..... | 591 |
| show mpls lsp command..... | 592 |
| show mpls path command..... | 600 |
| show msdp command..... | 128 |
| show msdp source command..... | 130 |
| show msdp source-active command..... | 132 |
| show msdp statistics command..... | 134 |
| show multicast flow-map command..... | 136 |
| show multicast interface command..... | 138 |
| show multicast minfo command..... | 140 |
| show multicast next-hops command..... | 142 |
| show multicast pim-to-igmp-proxy command..... | 144 |
| show multicast pim-to-mld-proxy command..... | 145 |
| show multicast route command..... | 146 |
| show multicast rpf command..... | 150 |
| show multicast scope command..... | 154 |
| show multicast sessions command..... | 156 |
| show multicast snooping route command..... | 158 |
| show multicast snooping statistics command..... | 160 |
| show multicast statistics command..... | 163 |
| show multicast usage command..... | 166 |
| show ospf database command..... | 279 |
| show ospf3 database command..... | 286 |
| show pgm negative-acknowledgments command..... | 168 |
| show pgm source-path-messages command..... | 170 |
| show pgm statistics command..... | 171 |
| show pim bootstrap command..... | 174 |
| show pim interfaces command..... | 176 |
| show pim join command..... | 179 |
| show pim mdt command..... | 184, 187 |
| show pim neighbors command..... | 188 |
| show pim rps command..... | 192 |
| show pim source command..... | 197 |
| show pim statistics command..... | 199 |
| show policer command..... | 513 |
| show policy command..... | 532 |
| show policy conditions command..... | 534 |
| show policy damping command..... | 54 |
| show rip general-statistics command..... | 490 |

| | | | |
|-------------------------------------------------------|-----|------------------------------------------|-----|
| show rip neighbor command..... | 491 | show vpls connections command..... | 706 |
| show rip statistics command..... | 493 | show vpls flood event-queue command..... | 711 |
| show ripng general-statistics command..... | 498 | show vpls flood instance command..... | 712 |
| show ripng neighbor command..... | 499 | show vpls flood route command..... | 714 |
| show ripng statistics command..... | 500 | show vpls mac-table command..... | 716 |
| show route active-path command..... | 335 | show vpls statistics command..... | 720 |
| show route advertising-protocol command..... | 339 | | |
| show route all command..... | 343 | T | |
| show route aspath-regex command..... | 344 | test msdp command..... | 202 |
| show route best command..... | 346 | test policy command..... | 536 |
| show route brief command..... | 349 | | |
| show route ccc command..... | 351 | | |
| show route command..... | 331 | | |
| show route community command..... | 352 | | |
| show route community-name command..... | 353 | | |
| show route damping command..... | 355 | | |
| show route detail command..... | 360 | | |
| show route exact command..... | 374 | | |
| show route export command..... | 376 | | |
| show route export vrf-target command..... | 378 | | |
| show route extensive command..... | 380 | | |
| show route flow validation command..... | 392 | | |
| show route forwarding-table command..... | 394 | | |
| show route hidden command..... | 405 | | |
| show route inactive-path command..... | 408 | | |
| show route inactive-prefix command..... | 411 | | |
| show route instance command..... | 413 | | |
| show route label command..... | 420 | | |
| show route label-switched-path command..... | 422 | | |
| show route martians command..... | 423 | | |
| show route next-hop command..... | 425 | | |
| show route no-community command..... | 431 | | |
| show route output command..... | 434 | | |
| show route protocol command..... | 438 | | |
| show route range command..... | 446 | | |
| show route receive-protocol command..... | 450 | | |
| show route resolution command..... | 457 | | |
| show route snooping command..... | 460 | | |
| show route source-gateway command..... | 468 | | |
| show route summary command..... | 474 | | |
| show route table command..... | 476 | | |
| show route terse command..... | 483 | | |
| show rsvp interface command..... | 613 | | |
| show rsvp neighbor command..... | 618 | | |
| show rsvp session command..... | 622 | | |
| show rsvp statistics command..... | 629 | | |
| show rsvp version command..... | 632 | | |
| show sap listen command..... | 201 | | |
| show spanning-tree bridge command..... | 666 | | |
| show spanning-tree interface command..... | 671 | | |
| show spanning-tree mstp configuration command..... | 676 | | |
| show spanning-tree statistics command..... | 678 | | |
| show ted database command..... | 601 | | |
| show ted link command..... | 605 | | |
| show ted protocol command..... | 607 | | |

