



JUNOS® Software

MPLS Applications Configuration Guide

Release 9.6

Juniper Networks, Inc.

1194 North Mathilda Avenue
Sunnyvale, California 94089
USA

408-745-2000

www.juniper.net

Published: 2009-07-22

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 MPLS Applications Configuration Guide

Release 9.6

Copyright © 2009, Juniper Networks, Inc.

All rights reserved. Printed in USA.

Writing: Albert Statti

Editing: Sonia Saruba, Laura Singer

Illustration: Nathaniel Woodward

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 NON-INFRINGEMENT. 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	xxxi
Part 1	Overview	
Chapter 1	Traffic Engineering Overview	3
Chapter 2	Complete MPLS Applications Configuration Statements	9
Part 2	MPLS	
Chapter 3	MPLS Overview	21
Chapter 4	MPLS Router Configuration Guidelines	51
Chapter 5	MPLS-Signaled LSP Configuration Guidelines	69
Chapter 6	DiffServ-Aware Traffic Engineering Configuration Guidelines	109
Chapter 7	Static and Explicit-Path LSP Configuration Guidelines	133
Chapter 8	Point-to-Multipoint LSP Configuration Guidelines	141
Chapter 9	Miscellaneous MPLS Properties Configuration Guidelines	149
Chapter 10	Summary of MPLS Configuration Statements	173
Part 3	RSVP	
Chapter 11	RSVP Overview	255
Chapter 12	RSVP Configuration Guidelines	273
Chapter 13	Summary of RSVP Configuration Statements	299
Part 4	LDP	
Chapter 14	LDP Overview	329
Chapter 15	LDP Configuration Guidelines	337
Chapter 16	Summary of LDP Configuration Statements	367
Part 5	CCC and TCC	
Chapter 17	CCC and TCC Overview	401
Chapter 18	CCC and TCC Configuration Guidelines	405
Chapter 19	Summary of CCC and TCC Configuration Statements	427
Part 6	GMPLS	
Chapter 20	GMPLS Overview	443

Chapter 21	GMPLS Configuration Guidelines	449
Chapter 22	Hierarchy of RSVP LSPs Configuration Guidelines	465
Chapter 23	Summary of GMPLS Configuration Statements	471
Part 7	Indexes	
	Index	491
	Index of Statements and Commands	507

Table of Contents

	About This Guide	xxxi
	JUNOS Documentation and Release Notes	xxxi
	Objectives	xxxii
	Audience	xxxii
	Supported Platforms	xxxii
	Using the Indexes	xxxiii
	Using the Examples in This Manual	xxxiii
	Merging a Full Example	xxxiii
	Merging a Snippet	xxxiv
	Documentation Conventions	xxxiv
	Documentation Feedback	xxxvi
	Requesting Technical Support	xxxvii
Part 1	Overview	
Chapter 1	Traffic Engineering Overview	3
	Traffic Engineering Capabilities	3
	Components of Traffic Engineering	4
	Packet Forwarding Component	4
	Packet Forwarding Based on Label Swapping	4
	How a Packet Traverses an MPLS Backbone	5
	Information Distribution Component	5
	Path Selection Component	6
	Offline Planning and Analysis	6
	Signaling Component	7
	Flexible LSP Calculation and Configuration	7
Chapter 2	Complete MPLS Applications Configuration Statements	9
	[edit logical-systems] Hierarchy Level	9
	[edit protocols connections] Hierarchy Level	10
	[edit protocols ldp] Hierarchy Level	10
	[edit protocols link-management] Hierarchy Level	12
	[edit protocols mpls] Hierarchy Level	13
	[edit protocols rsvp] Hierarchy Level	17

Part 2**MPLS****Chapter 3****MPLS Overview****21**

MPLS Introduction	22
MPLS Standards	22
Link-Layer Support	24
MPLS and Traffic Engineering	24
Label Description	25
Special Labels	25
Label Allocation	26
Operations on Labels	27
Routers in an LSP	28
How a Packet Travels Along an LSP	28
Types of LSPs	29
Scope of LSPs	29
Constrained-Path LSP Computation	29
How CSPF Selects a Path	31
Path Selection Tie-Breaking	32
Computing Paths Offline	32
LSPs on an Overloaded Router	33
Fate Sharing	33
IGP Shortcuts	34
Enabling IGP Shortcuts	35
LSPs Qualified in Shortcut Computations	36
IGP Shortcut Applications	36
IGP Shortcuts and Routing Table	37
IGP Shortcuts and VPN Environments	37
Advertising LSPs into IGP	37
IP and MPLS Packets on Aggregated Interfaces	38
MPLS Applications	39
BGP Destinations	39
IGP and BGP Destinations	41
Selecting a Forwarding LSP Next Hop	41
MPLS and Routing Tables	41
MPLS and Traffic Protection	43
Fast Reroute	44
Fast Reroute Overview	44
Detour Merging Process	46
Detour Computations	47
Fast Reroute Path Optimization	47
Automatic Bandwidth Allocation	48
Point-to-Multipoint LSPs	48
MPLS Load Balancing Based on the IP Header and MPLS Labels	50

Chapter 4**MPLS Router Configuration Guidelines****51**

Minimum MPLS Configuration	51
Configuring the Ingress Router for MPLS-Signaled LSPs	52
Creating Named Paths	52
Examples: Creating Named Paths	54
Configuring Alternate Backup Paths Using Fate Sharing	54
Configuring Fate Sharing	55
Implications for CSPF	56
Example: Configuring Fate Sharing	56
Examples: Configuring Ingress Routers for MPLS-Signaled LSPs	57
Example: Configuring a Constrained-Path LSP for Which JUNOS Makes All Forwarding Decisions	57
Example: Configuring an Explicit-Path LSP	57
Example: Configuring a Constrained-Path LSP for Which JUNOS Makes Most Forwarding Decisions and Considers Hop Constraints	58
Example: Configuring a Constrained-Path LSP for Which JUNOS Makes Most Forwarding Decisions and the Secondary Path Is Explicit	59
Configuring the Intermediate and Egress Routers for MPLS-Signaled LSPs	60
Configuring a Default Route for Unlabeled MPLS Packets	60
Improving Traffic Engineering Database Accuracy with RSVP PathErr Messages	60
PathErr Messages	61
Identifying the Problem Link	62
Configuring the Router to Improve Traffic Engineering Database Accuracy	62
Configuring MPLS-Signaled LSPs to Use GRE Tunnels	62
Example: Configuring MPLS-Signaled LSPs to Use GRE Tunnels	63
Tunneling IPv6 Traffic over MPLS IPv4 Networks	64
IPv6 over MPLS Standards	66
Configuring IPv4 MPLS Tunnels to Carry IPv6 Traffic	66
Configuring IPv6 on Both Core-Facing and CE Router-Facing Interfaces	66
Configuring MPLS and RSVP Between PE Routers	66
Enabling IPv6 Tunneling on PE Routers	67
Configuring Multiprotocol BGP to Carry IPv6 Traffic	67
Configuring ICMP Message Tunneling	67

Chapter 5**MPLS-Signaled LSP Configuration Guidelines****69**

LSP Configuration Overview	70
Configuring the Ingress and Egress Router Addresses for LSPs	73
Configuring the Ingress Router Address for LSPs	73
Configuring the Egress Router Address for LSPs	74
Preventing the Addition of Egress Router Addresses to Routing Tables	74
Configuring Primary and Secondary LSPs	75
Configuring Primary and Secondary Paths for an LSP	75
Configuring the Revert Timer for LSPs	76
Specifying the Conditions for Path Selection	77
Configuring a Text Description for LSPs	78
Configuring Fast Reroute	78
Configuring the Optimization Interval for Fast Reroute Paths	79
Adding LSP-Related Routes to the inet.3 Routing Table	80
Configuring the Connection Between Ingress and Egress Routers	81
Configuring LSP Metrics	82
Configuring Dynamic LSP Metrics	82
Configuring Static LSP Metrics	82
Configuring CSPF Tie Breaking	83
Configuring Load Balancing for MPLS LSPs	84
Using the First MPLS Label in the Hash Key	84
Using the Second MPLS Label in the Hash Key	84
Using the Third MPLS Label in the Hash Key	85
Using the IP Payload in the Hash Key	85
Using the First Two Labels and the IP Payload in the Hash Key	85
Configuring Load Balancing for MPLS LSPs Without CSPF	86
Disabling Normal TTL Decrementing	86
Configuring MPLS Soft Preemption	87
Configuring Automatic Bandwidth Allocation for LSPs	88
Configuring MPLS Statistics for Automatic Bandwidth Allocation	89
Configuring Automatic Bandwidth Allocation on LSPs	89
Configuring the Automatic Bandwidth Allocation Interval	90
Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth	90
Configuring the Automatic Bandwidth Adjustment Threshold	91
Configuring a Limit on Bandwidth Overflow Samples	91
Configuring Passive Bandwidth Utilization Monitoring	93
Requesting Automatic Bandwidth Allocation Adjustment	94
Disabling Constrained-Path LSP Computation	95
Configuring Administrative Groups	96
Configuring Preference Values for LSPs	98
Disabling Path Route Recording	98
Configuring Class of Service for MPLS LSPs	98
Class of Service for MPLS Overview	99
Configuring the MPLS CoS Bits	99
Rewriting IEEE 802.1p Packet Headers with the MPLS CoS Value	100
Configuring Adaptive LSPs	101

Configuring Priority and Preemption for LSPs	102
Optimizing Signaled LSPs	103
Configuring the Smart Optimize Timer	105
Limiting the Number of Hops in LSPs	105
Configuring the Bandwidth Value for LSPs	106
Configuring Hot Standby of Secondary Paths	106
Damping Advertisement of LSP State Changes	107

Chapter 6

DiffServ-Aware Traffic Engineering Configuration Guidelines 109

DiffServ-Aware Traffic Engineering Introduction	110
DiffServ-Aware Traffic Engineering Standards	110
DiffServ-Aware Traffic Engineering Overview	112
DiffServ-Aware Traffic Engineering Features	112
DiffServ-Aware Traffic Engineered LSPs	113
DiffServ-Aware Traffic Engineered LSPs Overview	113
DiffServ-Aware Traffic Engineered LSPs Operation	113
Multiclass LSPs	114
Multiclass LSP Overview	114
Establishing a Multiclass LSP on the Differentiated Services Domain	115
Configuring Routers for DiffServ-Aware Traffic Engineering	115
Configuring the Bandwidth Model	116
Configuring Traffic Engineering Classes	117
Requirements and Limitations for the Traffic Engineering Class Matrix	118
Configuring Class of Service for Diffserv-Aware Traffic Engineering	119
Bandwidth Oversubscription Overview	119
LSP Size Oversubscription	120
Link Size Oversubscription	120
Class Type Oversubscription and Local Oversubscription Multipliers	121
Class Type Bandwidth and the LOM	121
LOM Calculation for the MAM and Extended MAM Bandwidth Models	122
LOM Calculation for the Russian Dolls Bandwidth Model	122
Example: LOM Calculation	122
Configuring the Bandwidth Subscription Percentage for LSPs	124
Constraints on Configuring Bandwidth Subscription	125
Configuring LSPs for DiffServ-Aware Traffic Engineering	125
Configuring Class of Service for the Interfaces	126
Configuring IGP	126
Configuring Traffic-Engineered LSPs	127
Configuring Policing for LSPs	127
Configuring Fast Reroute for Traffic-Engineered LSPs	128
Configuring Multiclass LSPs	128
Configuring Class of Service for the Interfaces	129
Configuring the IGP	129
Configuring Class-Type Bandwidth Constraints for Multiclass LSPs	129
Configuring Policing for Multiclass LSPs	131
Configuring Fast Reroute for Multiclass LSPs	131

Chapter 7	Static and Explicit-Path LSP Configuration Guidelines	133
	Configuring Static LSPs	133
	Configuring the Ingress Router for Static LSPs	133
	Example: Configuring the Ingress Router	135
	Configuring the Intermediate and Egress Routers for Static LSPs	136
	Example: Configuring an Intermediate Router	137
	Example: Configuring an Egress Router	138
	Configuring Static Unicast Routes for Point-to-Multipoint LSPs	138
	Configuring Explicit-Path LSPs	140
Chapter 8	Point-to-Multipoint LSP Configuration Guidelines	141
	Configuring Primary and Branch LSPs for Point-to-Multipoint LSPs	141
	Configuring the Primary Point-to-Multipoint LSP	141
	Configuring a Branch LSP for Point-to-Multipoint LSPs	142
	Configuring the Branch LSP as a Dynamic Path	142
	Configuring the Branch LSP as a Static Path	143
	Example: Configuring a Point-to-Multipoint LSP	143
	Configuring Link Protection for Point-to-Multipoint LSPs	144
	Configuring Graceful Restart for Point-to-Multipoint LSPs	144
	Configuring a Multicast RPF Check Policy for Point-to-Multipoint LSPs	145
	Example: Configuring Multicast RPF Check Policy for a Point-to-Multipoint LSP	146
	Configuring Ingress PE Router Redundancy for Point-to-Multipoint LSPs	146
	Enabling Point-to-Point LSPs to Monitor Egress PE Routers	147
	Preserving Point-to-Multipoint LSP Functioning with Different JUNOS Software Releases	147
Chapter 9	Miscellaneous MPLS Properties Configuration Guidelines	149
	Configuring MPLS to Pop the Label on the Ultimate-Hop Router	149
	Configuring Traffic Engineering for LSPs	150
	Using RSVP and LDP Routes for Traffic Forwarding	150
	Using RSVP and LDP Routes for Forwarding in Virtual Private Networks	151
	Using RSVP and LDP Routes for Forwarding But Not Route Selection	151
	Advertising the LSP Metric in Summary LSAs	152
	Enabling Interarea Traffic Engineering	152
	Enabling Inter-AS Traffic Engineering for LSPs	153
	Inter-AS Traffic Engineering Requirements	153
	Inter-AS Traffic Engineering Limitations	154
	Configuring OSPF Passive TE Mode	155
	Configuring MPLS to Gather Statistics	156
	Configuring System Log Messages and SNMP Traps for LSPs	157
	Configuring MPLS Firewall Filters and Policers	158
	Configuring MPLS Firewall Filters	158
	Examples: Configuring MPLS Firewall Filters	159

Configuring Policers for LSPs	160
LSP Policer Limitations	161
Example: Configuring an LSP Policer	161
Configuring Automatic Policers	162
Configuring Automatic Policers for LSPs	163
Configuring Automatic Policers for DiffServ-Aware Traffic Engineering LSPs	164
Configuring Automatic Policers for Point-to-Multipoint LSPs	164
Disabling Automatic Policing on an LSP	164
Example: Configuring Automatic Policing for an LSP	165
Writing Different DSCP and EXP Values in MPLS-Tagged IP Packets	165
Configuring MPLS Rewrite Rules	165
Rewriting the EXP Bits of All Three Labels of an Outgoing Packet	166
Rewriting MPLS and IPv4 Packet Headers	166
Configuring BFD for MPLS IPv4 LSPs	167
Configuring BFD for RSVP-Signaled LSPs	167
Configuring a Failure Action for the BFD Session on an RSVP LSP	169
Pinging LSPs	170
Pinging MPLS LSPs	170
Pinging Point-to-Multipoint LSPs	170
Pinging the Endpoint Address of MPLS LSPs	171
Pinging CCC LSPs	171
Pinging Layer 3 VPNs	171
Support for LSP Ping and Traceroute Commands Based on RFC 4379	171
Tracing MPLS and LSP Packets and Operations	171

Chapter 10

Summary of MPLS Configuration Statements **173**

adaptive	173
adjust-interval	174
adjust-threshold	174
adjust-threshold-overflow-limit	175
admin-down	175
admin-group	176
admin-group (for Interfaces)	176
admin-group (for LSPs)	177
admin-groups	178
advertisement-hold-time	178
allow-fragmentation	179
associate-backup-pe-groups	179
auto-bandwidth	180
auto-policing	181
backup-pe-group	182
bandwidth	183
bandwidth-model	184
bandwidth-percent	185
bfd-liveness-detection	186
class-of-service	187
default-route	188

description	188
diffserv-te	189
disable	190
discard	190
double-push	191
encoding-type	191
exclude	192
exclude (for Administrative Groups)	192
exclude (for Fast Reroute)	193
expand-loose-hop	193
explicit-null	194
failure-action	195
fast-reroute	196
fate-sharing	197
from	198
gpip	199
hop-limit	200
icmp-tunneling	200
include-all	201
include-all (for Administrative Groups)	201
include-all (for Fast Reroute)	202
include-any	203
include-any (for Administrative Groups)	203
include-any (for Fast Reroute)	204
install	204
interface	205
ipv6-tunneling	206
label-map	206
label-switched-path	207
ldp-tunneling	209
least-fill	209
link-protection (MPLS)	210
log-updown	211
lsp-attributes	212
maximum-bandwidth	212
metric	213
minimum-bandwidth	213
monitor-bandwidth	214
most-fill	214
mpls	214
mtu-signaling	215
next-hop	215
no-cspf	216
no-decrement-ttl	217
no-exclude	217
no-include-all	217
no-include-any	217
no-install-to-address	218
no-propagate-ttl	218
no-record	219
no-trap	219

oam	220
optimize-aggressive	221
optimize-timer	222
p2mp	223
p2mp-lsp-next-hop	223
path	224
path-mtu	225
policing	225
pop	226
preference	227
primary	228
priority	229
push	230
random	231
record	232
reject	232
retry-limit	233
retry-timer	233
revert-timer	234
rpf-check-policy	235
rsvp-error-hold-time	236
secondary	237
select	238
signal-bandwidth	238
smart-optimize-timer	239
soft-preemption	239
standby	240
static-path	241
statistics	242
swap	243
swap-push	244
switching-type	245
te-class-matrix	246
to	247
traceoptions	248
traffic-engineering	250
triple-push	251

Part 3

RSVP

Chapter 11

RSVP Overview

255

RSVP Introduction	256
RSVP Standards	256
JUNOS Software RSVP Protocol Implementation	257
RSVP Operation	258
RSVP Operation Overview	258
RSVP Authentication	259

RSVP and IGP Hello Packets and Timers	259
RSVP Message Types	259
Path Messages	260
Resv Messages	260
PathTear Messages	260
ResvTear Messages	260
PathErr Messages	261
ResvErr Messages	261
ResvConfirm Messages	261
RSVP Reservation Styles	261
RSVP Refresh Reduction	262
MTU Signaling in RSVP	263
How the Correct MTU Is Signaled in RSVP	264
Determining an Outgoing MTU Value	265
MTU Signaling in RSVP Limitations	265
Link Protection	265
Fast Reroute, Node Protection, and Link Protection	266
Multiple Bypass LSPs	267
Node Protection	267
RSVP Graceful Restart	268
RSVP Graceful Restart Standard	269
RSVP Graceful Restart Operation	270
Processing the Restart Cap Object	271

Chapter 12

RSVP Configuration Guidelines

273

Minimum RSVP Configuration	273
Configuring RSVP and MPLS	274
Example: Configuring RSVP and MPLS	274
Configuring RSVP Interfaces	275
Configuring RSVP Refresh Reduction	275
Determining the Refresh Reduction Capability of RSVP	
Neighbors	277
Configuring the RSVP Hello Interval	278
Configuring RSVP Authentication	278
Configuring the Bandwidth Subscription for Class Types	279
Configuring the RSVP Update Threshold on an Interface	279
Configuring RSVP for Unnumbered Interfaces	280
Configuring Node Protection or Link Protection for LSPs	281
Configuring Inter-AS Node and Link Protection	281
Configuring Link Protection on Interfaces Used by LSPs	282
Configuring Bypass LSPs	283
Configuring the Next-Hop or Next-Next-Hop Node Address for Bypass	
LSPs	284
Configuring Administrative Groups for Bypass LSPs	284
Configuring the Bandwidth for Bypass LSPs	284
Configuring Class of Service for Bypass LSPs	285
Configuring the Hop Limit for Bypass LSPs	285
Configuring the Maximum Number of Bypass LSPs	286
Disabling CSPF for Bypass LSPs	287

Disabling Node Protection for Bypass LSPs	287
Configuring the Optimization Interval for Bypass LSPs	287
Configuring an Explicit Path for Bypass LSPs	288
Configuring the Amount of Bandwidth Subscribed for Bypass LSPs	288
Configuring Priority and Preemption for Bypass LSPs	289
Configuring RSVP Graceful Restart	289
Enabling Graceful Restart for All Routing Protocols	290
Disabling Graceful Restart for RSVP	290
Disabling RSVP Helper Mode	290
Configuring the Maximum Helper Recovery Time	290
Configuring the Maximum Helper Restart Time	290
Configuring Load Balancing Across RSVP LSPs	291
Configuring Timers for RSVP Refresh Messages	292
Preempting RSVP Sessions	293
Configuring MTU Signaling in RSVP	293
Enabling MTU Signaling in RSVP	294
Enabling Packet Fragmentation	294
Configuring RSVP to Pop the Label on the Ultimate-Hop Router	295
Disabling Adjacency Down and Neighbor Down Notification in IS-IS and OSPF	295
Enabling Ultimate-Hop Popping on Point-to-Multipoint LSPs	296
Tracing RSVP Protocol Traffic	297
Examples: Tracing RSVP Protocol Traffic	298

Chapter 13

Summary of RSVP Configuration Statements **299**

admin-group	299
aggregate	300
authentication-key	301
bandwidth	302
bypass	303
class-of-service	304
disable	305
fast-reroute optimize-timer	306
graceful-deletion-timeout	306
graceful-restart	307
hello-interval	308
hop-limit	308
interface	309
keep-multiplier	310
link-protection (RSVP)	311
load-balance	312
max-bypasses	313
no-adjacency-down-notification	313
no-aggregate	314
no-cspf	314
no-neighbor-down-notification	314
no-node-id-subobject	315
no-p2mp-sublsp	315
no-reliable	315

node-link-protection	316
optimize-timer	316
path	317
peer-interface	318
preemption	319
priority	320
refresh-time	321
reliable	321
rsvp	322
soft-preemption	322
subscription	323
traceoptions	324
tunnel-services	326
update-threshold	326

Part 4

LDP

Chapter 14

LDP Overview	329
LDP Introduction	329
LDP Standards	330
JUNOS Software LDP Protocol Implementation	330
LDP Operation	331
Tunneling LDP LSPs in RSVP LSPs	331
Tunneling LDP LSPs in RSVP LSPs Overview	331
Label Operations	332
LDP Message Types	333
Discovery Messages	333
Session Messages	334
Advertisement Messages	334
Notification Messages	334
LDP Graceful Restart	334

Chapter 15

LDP Configuration Guidelines	337
Minimum LDP Configuration	338
Enabling and Disabling LDP	338
Configuring the LDP Timer for Hello Messages	338
Configuring the LDP Timer for Link Hello Messages	339
Configuring the LDP Timer for Targeted Hello Messages	339
Configuring the Delay Before LDP Neighbors Are Considered Down	339
Configuring the LDP Hold Time for Link Hello Messages	340
Configuring the LDP Hold Time for Targeted Hello Messages	340
Enabling Strict Targeted Hello Messages for LDP	340
Configuring the Interval for LDP Keepalive Messages	341
Configuring the LDP Keepalive Timeout	341
Configuring LDP Route Preferences	341

Configuring LDP Graceful Restart	342
Enabling Graceful Restart	342
Disabling LDP Graceful Restart or Helper Mode	342
Configuring Recovery Time and Maximum Recovery Time	343
Filtering Inbound LDP Label Bindings	343
Examples: Filtering Inbound LDP Label Bindings	345
Filtering Outbound LDP Label Bindings	345
Examples: Filtering Outbound LDP Label Bindings	347
Specifying the Transport Address Used by LDP	347
Configuring the Prefixes Advertised into LDP from the Routing Table	348
Example: Configuring the Prefixes Advertised into LDP	348
Configuring FEC Deaggregation	349
Configuring Policers for LDP FECs	349
Configuring LDP IPv4 FEC Filtering	350
Configuring BFD for LDP LSPs	351
Configuring ECMP-Aware BFD for RSVP LSPs	353
Configuring a Failure Action for the BFD Session on an LDP LSP	353
Configuring the Holddown Interval for the BFD Session	354
Configuring OAM Ingress Policies for LDP	354
Configuring LDP LSP Traceroute	354
Collecting LDP Statistics	355
LDP Statistics Output	356
Disabling LDP Statistics on the Penultimate-Hop Router	357
LDP Statistics Limitations	357
Tracing LDP Protocol Traffic	358
Tracing LDP Protocol Traffic at the Protocol and Routing Instance Levels	358
Tracing LDP Protocol Traffic Within FECs	359
Examples: Tracing LDP Protocol Traffic	359
Configuring Miscellaneous LDP Properties	360
Configuring LDP to Use the IGP Route Metric	361
Preventing Addition of Ingress Routes to the inet.0 Routing Table	361
Multiple-Instance LDP and Carrier-of-Carriers VPNs	361
Configuring MPLS and LDP to Pop the Label on the Ultimate-Hop Router	362
Enabling LDP over RSVP-Established LSPs	362
Enabling LDP over RSVP-Established LSPs in Heterogeneous Networks	363
Configuring the TCP MD5 Signature for LDP Sessions	363
Disabling SNMP Traps for LDP	364
Configuring LDP Synchronization with the IGP on LDP Links	364
Configuring LDP Synchronization with the IGP on the Router	365
Configuring the Label Withdrawal Timer	365
Ignoring the LDP Subnet Check	366

Chapter 16

Summary of LDP Configuration Statements

367

allow-subnet-mismatch	367
authentication-key	368
bfd-liveness-detection	369

deaggregate	370
disable	371
ecmp	372
egress-policy	372
explicit-null	373
export	373
failure-action	374
graceful-restart	375
hello-interval	376
helper-disable	377
holddown-interval	377
hold-time	378
ignore-lsp-metrics	379
igp-synchronization	379
import	380
ingress-policy	380
interface	381
keepalive-interval	382
keepalive-timeout	382
l2-smart-policy	383
label-withdrawal-delay	383
ldp	384
ldp-synchronization	384
log-updown	385
maximum-neighbor-recovery-time	385
no-deaggregate	386
no-forwarding	386
oam	387
periodic-traceroute	389
policing	390
preference	391
recovery-time	391
session	392
strict-targeted-hellos	392
targeted-hello	393
traceoptions	394
track-igp-metric	396
traffic-statistics	397
transport-address	398

Part 5

CCC and TCC

Chapter 17

CCC and TCC Overview

401

CCC Overview	401
Transmitting Nonstandard BPDUs	402
TCC Overview	402
CCC and TCC Graceful Restart	403

Chapter 18**CCC and TCC Configuration Guidelines****405**

Configuring Layer 2 Switching Cross-Connects Using CCC	405
Configuring the CCC Encapsulation for Layer 2 Switching	
Cross-Connects	406
Configuring ATM Encapsulation for Layer 2 Switching	
Cross-Connects	406
Configuring Ethernet Encapsulation for Layer 2 Switching	
Cross-Connects	407
Configuring Ethernet VLAN Encapsulation for Layer 2 Switching	
Cross-Connects	407
Configuring Aggregated Ethernet Encapsulation for Layer 2 Switching	
Cross-Connects	408
Configuring Frame Relay Encapsulation for Layer 2 Switching	
Cross-Connects	409
Configuring PPP and Cisco HDLC Encapsulation for Layer 2 Switching	
Cross-Connects	410
Configuring the CCC Connection for Layer 2 Switching	
Cross-Connects	410
Configuring MPLS for Layer 2 Switching Cross-Connects	411
Example: Configuring a Layer 2 Switching Cross-Connect	411
Configuring MPLS LSP Tunnel Cross-Connects Using CCC	413
Configuring the CCC Encapsulation for LSP Tunnel Cross-Connects	414
Configuring the CCC Connection for LSP Tunnel Cross-Connects	415
Example: Configuring an LSP Tunnel Cross-Connect	416
Configuring LSP Stitching Cross-Connects Using CCC	417
Example: Configuring an LSP Stitching Cross-Connect	418
Configuring TCC	418
Configuring the Encapsulation for Layer 2 Switching TCCs	419
Configuring PPP and Cisco HDLC Encapsulation for Layer 2 Switching	
TCCs	419
Configuring ATM Encapsulation for Layer 2 Switching TCCs	420
Configuring Frame Relay Encapsulation for Layer 2 Switching	
TCCs	420
Configuring Ethernet Encapsulation for Layer 2 Switching TCCs	420
Configuring Ethernet Extended VLAN Encapsulation for Layer 2	
Switching TCCs	421
Configuring ARP for Ethernet and Ethernet Extended VLAN	
Encapsulations	422
Configuring the Connection for Layer 2 Switching TCCs	422
Configuring MPLS for Layer 2 Switching TCCs	423
Configuring CCC and TCC Graceful Restart	424
Configuring CCC Switching for Point-to-Multipoint LSPs	424
Configuring the Point-to-Multipoint LSP Switch on Ingress PE	
Routers	425
Configuring the Point-to-Multipoint LSP Switch on Egress PE	
Routers	425

Chapter 19	Summary of CCC and TCC Configuration Statements	427
	connections	428
	encapsulation	429
	encapsulation (Logical Interface)	430
	encapsulation (Physical Interface)	433
	interface-switch	435
	lsp-switch	436
	p2mp-receive-switch	437
	p2mp-transmit-switch	438
	remote-interface-switch	439
 Part 6	 GMPLS	
 Chapter 20	 GMPLS Overview	 443
	GMPLS Standards	443
	Overview	445
	GMPLS Operation	446
	GMPLS and OSPF	447
	GMPLS and CSPF	447
	GMPLS Features	447
 Chapter 21	 GMPLS Configuration Guidelines	 449
	LMP Configuration Overview	449
	Configuring LMP Traffic Engineering Links	450
	Configuring the Local IP Address for Traffic Engineering Links	451
	Configuring the Remote IP Address for Traffic Engineering Links	451
	Configuring the Remote ID for Traffic Engineering Links	452
	Configuring LMP Peers	452
	Configuring the ID for LMP Peers	453
	Configuring the Interface for Control Channels Between LMP Peers	453
	Configuring the LMP Control Channel Interface for the Peer	454
	Configuring the Remote IP Address for LMP Control Channels	454
	Configuring Hello Message Intervals for LMP Control Channels	455
	Controlling Message Exchange for LMP Control Channels	456
	Preventing the Local Peer from Initiating LMP Negotiation	456
	Associating Traffic Engineering Links with LMP Peers	456
	Disabling the Traffic Engineering Link for LMP Peers	457
	Configuring RSVP and OSPF for LMP Peer Interfaces	457
	Configuring RSVP Signaling for LMP Peer Interfaces	457
	Configuring OSPF Routing for LMP Peer Interfaces	458
	Configuring the Hello Interval for LMP Peer Interfaces	458
	Configuring MPLS Paths for GMPLS	458
	Tracing LMP Traffic	459

Configuring MPLS LSPs for GMPLS	460
Configuring the Encoding Type	460
Configuring the GPID	461
Configuring the Signal Bandwidth Type	461
Configuring GMPLS Bidirectional LSPs	461
Allowing Non-Packet GMPLS LSPs to Establish Paths Through Routers	
Running the JUNOS Software	462
Gracefully Tearing Down GMPLS LSPs	462
Temporarily Deleting GMPLS LSPs	462
Permanently Deleting GMPLS LSPs	463
Configuring the Graceful Deletion Timeout Interval	463

Chapter 22**Hierarchy of RSVP LSPs Configuration Guidelines 465**

Hierarchy of RSVP LSPs Standard	465
Hierarchy of RSVP LSPs Overview	465
Hierarchy of RSVP LSPs	466
Advertising the Forwarding Adjacency with OSPF	466
Configuring a Hierarchy of RSVP LSPs	466
Configuring an RSVP LSP on Ingress Routers	467
Configuring Forwarding Adjacencies	467
Configuring the Local IP Address for Forwarding Adjacencies	467
Configuring the Remote IP Address for Forwarding Adjacencies	468
Configuring the LSP for Forwarding Adjacencies	468
Configuring RSVP for Forwarding Adjacencies	468
Advertising Forwarding Adjacencies Using OSPF	469

Chapter 23**Summary of GMPLS Configuration Statements 471**

address	471
admin-down	471
control-channel	472
dead-interval	472
disable	473
disable (GMPLS)	473
disable (OSPF Peer Interface)	473
hello-dead-interval	474
hello-interval	475
hello-interval (LMP)	475
hello-interval (OSPF)	475
interface	476
label-switched-path	476
link-management	477
lmp-control-channel	477
lmp-protocol	478
local-address	478
passive	479
peer	479

peer-interface	480
peer-interface (OSPF)	480
peer-interface (for RSVP)	480
remote-address	481
remote-address (for LMP Control Channel)	481
remote-address (for LMP Traffic Engineering)	481
remote-id	482
retransmission-interval	482
retransmit-interval	483
retry-limit	483
te-link	484
traceoptions	485
transit-delay	487

Part 7

Indexes

Index	491
Index of Statements and Commands	507

List of Figures

Part 2

MPLS

Chapter 3	MPLS Overview	21
	Figure 1: Label Encoding	27
	Figure 2: Class-of-Service Bits	27
	Figure 3: CSPF Computation Process	31
	Figure 4: Aggregation Router A Dual-Homed on Core Routers B and C	33
	Figure 5: Typical SPF Tree, Sourced from Router A	34
	Figure 6: Modified SPF Tree, Using LSP A–D as a Shortcut	34
	Figure 7: Modified SPF Tree, Using LSP A–D and LSP A–E as Shortcuts	35
	Figure 8: IGP Shortcuts	36
	Figure 9: IGP Shortcuts in a Bigger Network	36
	Figure 10: SPF Computations with Advertised LSPs	38
	Figure 11: MPLS Application Topology	40
	Figure 12: How BGP Determines How to Reach Next-Hop Addresses	40
	Figure 13: Routing and Forwarding Tables, traffic-engineering bgp	42
	Figure 14: Routing and Forwarding Tables, traffic-engineering bgp-igp	43
	Figure 15: Detours Established for an LSP Using Fast Reroute	45
	Figure 16: Detour After the Link from Router B to Router C Fails	45
	Figure 17: Detours Merging into Other Detours	46
	Figure 18: Point-to-Multipoint LSPs	49
Chapter 4	MPLS Router Configuration Guidelines	51
	Figure 19: IPv6 Networks Linked by MPLS IPv4 Tunnels	64
Chapter 7	Static and Explicit-Path LSP Configuration Guidelines	133
	Figure 20: Static MPLS Configuration	135

Part 3

RSVP

Chapter 11	RSVP Overview	255
	Figure 21: Link Protection Creating a Bypass LSP for the Protected Interface	266
	Figure 22: Node Protection Creating a Next-Next-Hop Bypass LSP	268

Part 4

LDP

Chapter 14	LDP Overview	329
	Figure 23: Swap and Push When LDP LSPs Are Tunneled Through RSVP LSPs	332
	Figure 24: Double Push When LDP LSPs Are Tunneled Through RSVP LSPs	333

Part 5**CCC and TCC**

Chapter 17	CCC and TCC Overview	401
	Figure 25: TCC Example	402
	Figure 26: Remote Interface Switch Connecting Two CE Routers Using CCC	403
Chapter 18	CCC and TCC Configuration Guidelines	405
	Figure 27: Layer 2 Switching Cross-Connect	405
	Figure 28: Topology of a Frame Relay Layer 2 Switching Cross-Connect	411
	Figure 29: Sample Topology of a VLAN Layer 2 Switching Cross-Connect	412
	Figure 30: MPLS LSP Tunnel Cross-Connect	413
	Figure 31: Example Topology of MPLS LSP Tunnel Cross-Connect	416
	Figure 32: LSP Stitching Cross-Connect	417
	Figure 33: Example Topology of LSP Stitching Cross-Connect	418

List of Tables

	About This Guide	xxxi
	Table 1: Notice Icons	xxxv
	Table 2: Text and Syntax Conventions	xxxv
Part 2	MPLS	
Chapter 5	MPLS-Signaled LSP Configuration Guidelines	69
	Table 3: MPLS CoS Values	100
Chapter 6	DiffServ-Aware Traffic Engineering Configuration Guidelines	109
	Table 4: Default Values for the Traffic Engineering Class Matrix	117
Part 3	RSVP	
Chapter 11	RSVP Overview	255
	Table 5: One-to-One Backup Compared with Facility Backup	267
Chapter 12	RSVP Configuration Guidelines	273
	Table 6: RSVP Refresh Reduction Behavior	276
Part 4	LDP	
Chapter 15	LDP Configuration Guidelines	337
	Table 7: from Operators That Apply to LDP Received-Label Filtering	344
	Table 8: to Operators for LDP Outbound-Label Filtering	346

About This Guide

This preface provides the following guidelines for using the *JUNOS® Software MPLS Applications Configuration Guide*:

- JUNOS Documentation and Release Notes on page xxxi
- Objectives on page xxxii
- Audience on page xxxii
- Supported Platforms on page xxxii
- Using the Indexes on page xxxiii
- Using the Examples in This Manual on page xxxiii
- Documentation Conventions on page xxxiv
- Documentation Feedback on page xxxvi
- Requesting Technical Support on page xxxvii

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 an overview of the MPLS applications functions of the JUNOS software and describes how to configure MPLS applications on the router.



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/>.

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 complete index that includes topic entries, and an index of statements and commands only.

In the index of statements and commands, an entry refers to a statement summary section only. In the complete index, the entry for a configuration statement or command contains at least two parts:

- The primary entry refers to the statement summary section.
- The secondary entry, *usage guidelines*, refers to the section in a configuration guidelines chapter that describes how to use the statement or command.

Using the Examples in This Manual

If you want to use the examples in this manual, you can use the **load merge** or the **load merge relative** command. These commands cause the software to merge the incoming configuration into the current candidate configuration. If the example configuration contains the top level of the hierarchy (or multiple hierarchies), the example is a *full example*. In this case, use the **load merge** command.

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

```
system {
  scripts {
    commit {
      file ex-script.xsl;
    }
  }
}
interfaces {
  fxp0 {
```

```

        disable;
        unit 0 {
            family inet {
                address 10.0.0.1/24;
            }
        }
    }
}

```

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

```

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

```

Merging a Snippet

To merge a snippet, follow these steps:

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

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

```

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

```

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

```

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

```

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

```

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

```

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

Documentation Conventions

Table 1 on page xxxv 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 xxxv 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 <code>configure</code> 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> >;

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

Convention	Description	Examples
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	broadcast multicast (string1 string2 string3)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only
[] (square brackets)	Enclose a variable for which you can substitute one or more values.	community name members [community-ids]
Indentation and braces ({ })	Identify a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop 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

Overview

- Traffic Engineering Overview on page 3
- Complete MPLS Applications Configuration Statements on page 9

Chapter 1

Traffic Engineering Overview

This chapter discusses the following topics:

- Traffic Engineering Capabilities on page 3
- Components of Traffic Engineering on page 4
- Packet Forwarding Component on page 4
- Packet Forwarding Based on Label Swapping on page 4
- How a Packet Traverses an MPLS Backbone on page 5
- Information Distribution Component on page 5
- Path Selection Component on page 6
- Offline Planning and Analysis on page 6
- Signaling Component on page 7
- Flexible LSP Calculation and Configuration on page 7

Traffic Engineering Capabilities

The task of mapping traffic flows onto an existing physical topology is called *traffic engineering*. Traffic engineering provides the ability to move traffic flow away from the shortest path selected by the interior gateway protocol (IGP) and onto a potentially less congested physical path across a network.

Traffic engineering provides the capabilities to do the following:

- Route primary paths around known bottlenecks or points of congestion in the network.
- Provide precise control over how traffic is rerouted when the primary path is faced with single or multiple failures.
- Provide more efficient use of available aggregate bandwidth and long-haul fiber by ensuring that subsets of the network do not become overutilized while other subsets of the network along potential alternate paths are underutilized.
- Maximize operational efficiency.

- Enhance the traffic-oriented performance characteristics of the network by minimizing packet loss, minimizing prolonged periods of congestion, and maximizing throughput.
- Enhance statistically bound performance characteristics of the network (such as loss ratio, delay variation, and transfer delay) required to support a multiservices Internet.

Components of Traffic Engineering

In the JUNOS Software, traffic engineering is implemented with MPLS and RSVP. Traffic engineering is composed of four functional components:

- Packet Forwarding Component on page 4
- Information Distribution Component on page 5
- Path Selection Component on page 6
- Signaling Component on page 7

Packet Forwarding Component

The packet forwarding component of the JUNOS traffic engineering architecture is MPLS, which is responsible for directing a flow of IP packets along a predetermined path across a network. This path is called a *label-switched path (LSP)*. LSPs are simplex; that is, the traffic flows in one direction from the head-end (ingress) router to a tail-end (egress) router. Duplex traffic requires two LSPs: one LSP to carry traffic in each direction. An LSP is created by the concatenation of one or more label-switched hops, allowing a packet to be forwarded from one router to another across the MPLS domain.

When an ingress router receives an IP packet, it adds an MPLS header to the packet and forwards it to the next router in the LSP. The labeled packet is forwarded along the LSP by each router until it reaches the tail end of the LSP, the egress router. At this point the MPLS header is removed, and the packet is forwarded based on Layer 3 information such as the IP destination address. The value of this scheme is that the physical path of the LSP is not limited to what the IGP would choose as the shortest path to reach the destination IP address.

This section discusses the following topics:

- Packet Forwarding Based on Label Swapping on page 4
- How a Packet Traverses an MPLS Backbone on page 5

Packet Forwarding Based on Label Swapping

The packet forwarding process at each router is based on the concept of label swapping. This concept is similar to what occurs at each Asynchronous Transfer Mode (ATM) switch in a permanent virtual circuit (PVC). Each MPLS packet carries a 4-byte encapsulation header that contains a 20-bit, fixed-length label field. When a packet containing a label arrives at a router, the router examines the label and

copies it as an index to its MPLS forwarding table. Each entry in the forwarding table contains an interface-inbound label pair mapped to a set of forwarding information that is applied to all packets arriving on the specific interface with the same inbound label.

How a Packet Traverses an MPLS Backbone

This section describes how an IP packet is processed as it traverses an MPLS backbone network.

At the entry edge of the MPLS backbone, the IP header is examined by the ingress router. Based on this analysis, the packet is classified, assigned a label, encapsulated in an MPLS header, and forwarded toward the next hop in the LSP. MPLS provides a high degree of flexibility in the way that an IP packet can be assigned to an LSP. For example, in the JUNOS traffic engineering implementation, all packets arriving at the ingress router that are destined to exit the MPLS domain at the same egress router are forwarded along the same LSP.

Once the packet begins to traverse the LSP, each router uses the label to make the forwarding decision. The MPLS forwarding decision is made independently of the original IP header: the incoming interface and label are used as lookup keys into the MPLS forwarding table. The old label is replaced with a new label, and the packet is forwarded to the next hop along the LSP. This process is repeated at each router in the LSP until the packet reaches the egress router.

When the packet arrives at the egress router, the label is removed and the packet exits the MPLS domain. The packet is then forwarded based on the destination IP address contained in the packet's original IP header according to the traditional shortest path calculated by the IP routing protocol.

Information Distribution Component

Traffic engineering requires detailed knowledge about the network topology as well as dynamic information about network loading. To implement the information distribution component, simple extensions to the IGPs are defined. Link attributes are included as part of each router's link-state advertisement. IS-IS extensions include the definition of new type length values (TLVs), whereas OSPF extensions are implemented with opaque link-state advertisements (LSAs). The standard flooding algorithm used by the link-state IGPs ensures that link attributes are distributed to all routers in the routing domain. Some of the traffic engineering extensions to be added to the IGP link-state advertisement include maximum link bandwidth, maximum reserved link bandwidth, current bandwidth reservation, and link coloring.

Each router maintains network link attributes and topology information in a specialized traffic engineering database. The traffic engineering database is used exclusively for calculating explicit paths for the placement of LSPs across the physical topology. A separate database is maintained so that the subsequent traffic engineering computation is independent of the IGP and the IGP's link-state database. Meanwhile, the IGP continues its operation without modification, performing the traditional shortest-path calculation based on information contained in the router's link-state database.

Path Selection Component

After network link attributes and topology information are flooded by the IGP and placed in the traffic engineering database, each ingress router uses the traffic engineering database to calculate the paths for its own set of LSPs across the routing domain. The path for each LSP can be represented by either a strict or loose explicit route. An explicit route is a preconfigured sequence of routers that should be part of the physical path of the LSP. If the ingress router specifies all the routers in the LSP, the LSP is said to be identified by a strict explicit route. If the ingress router specifies only some of the routers in the LSP, the LSP is described as a loose explicit route. Support for strict and loose explicit routes allows the path selection process to be given broad latitude whenever possible, but to be constrained when necessary.

The ingress router determines the physical path for each LSP by applying a Constrained Shortest Path First (CSPF) algorithm to the information in the traffic engineering database. CSPF is a shortest-path-first algorithm that has been modified to take into account specific restrictions when the shortest path across the network is calculated. Input into the CSPF algorithm includes:

- Topology link-state information learned from the IGP and maintained in the traffic engineering database
- Attributes associated with the state of network resources (such as total link bandwidth, reserved link bandwidth, available link bandwidth, and link color) that are carried by IGP extensions and stored in the traffic engineering database
- Administrative attributes required to support traffic traversing the proposed LSP (such as bandwidth requirements, maximum hop count, and administrative policy requirements) that are obtained from user configuration

As CSPF considers each candidate node and link for a new LSP, it either accepts or rejects a specific path component based on resource availability or whether selecting the component violates user policy constraints. The output of the CSPF calculation is an explicit route consisting of a sequence of router addresses that provides the shortest path through the network that meets the constraints. This explicit route is then passed to the signaling component, which establishes the forwarding state in the routers along the LSP.

Offline Planning and Analysis

Despite the reduced management effort resulting from online path calculation, an offline planning and analysis tool is still required to optimize traffic engineering globally. Online calculation takes resource constraints into account and calculates one LSP at a time. The challenge with this approach is that it is not deterministic. The order in which LSPs are calculated plays a critical role in determining each LSP's physical path across the network. LSPs that are calculated early in the process have more resources available to them than LSPs calculated later in the process because previously calculated LSPs consume network resources. If the order in which the LSPs are calculated is changed, the resulting set of physical paths for the LSPs also can change.

An offline planning and analysis tool simultaneously examines each link's resource constraints and the requirements of each LSP. Although the offline approach can take several hours to complete, it performs global calculations, compares the results of each calculation, and then selects the best solution for the network as a whole. The output of the offline calculation is a set of LSPs that optimizes utilization of network resources. After the offline calculation is completed, the LSPs can be established in any order because each is installed according to the rules for the globally optimized solution.

Signaling Component

An LSP is not known to be workable until it is actually established by the signaling component. The signaling component, which is responsible for establishing LSP state and distributing labels, relies on a number of extensions to RSVP:

- The Explicit Route object allows an RSVP path message to traverse an explicit sequence of routers that is independent of conventional shortest-path IP routing. The explicit route can be either strict or loose.
- The Label Request object permits the RSVP path message to request that intermediate routers provide a label binding for the LSP that it is establishing.
- The Label object allows RSVP to support the distribution of labels without changing its existing mechanisms. Because the RSVP Resv message follows the reverse path of the RSVP path message, the Label object supports the distribution of labels from downstream nodes to upstream nodes.

Flexible LSP Calculation and Configuration

Traffic engineering involves mapping traffic flow onto a physical topology. You can determine the paths online using constraint-based routing. Regardless of how the physical path is calculated, the forwarding state is installed across the network through RSVP.

The JUNOS Software supports the following ways to route and configure an LSP:

- You can calculate the full path for the LSP offline and individually configure each router in the LSP with the necessary static forwarding state. This is analogous to the way some Internet service providers (ISPs) configure their IP-over-ATM cores.
- You can calculate the full path for the LSP offline and statically configure the ingress router with the full path. The ingress router then uses RSVP as a dynamic signaling protocol to install a forwarding state in each router along the LSP.
- You can rely on constraint-based routing to perform dynamic online LSP calculation. You configure the constraints for each LSP; then the network itself determines the path that best meets those constraints. Specifically, the ingress router calculates the entire LSP based on the constraints and then initiates signaling across the network.
- You can calculate a partial path for an LSP offline and statically configure the ingress router with a subset of the routers in the path; then you can permit online calculation to determine the complete path.

For example, consider a topology that includes two east-west paths across the United States: one in the north through Chicago and one in the south through Dallas. If you want to establish an LSP between a router in New York and one in San Francisco, you can configure the partial path for the LSP to include a single loose-routed hop of a router in Dallas. The result is an LSP routed along the southern path. The ingress router uses CSPF to compute the complete path and RSVP to install the forwarding state along the LSP.

- You can configure the ingress router with no constraints whatsoever. In this case, normal IGP shortest-path routing is used to determine the path of the LSP. This configuration does not provide any value in terms of traffic engineering. However, it is easy and might be useful in situations when services such as virtual private networks (VPNs) are needed.

In all these cases, you can specify any number of LSPs as backups for the primary LSP, thus allowing you to combine more than one configuration approach. For example, you might explicitly compute the primary path offline, set the secondary path to be constraint-based, and have the tertiary path be unconstrained. If a circuit on which the primary LSP is routed fails, the ingress router notices the outage from error notifications received from a downstream router or by the expiration of RSVP soft-state information. Then the router dynamically forwards traffic to a hot-standby LSP or calls on RSVP to create a forwarding state for a new backup LSP.

Chapter 2

Complete MPLS Applications Configuration Statements

This chapter is organized as follows:

- [edit logical-systems] Hierarchy Level on page 9
- [edit protocols connections] Hierarchy Level on page 10
- [edit protocols ldp] Hierarchy Level on page 10
- [edit protocols link-management] Hierarchy Level on page 12
- [edit protocols mpls] Hierarchy Level on page 13
- [edit protocols rsvp] Hierarchy Level on page 17

[edit logical-systems] Hierarchy Level

The following MPLS protocol statements can be configured at the [edit logical-systems] hierarchy level. This is not a comprehensive list of statements available for logical systems. Only the statements that are also documented in this manual are listed here. For more information about logical systems, see the *JUNOS Routing Protocols Configuration Guide*.



NOTE: Beginning with JUNOS Software Release 9.3, the logical router feature has been renamed logical system.

All configuration statements, operational commands, **show** command outputs, error messages, log messages, and SNMP MIB objects that contain the string logical-router or logical-routers have been changed to logical-system and logical-systems, respectively.

```
logical-systems {  
  logical-system-name {  
    protocols {  
      connections {  
        connections-configuration;  
      }  
      ldp {  
        ldp-configuration;  
      }  
    }  
  }  
}
```

```

        link-management {
            link-management-configuration;
        }
        mpls {
            mpls-configuration;
        }
        rsvp {
            rsvp-configuration;
        }
    }
}

```

[edit protocols connections] Hierarchy Level

The following statements can also be configured at the [edit logical-systems *logical-system-name*] hierarchy level:

```

protocols {
  connections {
    interface-switch connection-name {
      interface first-interface-name.unit-number;
      interface second-interface-name.unit-number;
    }
    lsp-switch connection-name {
      transmit-lsp label-switched-path;
      receive-lsp label-switched-path;
    }
    p2mp-receive-switch {
      output-interface interface-name.unit-number;
      receive-p2mp-lsp receiving-point-to-multipoint-lsp;
    }
    p2mp-transmit-switch {
      input-interface input-interface-name.unit-number;
      transmit-p2mp-lsp transmitting-point-to-multipoint-lsp;
    }
    remote-interface-switch connection-name {
      interface interface-name.unit-number;
      transmit-lsp label-switched-path;
      receive-lsp label-switched-path;
    }
  }
}

```

[edit protocols ldp] Hierarchy Level

The following statements can also be configured at the [edit logical-systems *logical-system-name*] hierarchy level:

```

protocols {
  ldp {
    (deaggregate | no-deaggregate);
    egress-policy [ policy-names ];
    explicit-null;
  }
}

```



```

export [ policy-names ];
graceful-restart {
    disable;
    helper-disable;
    maximum-neighbor-recovery-time seconds;
    recovery-time seconds;
}
import [ policy-names];
interface (interface-name | all) {
    disable;
    hello-interval seconds;
    hold-time seconds;
    transport-address (interface | router-id);
}
keepalive-interval seconds;
keepalive-timeout seconds;
log-updown {
    trap disable;
}
no-forwarding;
oam {
    bfd-liveness-detection {
        detection-time threshold milliseconds;
        ecmp;
        failure-action {
            remove-nexthop;
            remove-route;
        }
        holddown-interval milliseconds;
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        minimum-transmit-interval milliseconds;
        multiplier detection-time-multiplier;
        no-adaptation;
        transmit-interval {
            minimum-interval milliseconds;
            threshold milliseconds;
        }
    }
}
fec fec-address;
ingress-policy ingress-policy-name;
periodic-traceroute {
    disable;
    exp exp-value;
    fanout fanout-value;
    frequency minutes;
    paths number-of-paths;
    retries retry-attempts;
    source address;
    ttl tvl-value;
    wait seconds;
}
}
policing {
    fec fec-address {
        ingress-traffic filter-name;
    }
}

```

```

        transit-traffic filter-name;
    }
}
preference preference;
session address {
    authentication-key md5-authentication-key;
}
strict-targeted-hellos;
traceoptions {
    file filename <files number <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
track-igp-metric;
traffic-statistics {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    interval interval;
    no-penultimate-hop;
}
transport-address (address | interface | router-id);
}
}

```

[edit protocols link-management] Hierarchy Level

The following statements can also be configured at the [edit logical-systems *logical-system-name*] hierarchy level:

```

protocols {
    link-management {
        peer peer-name {
            address address;
            control-channel [ control-channel-interfaces ];
            te-link [te-link-names];
        }
        te-link te-link-name {
            disable;
            interface interface-name {
                disable;
                local-address address;
                remote-address address;
                remote-id id-number;
            }
            label-switched-path label-switched-path-name;
            local-address address;
            remote-address address;
            remote-id id-number;
        }
        traceoptions {
            file filename <files number> <size size> <world-readable | no-world-readable>;
            flag flag <flag-modifier> <disable>;
        }
    }
}

```

[edit protocols mpls] Hierarchy Level

The following statements can also be configured at the [edit logical-systems *logical-system-name*] hierarchy level:

```

protocols {
  mpls {
    disable;
    admin-group {
      exclude [ group-names];
      include-all [ group-names];
      include-any [ group-names];
    }
    admin-groups {
      group-name group-value;
    }
    advertisement-hold-time seconds;
    auto-policing {
      class all (drop | loss-priority-high | loss-priority-low);
      class ctnumber (drop | loss-priority-high | loss-priority-low);
    }
    bandwidth bps {
      ct0 bps;
      ct1 bps;
      ct2 bps;
      ct3 bps;
    }
    class-of-service cos-value;
    diffserv-te {
      bandwidth-model {
        extended-mam;
        mam;
        rdm;
      }
      te-class-matrix {
        tenclass {
          priority priority;
          traffic-class ctnumber priority priority;
        }
      }
    }
    explicit-null;
    hop-limit number;
    icmp-tunneling;
    interface (interface-name | all) {
      disable;
      admin-group [group-names];
      label-map (default-route | in-label-name) {
        class-of-service cos-value;
        next-hop (address | interface-name | address/interface-name) |
          (discard | reject);
        (pop | swap out-label);
        preference preference;
        swap-push swap-label push-label;
      }
    }
  }
}

```

```

    }
  }
  ipv6-tunneling;
  label-switched-path lsp-name {
    disable;
    adaptive;
    admin-down;
    admin-group {
      exclude [ group-names ];
      include-all;
      include-any [ group-names ];
    }
    auto-bandwidth {
      adjust-interval seconds;
      adjust-threshold percent;
      maximum-bandwidth bps;
      minimum-bandwidth bps;
      monitor-bandwidth;
    }
    bandwidth bps {
      ct0 bps;
      ct1 bps;
      ct2 bps;
      ct3 bps;
    }
    class-of-service cos-value;
    description text;
    fast-reroute {
      (bandwidth bps | bandwidth-percent percent);
      (exclude [ group-names ] | no-exclude);
      hop-limit number;
      (include-all [ group-names ] | no-include-all);
      (include-any [ group-names ] | no-include-any);
    }
    from address;
    hop-limit number;
    install {
      destination-prefix/prefix-length <active>;
    }
    ldp-tunneling;
    link-protection;
    lsp-attributes {
      encoding-type (ethernet | packet | pdh | sonet-sdh);
      gp-id (ethernet | hdlc | ipv4 | ppp);
      signal-bandwidth type;
      switching-type (fiber | lambda | psc-1 | tdm);
    }
    metric number;
    no-cspf;
    no-decrement-ttl;
    node-link-protection;
    optimize-timer seconds;
    p2mp path-name;
    policing {
      filter filter-name;
      no-automatic-policing;
    }
  }
}

```

```

}
preference preference;
primary path-name {
    adaptive;
    admin-group {
        exclude [ group-names ];
        include-all [ group-names ];
        include-any [ group-names ];
    }
    bandwidth bps {
        ct0 bps;
        ct1 bps;
        ct2 bps;
        ct3 bps;
    }
    class-of-service cos-value;
    hop-limit number;
    no-cspf;
    no-decrement-ttl;
    optimize-timer seconds;
    preference preference;
    priority setup-priority reservation-priority;
    (record | no-record);
    select (manual | unconditional);
}
standby;
}
priority setup-priority reservation-priority;
(random | least-fill | most-fill);
(record | no-record);
retry-limit number;
retry-timer seconds;
revert-timer seconds;
secondary path-name {
    adaptive;
    admin-group {
        exclude [ group-names ];
        include-all [ group-names ];
        include-any [ group-names ];
    }
    bandwidth bps {
        ct0 bps;
        ct1 bps;
        ct2 bps;
        ct3 bps;
    }
}
class-of-service cos-value;
hop-limit number;
no-cspf;
no-decrement-ttl;
optimize-timer seconds;
preference preference;
priority setup-priority reservation-priority;
(record | no-record);
select (manual | unconditional);
standby;

```

```

    }
    soft-preemption;
    standby;
    to address;
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
}
log-updown {
    no-trap {
        mpls-lsp-traps;
        rfc3812-traps;
    }
    (syslog | no-syslog);
    trap;
    trap-path-down;
    trap-path-up;
}
no-cspf;
no-decrement-ttl;
no-propagate-ttl;
optimize-aggressive;
optimize-timer seconds;
path path-name {
    (address | hostname) <strict | loose>;
}
path-mtu {
    allow-fragmentation;
    rsvp {
        mtu-signaling;
    }
}
preference preference;
priority setup-priority reservation-priority;
(record | no-record );
revert-timer seconds;
rsvp-error-hold-time seconds;
smart-optimize-timer seconds;
standby;
static-path inet {
    prefix {
        class-of-service cos-value;
        double-push bottom-label top-label;
        next-hop (address | interface-name | address/interface-name);
        preference preference;
        push out-label;
        triple-push bottom-label middle-label top-label;
    }
}
statistics {
    auto-bandwidth;
    file filename <files number> <size size> <world-readable | no-world-readable>;
    interval seconds;
}
traceoptions {

```

```

    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag;
  }
  traffic-engineering (bgp | bgp-igp | bgp-igp-both-ribs | mpls-forwarding);
}

```

[edit protocols rsvp] Hierarchy Level

The following statements can also be configured at the [edit logical-systems *logical-system-name*] hierarchy level:

```

protocols {
  rsvp {
    disable;
    fast-reroute optimize-timer seconds;
    graceful-deletion-timeout seconds;
    graceful-restart {
      disable;
      helper-disable;
      maximum-helper-recovery-time seconds;
      maximum-helper-restart-time seconds;
    }
    interface interface-name {
      disable;
      (aggregate | no-aggregate);
      authentication-key key;
      bandwidth bps;
      hello-interval seconds;
      link-protection {
        disable;
        admin-group {
          exclude group-names;
          include-all group-names;
          include-any group-names;
        }
        bandwidth bandwidth;
        bypass bypass-name {
          bandwidth bps {
            ct0 bps;
            ct1 bps;
            ct2 bps;
            ct3 bps;
          }
          hop-limit number;
          no-cspf;
          path address <strict | loose>;
          priority setup-priority reservation-priority;
          to address;
        }
        class-of-service cos-value;
        hop-limit number;
        max-bypasses number;
        no-cspf;
        no-node-protection;
      }
    }
  }
}

```

```

        optimize-timer seconds;
        path address <strict | loose>;
        priority setup-priority reservation-priority;
        subscription percentage {
            ct0 percentage;
            ct1 percentage;
            ct2 percentage;
            ct3 percentage;
        }
    }
    (reliable | no-reliable);
    subscription percentage {
        ct0 percentage;
        ct1 percentage;
        ct2 percentage;
        ct3 percentage;
    }
    update-threshold percentage;
}
keep-multiplier number;
load-balance {
    bandwidth;
}
no-node-id-subobject;
no-p2mp-sublsp;
peer-interface peer-interface-name {
    (aggregate | no-aggregate);
    authentication-key key;
    disable;
    hello-interval seconds;
    (reliable | no-reliable);
}
preemption {
    (aggressive | disabled | normal);
    soft-preemption {
        cleanup-timer seconds;
    }
}
refresh-time seconds;
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
tunnel-services {
    devices device-names;
}
}
}

```


Part 2

MPLS

- MPLS Overview on page 21
- MPLS Router Configuration Guidelines on page 51
- MPLS-Signaled LSP Configuration Guidelines on page 69
- DiffServ-Aware Traffic Engineering Configuration Guidelines on page 109
- Static and Explicit-Path LSP Configuration Guidelines on page 133
- Point-to-Multipoint LSP Configuration Guidelines on page 141
- Miscellaneous MPLS Properties Configuration Guidelines on page 149
- Summary of MPLS Configuration Statements on page 173

Chapter 3

MPLS Overview

This chapter discusses the following topics:

- MPLS Introduction on page 22
- MPLS Standards on page 22
- Link-Layer Support on page 24
- MPLS and Traffic Engineering on page 24
- Label Description on page 25
- Special Labels on page 25
- Label Allocation on page 26
- Operations on Labels on page 27
- Routers in an LSP on page 28
- How a Packet Travels Along an LSP on page 28
- Types of LSPs on page 29
- Scope of LSPs on page 29
- Constrained-Path LSP Computation on page 29
- How CSPF Selects a Path on page 31
- Path Selection Tie-Breaking on page 32
- Computing Paths Offline on page 32
- LSPs on an Overloaded Router on page 33
- Fate Sharing on page 33
- IGP Shortcuts on page 34
- Enabling IGP Shortcuts on page 35
- LSPs Qualified in Shortcut Computations on page 36
- IGP Shortcut Applications on page 36
- IGP Shortcuts and Routing Table on page 37
- IGP Shortcuts and VPN Environments on page 37
- Advertising LSPs into IGP on page 37
- IP and MPLS Packets on Aggregated Interfaces on page 38
- MPLS Applications on page 39
- BGP Destinations on page 39

- IGP and BGP Destinations on page 41
- Selecting a Forwarding LSP Next Hop on page 41
- MPLS and Routing Tables on page 41
- MPLS and Traffic Protection on page 43
- Fast Reroute on page 44
- Fast Reroute Overview on page 44
- Detour Merging Process on page 46
- Detour Computations on page 47
- Fast Reroute Path Optimization on page 47
- Automatic Bandwidth Allocation on page 48
- Point-to-Multipoint LSPs on page 48
- MPLS Load Balancing Based on the IP Header and MPLS Labels on page 50

MPLS Introduction

MPLS provides a mechanism for engineering network traffic patterns that is independent of routing tables. MPLS assigns short labels to network packets that describe how to forward them through the network. MPLS is independent of any routing protocol and can be used for unicast packets.

In the traditional Level 3 forwarding paradigm, as a packet travels from one router to the next, an independent forwarding decision is made at each hop. The IP network layer header is analyzed, and the next hop is chosen based on this analysis and on the information in the routing table. In an MPLS environment, the analysis of the packet header is performed just once, when a packet enters the MPLS cloud. The packet is then assigned to a stream, which is identified by a *label*, which is a short (20-bit), fixed-length value at the front of the packet. Labels are used as lookup indexes for the label forwarding table. For each label, this table stores forwarding information. You can associate additional information with a label—such as class-of-service (CoS) values—that can be used to prioritize packet forwarding.

MPLS Standards

The JUNOS Software supports the following RFCs and Internet drafts related to MPLS:

- RFC 2702, *Requirements for Traffic Engineering Over MPLS*
- RFC 2858, *Multiprotocol Extensions for BGP-4*
- RFC 3063, *MPLS Loop Prevention Mechanism*
- RFC 3031, *Multiprotocol Label Switching Architecture* (provides a good overview of MPLS)
- RFC 3032, *MPLS Label Stack Encoding*
- RFC 3140, *Per Hop Behavior Identification Codes*
- RFC 3208, *PGM Reliable Transport Protocol Specification* (only the network element)

- RFC 3270, *Multi-Protocol Label Switching (MPLS) Support of Differentiated Services (E-LSPs only)*
- RFC 3443, *Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks*
- RFC 3469, *Framework for Multi-Protocol Label Switching (MPLS)-based Recovery*
- RFC 3478, *Graceful Restart Mechanism for LDP*
- RFC 3564, *Requirements for Support of Differentiated Services-aware MPLS Traffic Engineering*
- RFC 4124, *Protocol Extensions for Support of Differentiated-Service-aware MPLS Traffic Engineering*
- RFC 4125, *Maximum Allocation Bandwidth Constraints Model for Diff-Serv-aware MPLS Traffic Engineering*
- RFC 4127, *Russian Dolls Bandwidth Constraints Model for Diff-Serv-aware MPLS Traffic Engineering*
- RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)*
- RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures.*
- RFC 4950, *ICMP Extensions for Multiprotocol Label Switching*
- Internet draft draft-ietf-bfd-mpls-02.txt, *BFD for MPLS LSPs*
- Internet draft draft-ietf-mpls-rsvp-lsp-fastreroute-03.txt, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels* (except node protection in facility backup)
- Internet draft draft-ietf-mpls-label-encaps-07.txt, *MPLS Label Stack Encoding*
- Internet draft draft-ietf-mpls-p2mp-requirement-02.txt, *Requirements for Point to Multipoint Extensions to RSVP-TE* (no expiration; revised January 2004)
- Internet draft draft-ietf-mpls-soft-preemption-02.txt, *MPLS Traffic Engineering Soft Preemption* (no expiration; revised March 2004)
- Internet draft draft-ietf-ppvpn-rfc2547bis-03.txt, *BGP/MPLS VPNs*
- Internet draft draft-martini-l2circuit-encap-mpls-07.txt, *Encapsulation Methods for Transport of Layer 2 Frames Over IP and MPLS Networks*

The JUNOS Software has the following exceptions:

- A packet with a sequence number of 0 is treated as out of sequence.
- Any packet which does not have the next incremental sequence number is considered out of sequence.
- When out-of-sequence packets arrive, the expected sequence number for the neighbor is set to the sequence number in the Layer 2 circuit control word.
- Internet draft draft-martini-l2circuit-trans-mpls-14.txt, *Transport of Layer 2 Frames Over MPLS*
- Internet draft draft-raggarwa-mpls-p2mp-te-02.txt (except non-adjacent signaling for branch LSPs, make-before-break and fast reroute, and LSP hierarchy using P2P LSPs), *Establishing Point to Multipoint MPLS TE LSPs* (no expiration; revised January 2004)

The following documents provide information about traffic engineering:

- RFC 2702, *Requirements for Traffic Engineering Over MPLS*
- RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*
- Internet draft draft-ietf-isis-traffic-04.txt, *IS-IS Extensions for Traffic Engineering* (expires February 2002)
- Internet draft draft-katz-yeung-ospf-traffic-06.txt, *Traffic Engineering Extensions to OSPF* (expires April 2002)
- Internet draft draft-ietf-mpls-rsvp-te-p2mp-01.txt, *Extensions to RSVP-TE for Point to Multipoint TE LSPs* (expires June 2005)

To access Internet RFCs and drafts, go to the IETF website at <http://www.ietf.org>.

The JUNOS Software supports a proprietary MIB for MPLS objects; see the *JUNOS Network Management Configuration Guide* for more information.

Link-Layer Support

MPLS supports the following link-layer protocols, which are all supported in the JUNOS MPLS implementation:

- Point-to-Point Protocol (PPP)—Protocol ID 0x0281, Network Control Protocol (NCP) protocol ID 0x8281.
- Ethernet/Cisco High-level Data Link Control (HDLC)—Ethernet type 0x8847.
- Asynchronous Transfer Mode (ATM)—Subnetwork attachment point encoded (SNAP-encoded) Ethernet type 0x8847. Support is included for both point-to-point mode or nonbroadcast multiaccess (NBMA) mode. Support is not included for encoding MPLS labels as part of ATM virtual path identifier/virtual circuit identifier (VPI/VCI).
- Frame Relay—SNAP-encoded, Ethernet type 0x8847. Support is not included for encoding MPLS labels as part of Frame Relay data-link connection identifier (DLCI).
- Generic routing encapsulation (GRE) tunnel—Ethernet type 0x8847.

MPLS and Traffic Engineering

Traffic engineering allows you to control the path that data packets follow, bypassing the standard routing model, which uses routing tables. Traffic engineering moves flows from congested links to alternate links that would not be selected by the automatically computed destination-based shortest path. With traffic engineering, you can:

- Make more efficient use of expensive long-haul fibers.
- Control how traffic is rerouted in the face of single or multiple failures.
- Classify critical and regular traffic on a per-path basis.

The core of the traffic engineering design is based on building label-switched paths (LSPs) among routers. An LSP is connection-oriented, like a virtual circuit in Frame Relay or ATM. LSPs are not reliable: Packets entering an LSP do not have delivery guarantees, although preferential treatment is possible. LSPs also are similar to unidirectional tunnels in that packets entering a path are encapsulated in an envelope and switched across the entire path without being touched by intermediate nodes. LSPs provide fine-grained control over how packets are forwarded in a network. To provide reliability, an LSP can use a set of primary and secondary paths.

LSPs can be configured for BGP traffic only (traffic whose destination is outside of an autonomous system [AS]). In this case, traffic within the AS is not affected by the presence of LSPs. LSPs can also be configured for both BGP and interior gateway protocol (IGP) traffic; therefore, both intra-AS and inter-AS traffic is affected by the LSPs.

This section discusses the following topics:

- Label Description on page 25
- Label Allocation on page 26
- Routers in an LSP on page 28
- How a Packet Travels Along an LSP on page 28
- Types of LSPs on page 29
- Scope of LSPs on page 29
- Constrained-Path LSP Computation on page 29
- LSPs on an Overloaded Router on page 33
- Fate Sharing on page 33
- IGP Shortcuts on page 34
- Advertising LSPs into IGP on page 37

Label Description

Packets traveling along an LSP are identified by a label—a 20-bit, unsigned integer in the range 0 through 1,048,575. Labels 0 through 999,999 are for internal use. Labels 1,000,000 through 1,048,575 are unassigned by the JUNOS Software and are available for static LSPs. When you configure static LSPs, you can use only this range of labels.

Special Labels

Some of the reserved labels (in the 0 through 15 range) have well-defined meanings. For more complete details, see RFC 3032, *MPLS Label Stack Encoding*.

- 0, IPv4 Explicit Null label—This value is legal only when it is the sole label entry (no label stacking). It indicates that the label must be popped upon receipt. Forwarding continues based on the IP version 4 (IPv4) packet.
- 1, Router Alert label—When a packet is received with a top label value of 1, it is delivered to the local software module for processing.
- 2, IPv6 Explicit Null label—This value is legal only when it is the sole label entry (no label stacking). It indicates that the label must be popped on receipt. Forwarding continues based on the IP version 6 (IPv6) packet.
- 3, Implicit Null label—This label is used in the control protocol (LDP or RSVP) only to request label popping by the downstream router. It never actually appears in the encapsulation. Labels with a value of 3 should not be used in the data packet as real labels. No payload type (IPv4 or IPv6) is implied with this label.
- 4 through 15—Unassigned.

Special labels are commonly used between the egress and penultimate routers of an LSP. If the LSP is configured to carry IPv4 packets only, the egress router might signal the penultimate router to use 0 as a final-hop label. If the LSP is configured to carry IPv6 packets only, the egress router might signal the penultimate router to use 2 as a final-hop label.

The egress router might simply signal the penultimate router to use 3 as the final label, which is a request to perform penultimate-hop label popping. The egress router will not process a labeled packet; rather, it receives the payload (IPv4, IPv6, or others) directly, reducing one MPLS lookup at egress.

For label-stacked packets, the egress router receives an MPLS label packet with its top label already popped by the penultimate router. The egress router cannot receive label-stacked packets that use label 0 or 2. It typically requests label 3 from the penultimate router.

Label Allocation

In the JUNOS Software, label values are allocated per router. The display output shows only the label (for example, **01024**). Labels for multicast packets are independent of those for unicast packets. Currently, the JUNOS Software does not support multicast labels.

Labels are assigned by downstream routers relative to the flow of packets. A router receiving labeled packets (the next-hop router) is responsible for assigning incoming labels. A received packet containing a label that is unrecognized (unassigned) is dropped. For unrecognized labels, the router does not attempt to unwrap the label to analyze the network layer header, nor does it generate an Internet Control Message Protocol (ICMP) destination unreachable message.

A packet can carry a number of labels, organized as a last-in, first-out stack. This is referred to as a *label stack*. At a particular router, the decision about how to forward a labeled packet is based exclusively on the label at the top of the stack.

Figure 1 on page 27 shows the encoding of a single label. The encoding appears after data link layer headers, but before any network layer header.



Figure 1: Label Encoding

Figure 2 on page 27 illustrates the purpose of the class-of-service bits (also known as the EXP or experimental bits). Bits 20 and 21 specify the queue number. Bit 22 is the packet loss priority (PLP) bit used to specify the random early detection (RED) drop profile. For more information about class of service and the class-of-service bits, see “Configuring Class of Service for MPLS LSPs” on page 98.

Figure 2: Class-of-Service Bits

Operations on Labels

The router supports the following label operations:

- **Push**—Add a new label to the top of the packet. For IPv4 packets, the new label is the first label. The TTL and s bits are derived from the IP packet header. The MPLS CoS is derived from the queue number. If the push operation is performed on an existing MPLS packet, you will have a packet with two or more labels. This is called label stacking. The top label must have its s bit set to 0, and might derive CoS and time to live (TTL) from lower levels. The new top label in a label stack always initializes its TTL to 255, regardless of the TTL value of lower labels.
- **Pop**—Remove the label from the beginning of the packet. Once the label is removed, the TTL is copied from the label into the IP packet header, and the

underlying IP packet is forwarded as a native IP packet. In the case of multiple labels in a packet (label stacking), removal of the top label yields another MPLS packet. The new top label might derive CoS and TTL from a previous top label. The popped TTL value from the previous top label is not written back to the new top label.

- **Swap**—Replace the label at the top of the label stack with a new label. The S and CoS bits are copied from the previous label, and the TTL value is copied and decremented (unless the `no-decrement-ttl` or `no-propagate-ttl` statement is configured). A transit router supports a label stack of any depth.
- **Multiple Push**—Add multiple labels (up to three) on top of existing packets. This operation is equivalent to pushing multiple times.
- **Swap and Push**—Replace the existing top of the label stack with a new label, and then push another new label on top.

Routers in an LSP

Each router in an LSP performs one of the following functions:

- **Ingress router**—The router at the beginning of an LSP. This router encapsulates IP packets with an MPLS Layer 2 frame and forwards it to the next router in the path. Each LSP can have only one ingress router.
- **Egress router**—The router at the end of an LSP. This router removes the MPLS encapsulation, thus transforming it from an MPLS packet to an IP packet, and forwards the packet to its final destination using information in the IP forwarding table. Each LSP can have only one egress router. The ingress and egress routers in an LSP cannot be the same router.
- **Transit router**—Any intermediate router in the LSP between the ingress and egress routers. A transit router forwards received MPLS packets to the next router in the MPLS path. An LSP can contain zero or more transit routers, up to a maximum of 253 transit routers in a single LSP.

A single router can be part of multiple LSPs. It can be the ingress or egress router for one or more LSPs, and it also can be a transit router in one or more LSPs. The functions that each router supports depend on your network design.

How a Packet Travels Along an LSP

When an IP packet enters an 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.

The packet is then forwarded to the next router in the LSP. This router and all subsequent routers in the LSP do not examine any of the IP routing information in the labeled packet. Rather, they use the label to look up information in their label forwarding table. They then replace the old label with a new label and forward the packet to the next router in the path.

When the packet reaches the egress router, the label is removed, and the packet again becomes a native IP packet and is again forwarded based on its IP routing information.

Types of LSPs

There are three types of LSPs:

- Static LSPs—For static paths, you must manually assign labels on all routers involved (ingress, transit, and egress). No signaling protocol is needed. This procedure is similar to configuring static routes on individual routers. Like static routes, there is no error reporting, liveliness detection, or statistics reporting.
- LDP-signaled LSPs—See “LDP Overview” on page 329.
- RSVP-signaled LSPs—For signaled paths, RSVP is used to set up the path and dynamically assign labels. (RSVP signaling messages are used to set up signaled paths.) You configure only the ingress router. The transit and egress routers accept signaling information from the ingress router, and they set up and maintain the LSP cooperatively. Any errors encountered while establishing an LSP are reported to the ingress router for diagnostics. For signaled LSPs to work, a version of RSVP that supports tunnel extensions must be enabled on all routers.

There are two types of RSVP-signaled LSPs:

- Explicit-path LSPs—All intermediate hops of the LSP are manually configured. The intermediate hops can be strict, loose, or any combination of the two. Explicit path LSPs provide you with complete control over how the path is set up. They are similar to static LSPs but require much less configuration.
- Constrained-path LSPs—The intermediate hops of the LSP are automatically computed by the software. The computation takes into account information provided by the topology information from the IS-IS or OSPF link-state routing protocol, the current network resource utilization determined by RSVP, and the resource requirements and constraints of the LSP. For signaled constrained-path LSPs to work, either the IS-IS or OSPF protocol and the IS-IS or OSPF traffic engineering extensions must be enabled on all routers.

Scope of LSPs

For constrained-path LSPs, the LSP computation is confined to one IGP domain, and cannot cross any AS boundary. This prevents an AS from extending its IGP into another AS.

Explicit-path LSPs, however, can cross as many AS boundaries as necessary. Because intermediate hops are manually specified, the LSP does not depend on the IGP topology or a local forwarding table.

Constrained-Path LSP Computation

The Constrained Shortest Path First (CSPF) algorithm is an advanced form of the shortest-path-first (SPF) algorithm used in OSPF and IS-IS route computations. CSPF

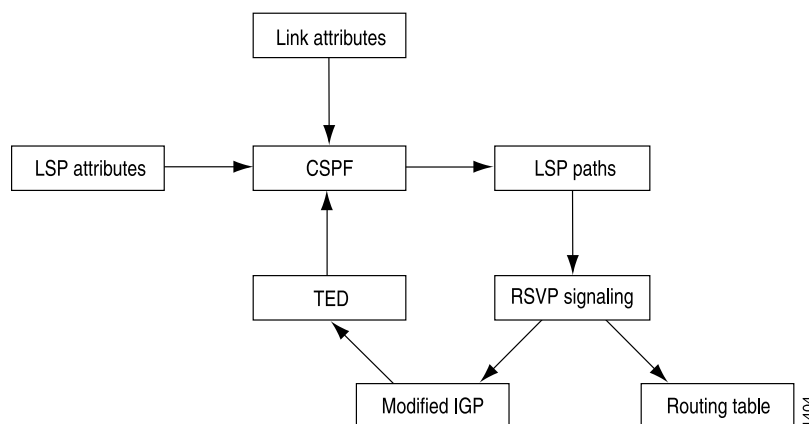
is used in computing paths for LSPs that are subject to multiple constraints. When computing paths for LSPs, CSPF considers not only the topology of the network, but also the attributes of the LSP and the links, and it attempts to minimize congestion by intelligently balancing the network load.

The constraints that CSPF considers include:

- LSP attributes
 - Administrative groups (that is, link color requirements)
 - Bandwidth requirements
 - Explicit route (strict or loose)
 - Hop limitations
 - Priority (setup and hold)
- Link attributes
 - Administrative groups (that is, link colors assigned to the link)
 - Reservable bandwidth of the links (static bandwidth minus the currently reserved bandwidth)

The data that CSPF considers comes from the following sources:

- Traffic engineering database—Provides CSPF with up-to-date topology information, the current reservable bandwidth of links, and the link colors. For the CSPF algorithm to perform its computations, a link-state IGP (such as OSPF or IS-IS) with special extensions is needed. For CSPF to be effective, the link-state IGP on all routers must support the special extensions. While building the topology database, the extended IGP must take into consideration the current LSPs and must flood the route information everywhere. Because changes in the reserved link bandwidth and link color cause database updates, an extended IGP tends to flood more frequently than a normal IGP. See Figure 3 on page 31 for a diagram of the relationships between these components.
- Currently active LSPs—Includes all the LSPs that should originate from the router and their current operational status (up, down, or timeout).

Figure 3: CSPF Computation Process

This section discusses the following topics:

- How CSPF Selects a Path on page 31
- Path Selection Tie-Breaking on page 32
- Computing Paths Offline on page 32

How CSPF Selects a Path

To select a path, CSPF follows these steps:

1. Computes LSPs one at a time, beginning with the highest priority LSP (the one with the lowest setup priority value). Among LSPs of equal priority, CSPF starts with those that have the highest bandwidth requirement.
2. Prunes the traffic engineering database of all the links that are not full duplex and do not have sufficient reservable bandwidth.
3. If the LSP configuration includes the **include** statement, prunes all links that do not share any included colors.
4. If the LSP configuration includes the **exclude** statement, prunes all links that contain excluded colors. If the link does not have a color, it is accepted.
5. Finds the shortest path toward the LSP's egress router, taking into account explicit-path constraints. For example, if the path must pass through Router A, two separate SPF's are computed, one from the ingress router to Router A, the other from Router A to the egress router.
6. If several paths have equal cost, chooses the one whose last-hop address is the same as the LSP's destination.
7. If several equal-cost paths remain, selects the one with the fewest number of hops.
8. If several equal-cost paths remain, applies the CSPF load-balancing rule configured on the LSP (least fill, most fill, or random).

Path Selection Tie-Breaking

If more than one path is available after the rules from the previous section have been applied, a tie-breaking rule is applied to choose the path for the LSP. There are three tie-breaking rules:

- **Random**—One of the remaining paths is picked at random. This rule tends to place an equal number of LSPs on each link, regardless of the available bandwidth ratio.
- **Least fill**—The path with the largest minimum available bandwidth ratio is preferred. This rule tries to equalize the reservation on each link.
- **Most fill**—The path with the smallest minimum available bandwidth ratio is preferred. This rule tries to fill a link before moving on to alternative links.

The rule used depends on the configuration. Random is the default rule.

For the other rules, the following definitions are needed:

- $\text{Reservable bandwidth} = \text{bandwidth of link} \times \text{subscription factor of link}$
- $\text{Available bandwidth} = \text{reservable bandwidth} - (\text{sum of the bandwidths of the LSPs traversing the link})$
- $\text{Available bandwidth ratio} = \text{available bandwidth} / \text{reservable bandwidth}$
- $\text{Minimum available bandwidth ratio (for a path)} = \text{the smallest available bandwidth ratio of the links in a path}$

Computing Paths Offline

The JUNOS Software provides online, real-time CSPF computation only; each router performs CSPF calculations independent of the other routers in the network. These calculations are based on currently available topology information—information that is usually recent, but not completely accurate. LSP placements are locally optimized, based on current network status.

To optimize links globally across the network, you can use an offline tool to perform the CSPF calculations and determine the paths for the LSPs. You can create such a tool yourself, or you can modify an existing network design tool to perform these calculations. You should run the tool periodically (daily or weekly) and download the results into the router. An offline tool should take the following into account when performing the optimized calculations:

- All the LSP's requirements
- All link attributes
- Complete network topology

LSPs on an Overloaded Router

An overloaded router is a router running IS-IS with its overload bit set in its IS-IS configuration. In this case, an MPLS LSP specifically refers to an RSVP-signaled or LDP-signaled LSP. In the case of RSVP, it applies to both CSPF and non-CSPF LSPs.


You cannot establish transit LSPs through an overloaded router. However, you can configure ingress and egress LSPs through an overloaded router.



NOTE: When you set the overload bit on an IS-IS router, all LSPs transiting through it are recomputed and rerouted away from it. If the recomputation fails, no additional attempt to reconfigure the LSP is made, and the affected LSPs are disconnected.

An example of when you might want to establish transit LSPs through an overloaded router is illustrated in Figure 4 on page 33, which shows an aggregation router (Router A) dual-homed on two core routers (Router B and Router C). You want to include the aggregation router in the LSP mesh, but transit LSPs should not pass through it, because it is a less capable router with relatively low-bandwidth uplinks to the core. Certain failure and rerouting scenarios could make it impossible for the aggregation router to establish some of its LSPs. Consequently, you run the router in a steady state with the overload bit set, but you are still able to establish ingress and egress LSPs through it.

Figure 4: Aggregation Router A Dual-Homed on Core Routers B and C



Fate Sharing

Fate sharing allows you to create a database of information that CSPF uses to compute one or more backup paths to use in case the primary path becomes unstable. The database describes the relationships between elements of the network, such as routers and links. You can specify one or more elements within a group.

Through fate sharing, you can configure backup paths that minimize the number of shared links and fiber paths with the primary paths as much as possible, to ensure

that if a fiber is cut, the minimum amount of data is lost and a path still exists to the destination.


For a backup path to work optimally, it must not share links or physical fiber paths with the primary path, ensuring that a single point of failure will not affect the primary and backup paths simultaneously. For more information about fate sharing, see the *JUNOS Routing Protocols Configuration Guide*.

IGP Shortcuts

Link-state protocols, such as OSPF and IS-IS, use the SPF algorithm to compute the shortest-path tree to all nodes in the network. The results of such computations can be represented by the destination node, next-hop address, and output interface, where the output interface is a physical interface. LSPs can be used to augment the SPF algorithm, for the purposes of resolving BGP next hops. On the node performing the calculations, LSPs appear to be logical interfaces directly connected to remote nodes in the network. If you configure the IGP to treat LSPs the same as a physical interface and use the LSPs as a potential output interface, the SPF computation results are represented by the destination node and output LSP, effectively using the LSP as a shortcut through the network to the destination.


As an illustration, begin with a typical SPF tree (see Figure 5 on page 34).

Figure 5: Typical SPF Tree, Sourced from Router A



If an LSP connects Router A to Router D and if IGP shortcuts are enabled on Router A, you might have the SPF tree shown in Figure 6 on page 34.

Figure 6: Modified SPF Tree, Using LSP A-D as a Shortcut



Note that Router D is now reachable through LSP A-D. When computing the shortest path to reach Router D, Router A has two choices:


- Use IGP path A-B-D.
- Use LSP A-D.

Router A decides between the two choices by comparing the IGP metrics for path A-B-D with the LSP metrics for LSP A-D. If the IGP metric is lower, path A-B-D is chosen (Figure 5 on page 34). If the LSP metric is lower, LSP A-D is used (Figure 6 on page 34). If both metrics are equal, LSP A-D is chosen because LSP paths are preferred over IGP paths.

Note that Routers E and F are also reachable through LSP A-D, because they are downstream from Router D in the SPF tree.

Assuming that another LSP connects Router A to Router E, you might have the SPF tree shown in Figure 7 on page 35.

Figure 7: Modified SPF Tree, Using LSP A-D and LSP A-E as Shortcuts



This section discusses the following topics:

- Enabling IGP Shortcuts on page 35
- LSPs Qualified in Shortcut Computations on page 36
- IGP Shortcut Applications on page 36
- IGP Shortcuts and Routing Table on page 37
- IGP Shortcuts and VPN Environments on page 37

Enabling IGP Shortcuts

IGP shortcuts are supported for both IS-IS and OSPF. A link-state protocol is required for IGP shortcuts. Shortcuts are disabled by default. For information about enabling IGP shortcuts for IS-IS and OSPF, see the *JUNOS Routing Protocols Configuration Guide*. You can enable IGP shortcuts on a per-router basis; you do not need to enable shortcuts globally. A router's shortcut computation does not depend on another router performing similar computations, and shortcuts performed by other routers are irrelevant.

LSPs Qualified in Shortcut Computations

Not all LSPs are used in IGP shortcuts. Only those LSPs whose egress point (using the `to` statement) matches the router ID of the egress node are considered. Other LSPs, whose egress point matches the egress node interface address, are ignored in IGP shortcuts.


There are exceptions, however. If an LSP has an alias egress point (using the `install` statement) and it matches certain router IDs, it is included in the shortcut computation as well. If multiple equal metric LSPs destined to the same router ID exist, traffic can load-share among them.

IGP Shortcut Applications

You can use shortcuts to engineer traffic traveling toward destination nodes that do not support MPLS LSPs. For example, in Figure 7 on page 35, traffic traveling toward Router F enters LSP A–E. You can control traffic between Router A and Router F by manipulating LSP A–E; you do not need to explicitly set up an LSP between Router A and Router F.


In Figure 8 on page 36, all traffic from Region 1 to Region 2 traverses LSP A–B if IGP shortcuts are enabled on the ingress router (Router A), permitting aggregation of interregional traffic into one LSP. To perform traffic engineering on the interregional traffic, you have to manipulate LSP A-B only, which avoids creating n^2 LSPs from all routers in Region 1 to all routers in Region 2 and allows efficient resource controls on the backbone network.

Figure 8: IGP Shortcuts



Shortcuts allow you to deploy LSPs into a network in an incremental, hierarchical fashion. In Figure 9 on page 36, each region can choose to implement traffic engineering LSPs independently, without requiring cooperation from other regions. Each region can choose to deploy intraregion LSPs to fit the region's bandwidth needs, at the pace appropriate for the region.

Figure 9: IGP Shortcuts in a Bigger Network



When intraregion LSPs are in place, interregional traffic automatically traverses the intraregion LSPs as needed, eliminating the need for a full mesh of LSPs between edge routers. For example, traffic from Router A to Router D traverses LSPs A–B, B–C, and C–D.

IGP Shortcuts and Routing Table

IGP typically performs two independent computations. The first is performed without considering any LSP. The result of the computation is stored in the `inet.0` table. This step is no different from traditional SPF computations and is always performed even if IGP shortcut is disabled.

The second computation is performed considering only LSPs as a logical interface. Each LSP's egress router is considered. The list of destinations whose shortest path traverses the egress router (established during the first computation) is placed in the `inet.3` routing table. These destinations are given the egress router of the LSP as a next hop, enabling BGP on the local router to use these LSPs to access BGP next hops beyond the egress router. Normally, BGP can use only LSPs that terminate at the BGP next hop. Note that BGP is the only protocol that uses the `inet.3` routing table. Other protocols will not route traffic through these LSPs.

If traffic engineering for IGP and BGP is enabled (see “IGP and BGP Destinations” on page 41), IGP moves all routes in `inet.3` into `inet.0`, merging all routes while emptying the `inet.3` table. The number of routes in `inet.0` will be exactly the same as before. Route next-hops can traverse a physical interface, an LSP, or the combination of the two if the metrics are equal.

IGP shortcuts are enabled on a per-node basis. You do not need to coordinate with other nodes.


IGP Shortcuts and VPN Environments

IGP shortcuts (configured under the `[edit protocols mpls traffic-engineering bgp-igp]` hierarchy level and under the `[edit protocols ospf traffic-engineering shortcuts]` hierarchy level) do not work in VPN environments. IGP shortcuts move routes in the `inet.3` routing table to the `inet.0` routing table. Virtual private network (VPN) IBGP (which belongs to family `inet-vpn`) relies on next hops that are in the `inet.3` table, so IGP shortcuts are incompatible with VPNs.

Advertising LSPs into IGP

You can configure your IGP to treat an LSP as a link. IGP shortcuts allow only the ingress router of an LSP to use the LSP in its SPF computation. However, other routers on the network do not know of the existence of that LSP, so they cannot use it. This can lead to suboptimal traffic engineering. In addition, only BGP can use an IGP shortcut to an LSP. When you advertise an LSP as a link into the IGP, all traffic can traverse it, and all routers know about it.

As an example, consider the network shown in Figure 10 on page 38.

Figure 10: SPF Computations with Advertised LSPs

Assume that Router A is computing a path to Router D. The link between Router E and Router F has a metric of 20; all other links have a metric of 10. Here, the path chosen by Router A is A-B-C-D, which has a metric of 30, instead of A-E-F-D, which has a metric of 40.

If Router E has an LSP to Router D with a metric of 15, you want traffic from Router A to Router D to use the path A-E-D, which has a metric of 25, instead of the path A-B-C-D. However, because Router A does not know about the LSP between Router E and Router D, it cannot route traffic through this path.

For all routers on the network to know about the LSP between Router E and Router D, you need to advertise it. This advertisement announces the LSP as a unidirectional, point-to-point link in the link-state database, and all routers can compute paths using the LSP. The link-state database maintains information about the AS topology and contains information about the router's local state (for example, the router's usable interfaces and reachable neighbors). In Figure 10 on page 38, Router A will see the link from Router E to Router D and route traffic along this lower-metric path.

Because an LSP is announced as a unidirectional link, you might need to configure a reverse LSP (one that starts at the egress router and ends at the ingress router) so that the SPF bidirectionality check succeeds. As a step in the SPF computation, IS-IS considers a link from Router E to Router D. Before IS-IS uses any link, it verifies that there is a link from Router D to Router E (there is bidirectional connectivity between router E and D). Otherwise, the SPF computation will not use an announced LSP.

When an LSP is advertised to the IGP, the advertising router uses the LSP as the forwarding path for regular routes after installing them in the `inet.0` routing table. All packets traversing the router could be forwarded through the LSP. Conversely, IGP shortcuts are used only to forward packets that are following BGP routes.



NOTE: Do not configure IGP shortcuts and advertise LSPs to the IGP at the same time.

IP and MPLS Packets on Aggregated Interfaces

You can send IP and MPLS packets over aggregated interfaces. To the IP or MPLS session, there is a single LSP composed of the aggregated interfaces. Packets sent to an LSP that is part of an aggregated interface are redistributed over the aggregated member interfaces.

Sending IP and MPLS packets over aggregated interfaces has the following benefits:

- Bandwidth aggregation—You can increase the number of MPLS packet flows sent over each connection. In MPLS, a set of packets sharing the same label is considered a part of the same flow.
- Link redundancy—If a link or a line card failure affects an aggregate member link, the traffic flowing across that link is immediately forwarded across one of the remaining links.

JUNOS supports aggregated SONET and Ethernet interfaces.

Note that the JUNOS implementation of IP and MPLS over aggregated interfaces (aggregated Ethernet devices only) complies with IEEE 802.3ad.

For information about how to configure aggregated Ethernet or aggregated SONET interfaces, see the *JUNOS Network Interfaces Configuration Guide*.

MPLS Applications

In the JUNOS Software implementation of MPLS, establishing an LSP installs on the ingress router a host route (a 32-bit mask) toward the egress router. The address of the host route is the destination address of the LSP. By default, the route has a preference value of 7, a value that is higher than all routes except direct interface and static routes. The 32-bit mask ensures that the route is more specific (that is, a longer match) than all other subnet routes. The host routes can be used to traffic-engineer BGP destinations only, or both IGP and BGP destinations.

This section discusses the following topics:

- BGP Destinations on page 39
- IGP and BGP Destinations on page 41
- Selecting a Forwarding LSP Next Hop on page 41

BGP Destinations

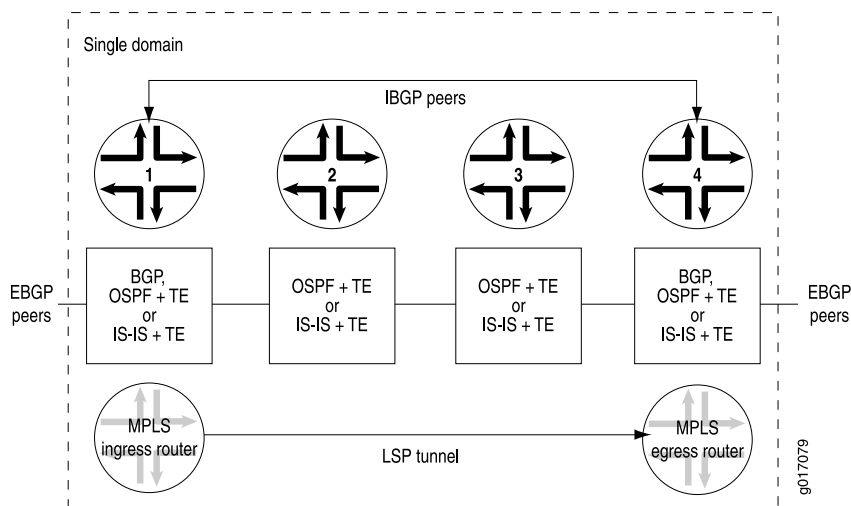
You can configure MPLS to control the paths that traffic takes to destinations outside an AS.

Both IBGP and EBGP take advantage of the LSP host routes without requiring extra configuration. BGP compares the BGP next-hop address with the LSP host route. If a match is found, the packets for the BGP route are label-switched over the LSP. If multiple BGP routes share the same next-hop address, all the BGP routes are mapped to the same LSP route, regardless of which BGP peer the routes are learned from. If the BGP next-hop address does not match an LSP host route, BGP routes continue to be forwarded based on the IGP routes within the routing domain. In general, when both an LSP route and an IGP route exist for the same BGP next-hop address, the one with the lowest preference is chosen.

Figure 11 on page 40 shows an MPLS topology that illustrates how MPLS and LSPs work. This topology consists of a single domain with four routers. The two routers

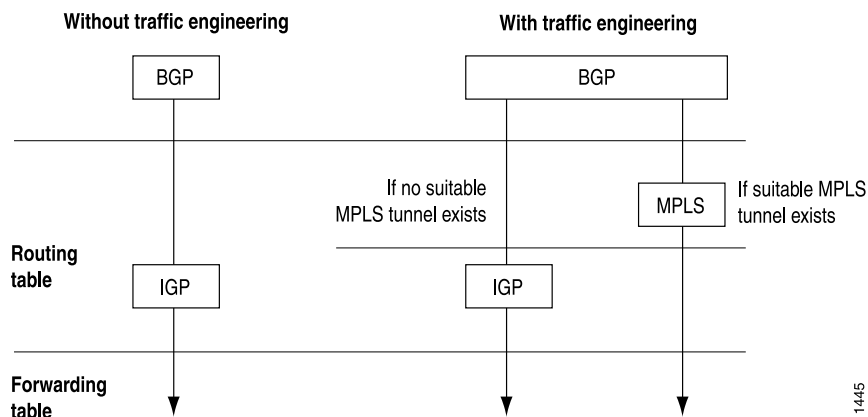
at the edges of the domain, Router 1 and Router 4, are running EBGP to communicate with peers outside the domain and IBGP to communicate between themselves. For intradomain communication, all four routers are running an IGP. Finally, an LSP tunnel exists from Router 1 to Router 4.

Figure 11: MPLS Application Topology



When BGP on Router 1 receives prefixes from Router 4, it must determine how to reach a BGP next-hop address. Typically, when traffic engineering is not enabled, BGP uses IGP routes to determine how to reach next-hop addresses. (See the left side of Figure 12 on page 40.) However, when traffic engineering is enabled, if the BGP next-hop matches the LSP tunnel endpoint (that is, the MPLS egress router), those prefixes enter the LSP tunnel. (To track these prefixes, look at the **Active Route** field in the `show mpls lsp` command output or at the output of the `show route label-switched-path path-name` command.) If the BGP next hop does not match an LSP tunnel endpoint, those prefixes are sent following the IGP's shortest path. (See Figure 12 on page 40.)

Figure 12: How BGP Determines How to Reach Next-Hop Addresses



IGP and BGP Destinations

You can configure MPLS to control the paths that traffic takes to destinations within an AS.

When traffic engineering is for BGP destinations only, the MPLS host routes are installed in the `inet.3` routing table (see Figure 13 on page 42), separate from the routes learned from other routing protocols. Not all `inet.3` routes are downloaded into the forwarding table. Packets directly addressed to the egress router do not follow the LSP, which prevents routes learned from LSPs from overriding routes learned from IGPs or other sources.

Traffic within a domain, including BGP control traffic between BGP peers, is not affected by LSPs. MPLS affects interdomain traffic only; that is, it affects only those BGP prefixes that are learned from an external domain. MPLS does not disrupt intradomain traffic, so IS-IS or OSPF routes remain undisturbed. If you issue a `ping` or `traceroute` command to any destination within the domain, the `ping` or `traceroute` packets follow the IGP path. However, if you issue a `ping` or `traceroute` command from Router 1 in Figure 11 on page 40 (the LSP ingress router) to a destination outside of the domain, the packets use the LSP tunnel.

When traffic engineering for IGP and BGP destinations is enabled, the MPLS host routes are installed in the `inet.0` table (see Figure 14 on page 43) and downloaded into the forwarding table. Any traffic destined to the egress router could enter the LSP. In effect, it moves all the routes in `inet.3` into `inet.0`, causing the `inet.3` table to be emptied.


RSVP packets automatically avoid all MPLS LSPs, including those established by RSVP or LDP. This prevents placing one RSVP session into another LSP, or in other words, nesting one LSP into another.

Selecting a Forwarding LSP Next Hop

If more than one LSP tunnel to a BGP next hop exists, the prefixes learned from the BGP next hop are randomly divided among the LSP tunnels. To control which LSP BGP uses to forward data for a given prefix, use the `install-nexthop` statement in the export policy applied to the forwarding table. For more information, see the *JUNOS Routing Protocols Configuration Guide*.

MPLS and Routing Tables

The IGPs and BGP store their routing information in the `inet.0` routing table, the main IP routing table. If `traffic-engineering bgp` is configured, thereby allowing only BGP to use MPLS paths for forwarding traffic, MPLS path information is stored in a separate routing table, `inet.3`. Only BGP accesses the `inet.3` routing table. BGP uses both `inet.0` and `inet.3` to resolve next-hop addresses. If `traffic-engineering bgp-igp` is configured, thereby allowing the IGPs to use MPLS paths for forwarding traffic, MPLS path information is stored in the `inet.0` routing table. (Figure 13 on page 42 and Figure 14 on page 43 illustrate the routing tables in the two traffic engineering configurations.)

Figure 13: Routing and Forwarding Tables, traffic-engineering bgp

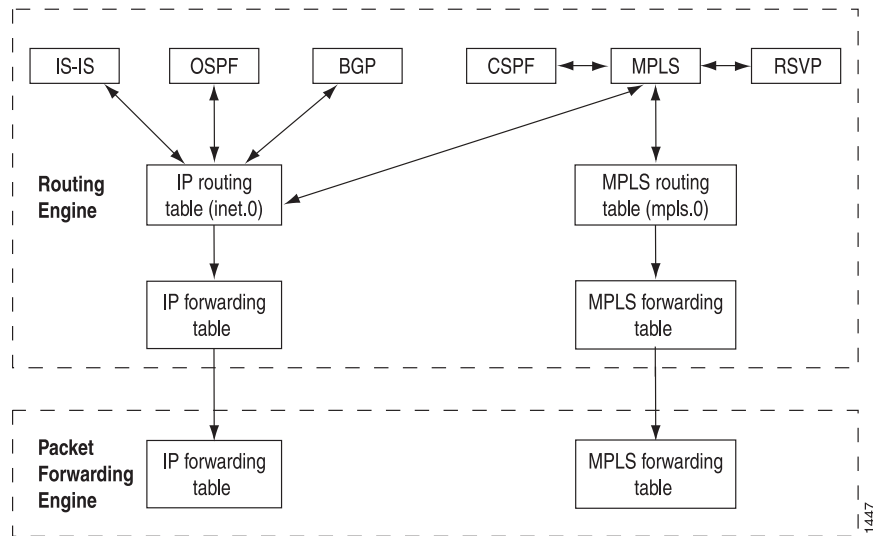
The `inet.3` routing table contains the host address of each LSP's egress router. This routing table is used on ingress routers to route packets to the destination egress router. BGP uses the `inet.3` routing table on the ingress router to help in resolving next-hop addresses.

MPLS also maintains an MPLS path routing table (`mpls.0`), which contains a list of the next label-switched router in each LSP. This routing table is used on transit routers to route packets to the next router along an LSP.

Typically, the egress router in an LSP does not consult the `mpls.0` routing table. (This router does not need to consult `mpls.0` because the penultimate router in the LSP either changes the packet's label to a value of 0 or pops the label.) In either case, the egress router forwards it as an IPv4 packet, consulting the IP routing table, `inet.0`, to determine how to forward the packet.

When a transit or egress router receives an MPLS packet, information in the MPLS forwarding table is used to determine the next transit router in the LSP or to determine that this router is the egress router.

When BGP resolves a next-hop prefix, it examines both the `inet.0` and `inet.3` routing tables, seeking the next hop with the lowest preference. If it finds a next-hop entry with an equal preference in both routing tables, BGP prefers the entry in the `inet.3` routing table.

Figure 14: Routing and Forwarding Tables, traffic-engineering bgp-igp

Generally, BGP selects next-hop entries in the `inet.3` routing table because their preferences are always lower than OSPF and IS-IS next-hop preferences. When you configure LSPs, you can override the default preference for MPLS LSPs, which might alter the next-hop selection process.

When BGP selects a next-hop entry from the `inet.3` routing table, it installs that LSP into the forwarding table in the Packet Forwarding Engine, which causes packets destined for that next hop to enter and travel along the LSP. If the LSP is removed or fails, the path is removed from the `inet.3` routing table and from the forwarding table, and BGP reverts to using a next hop from the `inet.0` routing table.

MPLS and Traffic Protection

Typically, when an LSP fails, the router immediately upstream from the failure signals the outage to the ingress router. The ingress router calculates a new path to the egress router, establishes the new LSP, and then directs the traffic from the failed path to the new path. This rerouting process can be time-consuming and prone to failure. For example, the outage signals to the ingress router might get lost, or the new path might take too long to come up, resulting in significant packet drops. The JUNOS Software provides several complementary mechanisms for protecting against LSP failures:

- Standby secondary paths—You can configure primary and secondary paths. You configure secondary paths with the `standby` statement. To activate traffic protection, you need to configure these standby paths only on the ingress router. If the primary path fails, the ingress router immediately reroutes traffic from the failed path to the standby path, thereby eliminating the need to calculate a new route and signal a new path. For information about configuring standby LSPs, see “Configuring Hot Standby of Secondary Paths” on page 106.
- Fast reroute—You configure fast reroute on an LSP to minimize the effect of a failure in the LSP. Fast reroute enables a router upstream from the failure to route around the failure quickly to the router downstream of the failure. The

upstream router then signals the outage to the ingress router, thereby maintaining connectivity before a new LSP is established. For a detailed overview of fast reroute, see “Fast Reroute” on page 44. For information about configuring fast reroute, see “Configuring Fast Reroute” on page 78.

- Link protection—You can configure link protection to help ensure that traffic traversing a specific interface from one router to another can continue to reach its destination in the event that this interface fails. When link protection is configured for an interface and configured for an LSP that traverses this interface, a bypass LSP is created that handles this traffic if the interface fails. The bypass LSP uses a different interface and path to reach the same destination. For information about configuring link protection, see “Configuring Link Protection on Interfaces Used by LSPs” on page 282.

When standby secondary path, and fast reroute or link protection are configured on an LSP, full traffic protection is enabled. When a failure occurs in an LSP, the router upstream from the failure routes traffic around the failure and notifies the ingress router of the failure. This rerouting keeps the traffic flowing while waiting for the notification to be processed at the ingress router. After receiving the failure notification, the ingress router immediately reroutes the traffic from the patched primary path to the more optimal standby path.

Fast reroute and link protection provide a similar type of traffic protection. Both features provide a quick transfer service and employ a similar design. Fast reroute and link protection are both described in RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*. However, you need to configure only one or the other. Although you can configure both, there is little, if any, benefit in doing so.

Fast Reroute


The following sections provide an overview of how fast reroute works:

- Fast Reroute Overview on page 44
- Detour Merging Process on page 46
- Detour Computations on page 47
- Fast Reroute Path Optimization on page 47


Fast Reroute Overview

Fast rerouting is accomplished by precomputing and preestablishing a number of detours along the LSP. In case of a network failure on the current LSP path, the traffic is quickly routed to one of the detours. Figure 15 on page 45 illustrates an LSP from Router A to Router F, showing the established detours. Each detour is established by an upstream node to avoid the link toward the immediate downstream node and the immediate downstream node itself. Each detour might traverse through one or more label-switched routers that are not shown in the figure.

Fast reroute protects traffic against any single point of failure between the ingress and egress routers. If there are multiple failures along an LSP, fast reroute itself might fail. Also, fast reroute does not protect against failure of the ingress or egress routers.

Figure 15: Detours Established for an LSP Using Fast Reroute

If a node detects that a downstream link has failed (using a link-layer-specific liveness detection mechanism) or that a downstream node has failed (for example, using the RSVP neighbor hello protocol), the node quickly switches the traffic to the detour and, at the same time, signals the ingress router about the link or node failure. Figure 16 on page 45 illustrates the detour taken when the link between Router B and Router C fails.

Figure 16: Detour After the Link from Router B to Router C Fails

If the network topology is not rich enough (there are not enough routers with sufficient links to other routers), some of the detours might not succeed. For example, the detour from Router A to Router C in Figure 15 on page 45 cannot traverse link A-B and Router B. If such a path is not possible, the detour does not occur.

Note that after the node switches traffic to the detour, it might switch the traffic again to a newly calculated detour soon after. This is because the initial detour route might not be the best route. To make rerouting as fast as possible, the node switches traffic onto the initial detour without first verifying that the detour is valid. Once the switch is made, the node recomputes the detour. If the node determines that the initial detour is still valid, traffic continues to flow over this detour. If the node determines that the initial detour is no longer valid, it again switches the traffic to a newly computed detour.



NOTE: If you issue `show` commands after the node has switched traffic to the initial detour, the node might indicate that the traffic is still flowing over the original LSP. This situation is temporary and should correct itself quickly.

The time required for a fast-rerouting detour to take effect depends on two independent time intervals:

- Amount of time to detect that there is a link or node failure—This interval depends greatly on the link layer in use and the nature of the failure. For example, failure

detection on an SONET/SDH link typically is much faster than on a Gigabit Ethernet link, and both are much faster than detection of a router failure.


- Amount of time required to splice the traffic onto the detour—This operation is performed by the Packet Forwarding Engine, which requires little time to splice traffic onto the detour. The time needed can vary depending on the number of LSPs being switched to detours.

Fast reroute is a short-term patch to reduce packet loss. Because detour computation might not reserve adequate bandwidth, the detours might introduce congestion on the alternate links. The ingress router is the only router that is fully aware of LSP policy constraints and, therefore, is the only router able to come up with adequate long-term alternate paths.

Detours are created by use of RSVP and, like all RSVP sessions, they require extra state and overhead in the network. For this reason, each node establishes at most one detour for each LSP that has fast reroute enabled. Creating more than one detour for each LSP increases the overhead, but serves no practical purpose.

To reduce network overhead further, each detour attempts to merge back into the LSP as soon as possible after the failed node or link. If you can consider an LSP that travels through n router nodes, it is possible to create $n - 1$ detours. For instance, in Figure 17 on page 46, the detour tries to merge back into the LSP at Router D instead of at Router E or Router F. Merging back into the LSP makes the detour scalability problem more manageable. If topology limitations prevent the detour from quickly merging back into the LSP, detours merge with other detours automatically.

Figure 17: Detours Merging into Other Detours



Detour Merging Process

This section describes the process used by a router to determine which LSP to select when the router receives path messages from different interfaces with identical Session and Sender Template objects. When this occurs, the router needs to merge the path states.

The router employs the following process to determine when and how to merge path states:

- When all the path messages do not include a fast reroute or a detour object, or when the router is the egress of the LSP, no merging is required. The messages are processed according to RSVP traffic engineering.
- Otherwise, the router *must* record the path state in addition to the incoming interface. If the path messages do not share the same outgoing interface and

next-hop router, the router considers them to be independent LSPs and does not merge them.

- For all the path messages that share the same outgoing interface and next-hop router, the router uses the following process to select the final LSP:
 - If only one LSP originates from this node, select it as the final LSP.
 - If only one LSP contains a fast reroute object, select it as the final LSP.
 - If there are several LSPs and some of them have a detour object, eliminate those containing a detour object from the final LSP selection process.
 - If several final LSP candidates remain (that is, there are still both detour and protected LSPs), select the LSPs with fast reroute objects.
 - If none of the LSPs have fast reroute objects, select the ones without detour objects. If all the LSPs have detour objects, select them all.
 - Of the remaining LSP candidates, eliminate from consideration those that traverse nodes that other LSPs avoid.
 - If several candidate LSPs still remain, select the one with the shortest explicit route object (ERO) path length. If more than one LSP has the same path length, select one randomly.
- Once the final LSP has been identified, the router must transmit only the path messages that correspond to this LSP. All other LSPs are considered merged at this node.

Detour Computations

Computing and setting up detours is done independently at each node. On a node, if an LSP has fast reroute enabled and if a downstream link or node can be identified, the router performs a Constrained Shortest Path First (CSPF) computation using the information in the local traffic engineering database. For this reason, detours rely on your IGP supporting traffic engineering extensions. Without the traffic engineering database, detours cannot be established.

CSPF initially attempts to find a path that skips the next downstream node. Attempting to find this path provides protection against downstream failures in either nodes or links. If a node-skipping path is not available, CSPF attempts to find a path on an alternate link to the next downstream node. Attempting to find an alternate link provides protection against downstream failures in links only. Detour computations might not succeed the first time. If a computation fails, the router recomputes detours approximately once every refresh interval until the computation succeeds. The RSVP metric for each detour is set to a value in the range from 10,000 through 19,999.

Fast Reroute Path Optimization

A fast reroute protection path is nondeterministic. The actual protection path of a particular node depends on the history of the LSP and the network topology when the fast reroute path was computed. The lack of deterministic behavior can lead to operational difficulties and poorly optimized paths after multiple link flaps in a

network. Even in a small network, after a few link flaps fast reroute paths can traverse an arbitrarily large number of nodes and can remain in that state indefinitely. This is inefficient and makes the network less predictable.

Fast reroute optimization addresses this deficiency. It provides a global path optimization timer, allowing you to optimize all LSPs that have fast reroute enabled and a detour path up and running. The timer value can be varied depending on the expected RE processing load.

The fast reroute optimization algorithm is based on the IGP metric only. As long as the new path's IGP metric is lower than the old path's, the CSPF result is accepted, even if the new path might be more congested (higher bandwidth utilization) or traverses more hops.

In conformance with RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*, when a new path is computed and accepted for fast reroute optimization, the existing detour is destroyed first and then the new detour is established. To prevent traffic loss, detours actively protecting traffic are not optimized.

Automatic Bandwidth Allocation

Automatic bandwidth allocation allows an MPLS tunnel to automatically adjust its bandwidth allocation based on the volume of traffic flowing through the tunnel. You can configure an LSP with minimal bandwidth; this feature can dynamically adjust the LSP's bandwidth allocation based on current traffic patterns. The bandwidth adjustments do not interrupt traffic flow through the tunnel.

You set a sampling interval on an LSP configured with automatic bandwidth allocation. The average bandwidth is monitored during this interval. At the end of the interval, an attempt is made to signal a new path for the LSP with the bandwidth allocation set to the maximum average value for the preceding sampling interval. If the new path is successfully established and the original path is removed, the LSP is switched over to the new path. If a new path is not created, the LSP continues to use its current path until the end of the next sampling interval, when another attempt is made to establish a new path. Note that you can set minimum and maximum bandwidth values for the LSP.


During the automatic bandwidth allocation interval, the router might receive a steady increase in traffic (increasing bandwidth utilization) on an LSP, potentially causing congestion or packet loss. To prevent this, you can define a second trigger to prematurely expire the automatic bandwidth adjustment timer before the end of the current adjustment interval.

Point-to-Multipoint LSPs

A point-to-multipoint MPLS LSP is an RSVP-signaled LSP with a single source and multiple destinations. By taking advantage of the MPLS packet replication capability of the network, point-to-multipoint LSPs avoid unnecessary packet replication at the ingress router. Packet replication takes place only when packets are forwarded to two or more different destinations requiring different network paths.

This process is illustrated in Figure 18 on page 49. Router PE1 is configured with a point-to-multipoint LSP to Routers PE2, PE3, and PE4. When Router PE1 sends a packet on the point-to-multipoint LSP to Routers P1 and P2, Router P1 replicates the packet and forwards it to Routers PE2 and PE3. Router P2 sends the packet to Router PE4.

Figure 18: Point-to-Multipoint LSPs



The following are some of the properties of point-to-multipoint LSPs:

- A point-to-multipoint LSP allows you to use MPLS for point-to-multipoint data distribution. This functionality is similar to that provided by IP multicast.
- You can add and remove branch LSPs from a main point-to-multipoint LSP without disrupting traffic. The unaffected parts of the point-to-multipoint LSP continue to function normally.
- You can configure a node to be both a transit and an egress router for different branch LSPs of the same point-to-multipoint LSP.
- You can enable link protection on a point-to-multipoint LSP. Link protection can provide a bypass LSP for each of the branch LSPs that make up the point-to-multipoint LSP. If any of the primary paths fail, traffic can be quickly switched to the bypass.
- You can configure branch LSPs either statically, dynamically, or as a combination of static and dynamic LSPs.
- You can enable graceful restart on point-to-multipoint LSPs.

For information on how to configure point-to-multipoint LSPs, see “Point-to-Multipoint LSP Configuration Guidelines” on page 141. This feature is described in detail in the Internet drafts draft-raggarwa-mpls-p2mp-te-02.txt, *Establishing Point to Multipoint MPLS TE LSPs*, and draft-ietf-mpls-rsvp-te-p2mp-02.txt, *Extensions to RSVP-TE for Point to Multipoint TE LSPs*.

MPLS Load Balancing Based on the IP Header and MPLS Labels

Juniper Networks routers can load-balance on a per-packet basis in MPLS. Load balancing can be performed on information in both the IP header and on up to three MPLS labels, providing a more uniform distribution of MPLS traffic to next hops. This feature is enabled on supported platforms by default and requires no configuration.



NOTE: This feature is only available on M320 and T Series Core Routers with enhanced Flexible PIC Concentrators (FPCs).

The following information is extracted from the packet and used to load-balance the MPLS traffic:

- Interface index—24 bits
- MPLS label stack—bits 0 through 23 (the TTL bits are not examined)
- IPv4 header information:
 - Protocol—8 bits
 - Destination address—32 bits
 - Source address—32 bits
 - Source port—16 bits
 - Destination port—16 bits
- IPv6 header information:
 - Next header—8 bits
 - Least significant 4 bytes of destination address
 - Least significant 4 bytes of source address
 - Source port—16 bits
 - Destination port—16 bits

In summary, MPLS load balancing is performed using the following fields:

Interface index + MPLS label + IP header (IPv4 or IPv6)

Chapter 4

MPLS Router Configuration Guidelines

This chapter discusses the following topics:

- Minimum MPLS Configuration on page 51
- Configuring the Ingress Router for MPLS-Signaled LSPs on page 52
- Examples: Configuring Ingress Routers for MPLS-Signaled LSPs on page 57
- Example: Configuring a Constrained-Path LSP for Which JUNOS Makes All Forwarding Decisions on page 57
- Example: Configuring an Explicit-Path LSP on page 57
- Example: Configuring a Constrained-Path LSP for Which JUNOS Makes Most Forwarding Decisions and Considers Hop Constraints on page 58
- Example: Configuring a Constrained-Path LSP for Which JUNOS Makes Most Forwarding Decisions and the Secondary Path Is Explicit on page 59
- Configuring the Intermediate and Egress Routers for MPLS-Signaled LSPs on page 60
- Configuring a Default Route for Unlabeled MPLS Packets on page 60
- Improving Traffic Engineering Database Accuracy with RSVP PathErr Messages on page 60
- Configuring MPLS-Signaled LSPs to Use GRE Tunnels on page 62
- Tunneling IPv6 Traffic over MPLS IPv4 Networks on page 64
- Configuring ICMP Message Tunneling on page 67

Minimum MPLS Configuration

To enable MPLS on the router, you must include at least the following statements. This minimum configuration enables MPLS on a logical interface. All other MPLS configuration statements are optional. Note that this configuration does nothing more than enable MPLS on the router and on the specified interface.

Include the **family mpls** statement:

```
family mpls;
```

You can include this statement at the following hierarchy levels:

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

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

Include the interface in the MPLS and RSVP protocol configuration:

```
mpls {
  interface (interface-name | all); # Required to enable MPLS on the interface
}
rsvp { # Required for RSVP-signaled MPLS only
  interface interface-name;
}
```

You can configure these statements at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

For every interface you enable, two special routes are installed automatically in the MPLS forwarding table. One route has a label value of 0, and the second has a label value of 1. (For information about these labels, see “Special Labels” on page 25.)

Configuring the Ingress Router for MPLS-Signaled LSPs

MPLS-signaled label-switched paths (LSPs) run from a specific ingress router to a specific egress router. For basic MPLS-signaled LSP function, you must configure the ingress router, but do not have to configure any other routers.

To configure signaled LSPs, perform the following tasks on the ingress router:

- Creating Named Paths on page 52
- Configuring Alternate Backup Paths Using Fate Sharing on page 54

Creating Named Paths

To configure signaled LSPs, you must first create one or more named paths on the ingress router. For each path, you can specify some or all transit routers in the path, or you can leave it empty.

Each pathname can contain up to 32 characters and can include letters, digits, periods, and hyphens. The name must be unique within the ingress router. Once a named path is created, you can use the named path with the **primary** or **secondary** statement to configure LSPs at the [edit protocols mpls label-switched-path *label-path-name*] hierarchy level. You can specify the same named path on any number of LSPs.

To determine whether an LSP is associated with the primary or secondary path in an RSVP session, issue the **show rsvp session detail** command. For more information, see the *JUNOS Routing Protocols and Policies Command Reference*.

To create an empty path, create a named path by including the following form of the **path** statement. This form of the **path** statement is empty, which means that any

path between the ingress and egress routers is accepted. In actuality, the path used tends to be the same path as is followed by destination-based, best-effort traffic.

`path path-name;`

You can include this statement at the following hierarchy levels:

- `[edit protocols mpls]`
- `[edit logical-systems logical-system-name protocols mpls]`

To create a path in which you specify some or all transit routers in the path, include the following form of the `path` statement, specifying one address for each transit router:

```
path path-name {
  (address | hostname) <strict | loose>;
}
```

You can include this statement at the following hierarchy levels:

- `[edit protocols mpls]`
- `[edit logical-systems logical-system-name protocols mpls]`

In this form of the `path` statement, you specify one or more transit router addresses. Specifying the ingress and/or egress routers is optional. You can specify the address or hostname of each transit router, although you do not need to list each transit router if its type is `loose`. Specify the addresses in order, starting with the ingress router (optional) or the first transit router, and continuing sequentially along the path up to the egress router (optional) or the router immediately before the egress router. You need to specify only one address per router hop. If you specify more than one address for the same router, only the first address is used; the additional addresses are ignored and truncated.

For each router address, you specify the type, which can be one of the following:

- **strict**—(Default) The route taken from the previous router to this router is a direct path and cannot include any other routers. If **address** is an interface address, this router also ensures that the incoming interface is the one specified. Ensuring that the incoming interface is the one specified is important when there are parallel links between the previous router and this router. It also ensures that routing can be enforced on a per-link basis.

For strict addresses, you must ensure that the router immediately preceding the router you are configuring has a direct connection to that router. The address can be a loopback interface address, in which case the incoming interface is not checked.

- **loose**—The route taken from the previous router to this router need not be a direct path, can include other routers, and can be received on any interface. The address can be any interface address or the address of the loopback interface.

Examples: Creating Named Paths

Configure a path, **to-hastings**, to specify the complete strict path from the ingress to the egress routers through **14.1.1.1**, **13.1.1.1**, **12.1.1.1**, and **11.1.1.1**, in that order. There cannot be any intermediate routers except the ones specified. However, there can be intermediate routers between **11.1.1.1** and the egress router because the egress router is not specifically listed in the **path** statement. To prevent intermediate routers before egress, configure the egress router as the last router, with a **strict** type.

```
[edit protocols mpls]
path to-hastings {
  14.1.1.1 strict;
  13.1.1.1 strict;
  12.1.1.1 strict;
  11.1.1.1 strict;
}
```

Create a path, **alt-hastings**, to allow any number of intermediate routers between routers **14.1.1.1** and **11.1.1.1**. In addition, intermediate routers are permitted between **11.1.1.1** and the egress router.

```
[edit protocols mpls]
path alt-hastings {
  14.1.1.1 strict;
  11.1.1.1 loose;
}
```

Configuring Alternate Backup Paths Using Fate Sharing

You can create a database of information that Constrained Shortest Path First (CSPF) uses to compute one or more backup paths in case the primary path becomes unstable. The database describes the relationships between elements of the network, such as routers and links. Because these network elements share the same fate, this relationship is called fate sharing.

You can configure backup paths that minimize the number of shared links and fiber paths with the primary paths as much as possible to ensure that, if a fiber is cut, the minimum amount of data is lost and a path still exists to the destination.

For a backup path to work optimally, it must not share links or physical fiber paths with the primary path. This ensures that a single point of failure will not affect the primary and backup paths at the same time.

The following sections describe how to configure fate sharing and how it affects CSPF, and provides a fate sharing configuration example:

- Configuring Fate Sharing on page 55
- Implications for CSPF on page 56
- Example: Configuring Fate Sharing on page 56

Configuring Fate Sharing

To configure fate sharing, include the `fate-sharing` statement:

```
fate-sharing {
  group group-name {
    cost value;
    from address <to address>;
  }
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Each fate-sharing group must have a name, which can be up to 32 characters long and can contain letters, digits, periods (.) and hyphens (-). You can define up to 512 groups.

Fate-sharing groups contain three types of objects:

- Point-to-point links—Identified by the IP addresses at each end of the link. Unnumbered point-to-point links are typically identified by borrowing IP addresses from other interfaces. Order is not important; **from 1.2.3.4 to 1.2.3.5** and **from 1.2.3.5 to 1.2.3.4** have the same meaning.
- Non-point-to-point links—Include links on a LAN interface (such as Gigabit Ethernet interfaces) or nonbroadcast multiaccess (NBMA) interfaces (such as Asynchronous Transfer Mode [ATM] or Frame Relay). You identify these links by their individual interface address. For example, if the LAN interface **192.168.200.0/24** has four routers attached to it, each router link is individually identified:

```
from 192.168.200.1; # LAN interface of router 1
from 192.168.200.2; # LAN interface of router 2
from 192.168.200.3; # LAN interface of router 3
from 192.168.200.4; # LAN interface of router 4
```

You can list the addresses in any order.

- A router node—Identified by its configured router ID.

All objects in a group share certain similarities. For example, you can define a group for all fibers that share the same fiber conduit, all optical channels that share the same fiber, all links that connect to the same LAN switch, all equipment that shares the same power source, and so on. All objects are treated as /32 host addresses.

For a group to be meaningful, it should contain at least two objects. You can configure groups with zero or one object; these groups are ignored during processing.

An object can be in any number of groups, and a group can contain any number of objects. Each group has a configurable cost attributed to it, which represents the level of impact this group has on CSPF computations. The higher the cost, the less likely a backup path will share with the primary path any objects in the group. The cost is directly comparable to traffic engineering metrics. By default, the cost is 1.

Changing the fate-sharing database does not affect established LSPs until the next reoptimization of CSPF. The fate-sharing database does influence fast-reroute computations.

Implications for CSPF

When CSPF computes the primary paths of an LSP (or secondary paths when the primary path is not active), it ignores the fate-sharing information. You always want to find the best possible path (least IGP cost) for the primary path.

When CSPF computes a secondary path while the primary path (of the same LSP) is active, the following occurs:

1. CSPF identifies all fate-sharing groups that are associated with the primary path. CSPF does this by identifying all links and nodes that the primary path traverses and compiling group lists that contain at least one of the links or nodes. CSPF ignores the ingress and egress nodes in the search.
2. CSPF checks each link in the traffic engineering database against the compiled group list. If the link is a member of a group, the cost of the link is increased by the cost of the group. If a link is a member of multiple groups, all group costs are added together.
3. CSPF performs the check for every node in the traffic engineering database, except the ingress and egress node. Again, a node can belong to multiple groups, so costs are additive.
4. The router performs regular CSPF computation with the adjusted topology.

Example: Configuring Fate Sharing

Configure fate-sharing groups `east` and `west`. Because `west` has no objects, it is ignored during processing.

```
[edit routing-options]
fate-sharing {
  group east {
    cost 20; # Optional, default value is 1
    from 1.2.3.4 to 1.2.3.5; # A point-to-point link
    from 192.168.200.1; # LAN interface
    from 192.168.200.2; # LAN interface
    from 192.168.200.3; # LAN interface
    from 192.168.200.4; # LAN interface
    from 10.168.1.220; # Router ID of a router node
    from 10.168.1.221; # Router ID of a router node
  }
  group west {
    ....
  }
}
```

Examples: Configuring Ingress Routers for MPLS-Signaled LSPs

The following examples illustrate how to configure MPLS-signaled LSPs on the ingress router:

Example: Configuring a Constrained-Path LSP for Which JUNOS Makes All Forwarding Decisions

On the ingress router, create a constrained-path LSP in which the JUNOS Software makes all the forwarding decisions. When the LSP is successfully set up, a route toward 10.1.1.1/32 is installed in the `inet.3` table so that all BGP routes with matching BGP next-hop addresses can be forwarded through the LSP.

```
[edit]
interfaces {
  so-0/0/0 {
    unit 0 {
      family mpls;
    }
  }
}
protocols {
  rsvp {
    interface so-0/0/0;
  }
  mpls {
    label-switched-path to-hastings {
      to 10.1.1.1;
    }
    interface so-0/0/0;
  }
}
```

Example: Configuring an Explicit-Path LSP

On the ingress router, create an explicit-path LSP, and specify the transit routers between the ingress and egress routers. In this configuration, no constrained-path computation is performed. For the primary path, all intermediate hops are strictly specified so that its route cannot change. The secondary path must travel through router 14.1.1.1 first, then take whatever route is available to reach the destination. The remaining route taken by the secondary path is typically the shortest path computed by the IGP.

```
[edit]
interfaces {
  so-0/0/0 {
    unit 0 {
      family mpls;
    }
  }
}
```

```

protocols {
  rsvp {
    interface so-0/0/0;
  }
  mpls {
    path to-hastings {
      14.1.1.1 strict;
      13.1.1.1 strict;
      12.1.1.1 strict;
      11.1.1.1 strict;
    }
    path alt-hastings {
      14.1.1.1 strict;
      11.1.1.1 loose; # Any IGP route is acceptable
    }
    label-switched-path hastings {
      to 11.1.1.1;
      hop-limit 32;
      bandwidth 10m; # Reserve 10 Mbps
      no-cspf; # do not perform constrained-path computation
      primary to-hastings;
      secondary alt-hastings;
    }
    interface so-0/0/0;
  }
}

```

Example: Configuring a Constrained-Path LSP for Which JUNOS Makes Most Forwarding Decisions and Considers Hop Constraints

On the ingress router, create a constrained-path LSP in which the JUNOS Software makes most of the forwarding decisions, taking into account the hop constraints listed in the `path` statements. The LSP is adaptive so that no bandwidth double-counting occurs on links shared by primary and secondary paths. To acquire the necessary link bandwidth, this LSP is allowed to preempt lower priority sessions. Finally, this path always keeps the secondary path in hot-standby state for quick failover.

```

[edit protocols]
mpls {
  path to-hastings {
    14.1.1.1 loose;
  }
  path alt-hastings {
    12.1.1.1 loose;
    11.1.1.1 strict;
  }
  label-switched-path hastings {
    to 11.1.1.1;
    bandwidth 10m; # Reserve 10 Mbps
    priority 0 0; # Preemptive, but not preemptable
    adaptive; # Set adaptivity
    primary to-hastings;
    secondary alt-hastings {

```



```

        standby;
        bandwidth 1m; # Reserve only 1 Mbps for the secondary path
    }
}
interface all;
}

```

Example: Configuring a Constrained-Path LSP for Which JUNOS Makes Most Forwarding Decisions and the Secondary Path Is Explicit

On the ingress router, create a constrained-path LSP in which the JUNOS Software makes most of the forwarding decisions for the primary path, subject to constraints of the path to-hastings, and in which the secondary path is an explicit path. The primary path must transit green or yellow links and must stay away from red links. The primary path is periodically recomputed and reoptimized. Finally, this path always keeps the secondary path in hot-standby state for quick failover.

When the LSP is up—either because the primary or secondary path is up, or because both paths are up—the prefix 16.0.0.0/8 is installed in the inet.3 table so that all BGP routes whose BGP next hop falls within that range can use the LSP. Also, the prefix 17/8 is installed in the inet.0 table so that BGP can resolve only its next hop through that prefix. The route also can be reached with `traceroute` or `ping`. These two routes are in addition to the 11.1.1.1/32 route.

```

[edit protocols]
mpls {
  admin-groups {
    green 1;
    yellow 2;
    red 3;
  }
  path to-hastings {
    14.1.1.1 loose;
  }
  path alt-hastings {
    14.1.1.1 strict;
    13.1.1.1 strict;
    12.1.1.1 strict;
    11.1.1.1 strict;
  }
  label-switched-path hastings {
    to 11.1.1.1;
    bandwidth 100m;
    install 16.0.0.0/8; # in inet.3; cannot use to traceroute or ping
    install 17.0.0.0/8 active; # installed in inet.0; can use to traceroute or ping
    primary to-hastings {
      admin-group { # further constraints for path computation
        include-all [ green yellow ];
        exclude red;
      }
      optimize-timer 3600; # reoptimize every hour
    }
    secondary alt-hastings {
      standby;
    }
  }
}

```

```

        no-cspf; # do not perform constrained-path computation
    }
}
interface all;

```

Configuring the Intermediate and Egress Routers for MPLS-Signaled LSPs

To configure signaled LSPs on all MPLS routers that should participate in MPLS, you need to enable MPLS and RSVP on these routers, as described in “Minimum MPLS Configuration” on page 51 and “RSVP Configuration Guidelines” on page 273.

Configuring a Default Route for Unlabeled MPLS Packets

You can process MPLS packets that have not been assigned label values and have no corresponding entry in the `mpls.0` routing table. This allows you to assign a default route to unlabeled MPLS packets. The process is also referred to as *MPLS exception monitoring*.

To configure a default label value for MPLS packets, include the `default-route` statement:

```

default-route {
  class-of-service cos-value;
  (next-hop (address | interface-name | address/interface-name) |
  (discard | reject);
  (pop | (swap out-label);
  preference preference;
  swap-push swap-label push-label;
}

```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls interface *interface-name* label-map]
- [edit logical-systems *logical-system-name* protocols mpls interface *interface-name* label-map]

Improving Traffic Engineering Database Accuracy with RSVP PathErr Messages

An essential element of RSVP-based traffic engineering is the traffic engineering database. The traffic engineering database contains a complete list of all network nodes and links participating in traffic engineering, and a set of attributes each of those links can hold. (For more information about the traffic engineering database, see “Constrained-Path LSP Computation” on page 29.) One of the most important link attributes is bandwidth.

Bandwidth availability on links changes quickly as RSVP LSPs are established and terminated. It is likely that the traffic engineering database will develop inconsistencies relative to the real network. These inconsistencies cannot be fixed by increasing the rate of IGP updates.

Link availability can share the same inconsistency problem. A link that becomes unavailable can break all existing RSVP LSPs. However, its unavailability might not readily be known by the network.

When you configure the `rsvp-error-hold-time` statement, a source node (ingress of an RSVP LSP) learns from the failures of its LSP by monitoring PathErr messages transmitted from downstream nodes. Information from the PathErr messages is incorporated into subsequent LSP computations, which can improve the accuracy and speed of LSP setup. Some PathErr messages are also used to update traffic engineering database bandwidth information, reducing inconsistencies between the traffic engineering database and the network.

You can control the frequency of IGP updates by using the `update-threshold` statement. See “Configuring the RSVP Update Threshold on an Interface” on page 279.

This section discusses the following topics:

- PathErr Messages on page 61
- Identifying the Problem Link on page 62
- Configuring the Router to Improve Traffic Engineering Database Accuracy on page 62

PathErr Messages

PathErr messages report a wide variety of problems by means of different code and subcode numbers. You can find a complete list of these PathErr messages in RFC 2205, *Resource Reservation Protocol (RSVP), Version 1, Functional Specification* and RFC 3209, *RSVP-TE: Extensions to RSVP for LSP Tunnels*.

When you configure the `rsvp-error-hold-time` statement, two categories of PathErr messages, which specifically represent link failures, are examined:

- Link bandwidth is low for this LSP: Requested bandwidth unavailable—code 1, subcode 2

This type of PathErr message represents a global problem that affects all LSPs transiting the link. They indicate that the actual link bandwidth is lower than that required by the LSP, and that it is likely that the bandwidth information in the traffic engineering database is an overestimate.

When this type of error is received, the available link bandwidth is reduced in the local traffic engineering database, affecting all future LSP computations.

- Link unavailable for this LSP:
 - Admission Control failure—code 1, any subcode except 2
 - Policy Control failures—code 2
 - Service Preempted—code 12
 - Routing problem—no route available toward destination—code 24, subcode 5

These types of PathErr messages are generally pertinent to the specified LSP. The failure of this LSP does not necessarily imply that other LSPs could also fail. These errors can indicate maximum transfer unit (MTU) problems, service preemption (either manually initiated by the operator or by another LSP with a higher priority), that a next-hop link is down, that a next-hop neighbor is down, or service rejection because of policy considerations. It is best to route this particular LSP away from the link.

Identifying the Problem Link

Each PathErr message includes the sender's IP address. This information is propagated unchanged toward the ingress router. A lookup in the traffic engineering database can identify the node that originated the PathErr message.

Each PathErr message carries enough information to identify the RSVP session that triggered the message. If this is a transit router, it simply forwards the message. If this router is the ingress router (for this RSVP session), it has the complete list of all nodes and links the session should traverse. Coupled with the originating node information, the link can be uniquely identified.

Configuring the Router to Improve Traffic Engineering Database Accuracy

To improve the accuracy of the traffic engineering database, configure the `rsvp-error-hold-time` statement. When this statement is configured, a source node (ingress of an RSVP LSP) learns from the failures of its LSP by monitoring PathErr messages transmitted from downstream nodes. Information from the PathErr messages is incorporated into subsequent LSP computations, which can improve the accuracy and speed of LSP setup. Some PathErr messages also are used to update traffic engineering database bandwidth information, reducing inconsistencies between the traffic engineering database and the network.

To configure how long MPLS should remember RSVP PathErr messages and consider them in CSPF computation, include the `rsvp-error-hold-time` statement:

```
rsvp-error-hold-time seconds;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls]
- [edit logical-systems *logical-system-name* protocols mpls]

The time can be a value from 1 to 240 seconds. The default is 25 seconds. Configuring a value of 0 disables the monitoring of PathErr messages.

Configuring MPLS-Signaled LSPs to Use GRE Tunnels

MPLS LSPs can use generic routing encapsulation (GRE) tunnels to cross routing areas, autonomous systems, and ISPs. Bridging MPLS LSPs over an intervening IP domain is possible without disrupting the outlying MPLS domain.

LSPs can reach any destination that the GRE tunnels can reach. MPLS applications can be deployed without requiring all transit nodes to support MPLS, or requiring all transit nodes to support the same label distribution protocols (LDP or RSVP). If you use CSPF, you must configure OSPF or IS-IS through the GRE tunnel. Traffic engineering is not supported over GRE tunnels; for example, you cannot reserve bandwidth or set priority or preemption.



NOTE: Use the **no-control word** statement to disable the control word when the topology uses GRE as the connection mechanism between provider edge routers and one of the provider edge routers is an M Series Multiservice Edge Router.

For more information about GRE tunnels, see the *JUNOS Services Interfaces Configuration Guide*.

Example: Configuring MPLS-Signaled LSPs to Use GRE Tunnels

To configure MPLS over GRE tunnels:

1. Enable family mpls under the GRE interface configuration:

```
[edit interfaces]
interface gr-1/2/0 {
  unit 0 {
    tunnel {
      source 192.168.1.1;
      destination 192.168.1.2;
    }
    family inet {
      address 5.1.1.1/30;
    }
    family iso;
    family mpls;
  }
}
```

2. Enable RSVP and MPLS over the GRE tunnel:

```
[edit protocols]
rsvp {
  interface gr-1/2/0.0;
}
mpls {
  ...
  interface gr-1/2/0.0;
}
```

3. Configure LSPs to travel through the GRE tunnel endpoint address:

```
[edit protocols]
mpls {
  label-switched-path gre-tunnel {
    to 5.1.1.2;
    ...
  }
}
```

```

    }
}

```


Standard LSP configuration options apply. If the routing table specifies that a particular route will traverse a GRE tunnel, the RSVP packets will traverse the tunnel as well.

Tunneling IPv6 Traffic over MPLS IPv4 Networks

You can configure the JUNOS Software to tunnel IPv6 over an MPLS-based IPv4 network. This configuration allows you to interconnect a number of smaller IPv6 networks over an IPv4-based network core, giving you the ability to provide IPv6 service without having to upgrade the routers in your core network. Multiprotocol Border Gateway Protocol (MP-BGP) is configured to exchange routes between the IPv6 networks, and data is tunneled between these IPv6 networks by means of IPv4-based MPLS.

In Figure 19 on page 64, Routers PE1 and PE2 are dual-stack BGP routers, meaning they have both IPv4 and IPv6 stacks. The PE routers link the IPv6 networks through the customer edge (CE) routers to the IPv4 core network. The CE routers and the PE routers connect through a link layer that can carry IPv6 traffic. The PE routers use IPv6 on the CE router-facing interfaces and use IPv4 and MPLS on the core-facing interfaces. Note that one of the connected IPv6 networks could be the global IPv6 Internet.

Figure 19: IPv6 Networks Linked by MPLS IPv4 Tunnels



The two PE routers are linked through a MP-BGP session using IPv4 addresses. They use the session to exchange IPv6 routes with an IPv6 (value 2) address family indicator (AFI) and a subsequent AFI (SAFI) (value 4). Each PE router sets the next hop for the IPv6 routes advertised on this session to its own IPv4 address. Because MP-BGP requires the BGP next hop to correspond to the same address family as the network

layer reachability information (NLRI), this IPv4 address needs to be embedded within an IPv6 format.

The PE routers can learn the IPv6 routes from the CE routers connected to them via routing protocols Routing Information Protocol next generation (RIPng) or MP-BGP, or through static configuration. Note that if BGP is used as the PE-router-to-CE-router protocol, the MP-BGP session between the PE router and CE router could occur over an IPv4 or IPv6 Transmission Control Protocol (TCP) session. Also, the BGP routes exchanged on that session would have SAFI unicast. You must configure an export policy to pass routes between IBGP and EBGP, and between BGP and any other protocol.

The PE routers have MPLS LSPs routed to each others' IPv4 addresses. IPv4 provides signaling for the LSPs by means of either LDP or RSVP. These LSPs are used to resolve the next-hop addresses of the IPv6 routes learned from MP-BGP. The next hops use IPv4-mapped IPv6 addresses, while the LSPs use IPv4 addresses.

The PE routers always advertise IPv6 routes to each other using a label value of 2, the explicit null label for IPv6 as defined in RFC 3032, *MPLS Label Stack Encoding*. As a consequence, each of the forwarding next hops for the IPv6 routes learned from remote PE routers normally push two labels. The inner label is 2 (this label could be different if the advertising PE router is not a Juniper Networks routing platform), and the outer label is the LSP label. If the LSP is a single-hop LSP, then only label 2 is pushed.

It is also possible for the PE routers to exchange plain IPv6 routes using SAFI unicast. However, there is one major advantage in exchanging labeled IPv6 routes. The penultimate-hop router for an MPLS LSP can pop the outer label and then send the packet with the inner label as an MPLS packet. Without the inner label, the penultimate-hop router would need to discover whether the packet is an IPv4 or IPv6 packet to set the protocol field in the Layer 2 header correctly.

When the PE1 router in Figure 19 on page 64 receives an IPv6 packet from the CE1 router, it performs a lookup in the IPv6 forwarding table. If the destination matches a prefix learned from the CE2 router, then no labels need to be pushed and the packet is simply sent to the CE2 router. If the destination matches a prefix that was learned from the PE2 router, then the PE1 router pushes two labels onto the packet and sends it to the provider router. The inner label is 2 and the outer label is the LSP label for the PE2 router.

Each provider router in the service provider's network handles the packet as it would any MPLS packet, swapping labels as it passes from provider router to provider router. The penultimate-hop provider router for the LSP pops the outer label and sends the packet to the PE2 router. When the PE2 router receives the packet, it recognizes the IPv6 explicit null label on the packet (Label 2). It pops this label and treats it as an IPv6 packet, performing a lookup in the IPv6 forwarding table and forwarding the packet to the CE3 router.

This section discusses the following topics:

- IPv6 over MPLS Standards on page 66
- Configuring IPv4 MPLS Tunnels to Carry IPv6 Traffic on page 66

IPv6 over MPLS Standards

Detailed information about the Juniper Networks implementation of IPv6 over MPLS is described in the following Internet drafts:

- Internet draft draft-ietf-l3vpn-bgp-ipv6-07.txt, *BGP-MPLS IP VPN extension for IPv6 VPN* (expires January 2006)
- Internet draft draft-ooms-v6ops-bgp-tunnel-06.txt, *Connecting IPv6 Islands over IPv4 MPLS using IPv6 Provider Edge Routers* (expires July 2006)

These Internet drafts are available on the IETF website at <http://www.ietf.org/>.

Configuring IPv4 MPLS Tunnels to Carry IPv6 Traffic

You must perform the following tasks to allow IPv6 to be carried over an IPv4 MPLS tunnel:

- Configuring IPv6 on Both Core-Facing and CE Router-Facing Interfaces on page 66
- Configuring MPLS and RSVP Between PE Routers on page 66
- Enabling IPv6 Tunneling on PE Routers on page 67
- Configuring Multiprotocol BGP to Carry IPv6 Traffic on page 67

Configuring IPv6 on Both Core-Facing and CE Router-Facing Interfaces

In addition to configuring the `family inet6` statement on all the CE router-facing interfaces, you must also configure the statement on all the core-facing interfaces running MPLS. Both configurations are necessary because the router must be able to process any IPv6 packets it receives on these interfaces. You should not see any regular IPv6 traffic arrive on these interfaces, but you will receive MPLS packets tagged with label 2. Even though label 2 MPLS packets are sent in IPv4, these packets are treated as native IPv6 packets.

Include the `family inet6` statement:

```
family inet6 {
    address inet6-address;
}
```

You can include this statement at the following hierarchy levels:

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

Configuring MPLS and RSVP Between PE Routers

For information about how to configure MPLS and RSVP, see the following sections:

- Configuring the Ingress Router for MPLS-Signaled LSPs on page 52

- Configuring the Intermediate and Egress Routers for MPLS-Signaled LSPs on page 60
- RSVP Configuration Guidelines on page 273

Enabling IPv6 Tunneling on PE Routers

You enable IPv6 tunneling by including the `ipv6-tunneling` statement in the configuration for the PE routers. This statement allows IPv6 routes to be resolved over an MPLS network by converting all routes stored in the `inet.3` routing table to IPv4-compatible IPv6 addresses and then copying them into the `inet6.3` routing table. This routing table can be used to resolve next hops for both `inet6` and `inet6-vpn` routes.

To configure IPv6 tunneling, include the `ipv6-tunneling` statement on the PE routers:

```
ipv6-tunneling;
```

You can include this statement at the following hierarchy levels:

- `[edit protocols mpls]`
- `[edit logical-systems logical-system-name protocols mpls]`

You also need to configure IPv6 tunneling when you configure IPv6 VPNs. For more information, see the *JUNOS VPNs Configuration Guide*.

Configuring Multiprotocol BGP to Carry IPv6 Traffic

When you configure MP-BGP to carry IPv6 traffic, the IPv4 MPLS label is removed at the destination PE router. The remaining IPv6 packet without a label can then be forwarded to the IPv6 network. Include the `explicit-null` statement:

```
explicit-null;
```

You can include this statement at the following hierarchy levels:

- `[edit protocols bgp family inet6 labeled-unicast]`
- `[edit protocols bgp group group-name family inet6 labeled-unicast]`
- `[edit protocols bgp group group-name neighbor neighbor-name family inet6 labeled-unicast]`
- `[edit logical-systems logical-system-name protocols bgp family inet6 labeled-unicast]`
- `[edit logical-systems logical-system-name protocols bgp group group-name family inet6 labeled-unicast]`
- `[edit logical-systems logical-system-name protocols bgp group group-name neighbor neighbor-name family inet6 labeled-unicast]`

Configuring ICMP Message Tunneling

When you configure MPLS to tunnel through a routing domain, it is difficult to route a fragmented packet to its source address; for example, when the IP addresses carried

in a packet are private (not globally unique) and MPLS is used to tunnel the packets through a public backbone.

When you configure Internet Control Message Protocol (ICMP) message tunneling, an ICMP message is sent to the source of a packet. The label stack is copied from the original packet to the ICMP message. The ICMP message is then label switched across the network. This causes the message to go to the original packet destination, rather than its source. Unless the message is label switched all the way to the destination host, it ends up unlabeled in a router that does know the source of the original packet, at which point the message is sent in the proper direction.

ICMP message tunneling can be useful for debugging and tracing purposes if the message is an ICMP time exceeded message.

To configure ICMP message tunneling, include the `icmp-tunneling` statement:

```
icmp-tunneling;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls]
- [edit logical-systems *logical-system-name* protocols mpls]

Chapter 5

MPLS-Signaled LSP Configuration Guidelines

MPLS-signaled label-switched paths (LSPs) run from a specific ingress router to a specific egress router. This chapter describes how to configure LSPs. You can configure an LSP so that the JUNOS Software makes all forwarding decisions, or you can configure some or all routers in the path.

For instructions for configuring the ingress router and other routers on the path, see “MPLS Router Configuration Guidelines” on page 51.

- LSP Configuration Overview on page 70
- Configuring the Ingress and Egress Router Addresses for LSPs on page 73
- Configuring Primary and Secondary LSPs on page 75
- Configuring a Text Description for LSPs on page 78
- Configuring Fast Reroute on page 78
- Configuring the Optimization Interval for Fast Reroute Paths on page 79
- Adding LSP-Related Routes to the inet.3 Routing Table on page 80
- Configuring the Connection Between Ingress and Egress Routers on page 81
- Configuring LSP Metrics on page 82
- Configuring CSPF Tie Breaking on page 83
- Configuring Load Balancing for MPLS LSPs on page 84
- Disabling Normal TTL Decrementing on page 86
- Configuring MPLS Soft Preemption on page 87
- Configuring Automatic Bandwidth Allocation for LSPs on page 88
- Disabling Constrained-Path LSP Computation on page 95
- Configuring Administrative Groups on page 96
- Configuring Preference Values for LSPs on page 98
- Disabling Path Route Recording on page 98
- Configuring Class of Service for MPLS LSPs on page 98
- Configuring Adaptive LSPs on page 101
- Configuring Priority and Preemption for LSPs on page 102
- Optimizing Signaled LSPs on page 103

- Configuring the Smart Optimize Timer on page 105
- Limiting the Number of Hops in LSPs on page 105
- Configuring the Bandwidth Value for LSPs on page 106
- Configuring Hot Standby of Secondary Paths on page 106
- Damping Advertisement of LSP State Changes on page 107

LSP Configuration Overview

To configure an MPLS-signaled LSP, you define the properties associated with the LSP on the ingress router. Include the `label-switched-path` statement:

```
label-switched-path lsp-name {
  disable;
  adaptive;
  admin-down;
  admin-group {
    exclude [ group-names ];
    include-all [ group-names ];
    include-any [ group-names ];
  }
  auto-bandwidth {
    adjust-interval seconds;
    adjust-threshold percent;
    maximum-bandwidth bps;
    minimum-bandwidth bps;
    monitor-bandwidth;
  }
  bandwidth bps {
    ct0 bps;
    ct1 bps;
    ct2 bps;
    ct3 bps;
  }
  class-of-service cos-value;
  description text;
  fast-reroute {
    (bandwidth bps | bandwidth-percent percent);
    (exclude [ group-names ] | no-exclude);
    hop-limit number;
    (include-all [ group-names ] | no-include-all);
    (include-any [ group-names ] | no-include-any);
  }
  from address;
  hop-limit number;
  install {
    destination-prefix/prefix-length <active>;
  }
  ldp-tunneling;
  link-protection;
  lsp-attributes {
    encoding-type (ethernet | packet | pdh | sonet-sdh);
    gpid (ethernet | hdlc | ipv4 | ppp);
  }
}
```

```

    signal-bandwidth type;
    switching-type (fiber | lambda | psc-1 | tdm);
}
metric number;
no-cspf;
no-decrement-ttl;
node-link-protection;
optimize-timer seconds;
p2mp path-name;
policing {
    filter filter-name;
    no-automatic-policing;
}
preference preference;
primary path-name {
    adaptive;
    admin-group {
        exclude [ group-names ];
        include-all [ group-names ];
        include-any [ group-names ];
    }
    bandwidth bps {
        ct0 bps;
        ct1 bps;
        ct2 bps;
        ct3 bps;
    }
    class-of-service cos-value;
    hop-limit number;
    no-cspf;
    no-decrement-ttl;
    optimize-timer seconds;
    preference preference;
    priority setup-priority reservation-priority;
    (record | no-record);
    select (manual | unconditional);
    standby;
}
priority setup-priority reservation-priority;
(random | least-fill | most-fill);
(record | no-record);
retry-limit number;
retry-timer seconds;
revert-timer seconds;
secondary path-name {
    adaptive;
    admin-group {
        exclude [ group-names ];
        include-all [ group-names ];
        include-any [ group-names ];
    }
    bandwidth bps {
        ct0 bps;
        ct1 bps;
        ct2 bps;
        ct3 bps;
    }
}

```

```

    }
    class-of-service cos-value;
    hop-limit number;
    no-cspf;
    no-decrement-ttl;
    optimize-timer seconds;
    preference preference;
    priority setup-priority reservation-priority;
    (record | no-record);
    select (manual | unconditional);
    standby;
  }
  soft-preemption {
    cleanup-timer seconds;
  }
  standby;
  to address;
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
}

```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls]
- [edit logical-systems *logical-system-name* protocols mpls]

Each LSP must have a name, *lsp-name*, which can be up to 32 characters long and can contain letters, digits, periods (.), and hyphens (-). The name must be unique within the ingress router. For ease of management and identification, configure unique names across the entire domain.

When you configure LSPs, you can specify the following statements either for each LSP or for each path. For statements that you configure on a per-LSP basis, the value applies to all paths in the LSP. For statements that you configure on a per-path basis, the path value overrides the per-LSP value.

- adaptive
- admin-group
- auto-bandwidth
- bandwidth
- class-of-service
- hop-limit
- no-cspf
- optimize-timer
- preference
- priority

- record or no-record
- standby

For maintenance purposes, you can also configure the following attributes across all LSPs and any paths within those LSPs:

- admin-group
- bandwidth
- class-of-service
- no-decrement-ttl
- no-record
- optimize-timer
- preference
- priority
- smart-optimize-timer
- standby

Configuring the Ingress and Egress Router Addresses for LSPs

The following sections describe how to specify the addresses of an LSP's ingress and egress routers:

- Configuring the Ingress Router Address for LSPs on page 73
- Configuring the Egress Router Address for LSPs on page 74
- Preventing the Addition of Egress Router Addresses to Routing Tables on page 74

Configuring the Ingress Router Address for LSPs

The local router always is considered to be the ingress router, which is the beginning of the LSP. The software automatically determines the proper outgoing interface and IP address to use to reach the next router in an LSP.

By default, the router ID is chosen as the address of the ingress router. To override the automatic selection of the source address, specify a source address in the **from** statement:

```
from address;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls label-switched-path *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*]

The outgoing interface used by the LSP is not affected by the source address that you configure.

Configuring the Egress Router Address for LSPs

When configuring an LSP, you must specify the address of the egress router by including the **to** statement:

```
to address;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls label-switched-path *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*]

When you are setting up an LSP, the **to** statement is the only required statement. All other statements are optional.

After the LSP is established, the address of the egress router is installed as a host route in the routing table. This route can then be used by BGP to forward traffic.

To have the software send BGP traffic over an LSP, the address of the egress router is the same as the address of the BGP next hop. You can specify the egress router's address as any one of the router's interface addresses or as the BGP router ID. If you specify a different address, even if the address is on the same router, BGP traffic is not sent over the LSP.

To determine the address of the BGP next hop, use the **show route detail** command. To determine the destination address of an LSP, use the **show mpls lsp** command. To determine whether a route has gone through an LSP, use the **show route** or **show route forwarding-table** command. In the output of these last two commands, the **label-switched-path** or **push** keyword included with the route indicates it has passed through an LSP. Also, use the **traceroute** command to trace the actual path to which the route leads. This is another indication as to whether a route has passed through an LSP.

You also can manipulate the address of the BGP next hop by defining a BGP import policy filter that sets the route's next-hop address.

Preventing the Addition of Egress Router Addresses to Routing Tables

You must configure an address using the **to** statement for all RSVP LSPs. This address is always installed as a /32 prefix in the **inet.3** or **inet.0** routing tables. You can prevent the egress router address configured using the **to** statement from being added to the **inet.3** and **inet.0** routing tables by including the **no-install-to-address** statement.

Some reasons not to install the **to** statement address in the **inet.3** and **inet.0** routing tables include the following:

- Allow Constrained Shortest Path First (CSPF) RSVP LSPs to be mapped to traffic intended for secondary loopback addresses. If you configure an RSVP tunnel, including the **no-install-to-address** statement, and then configure an **install pfx/<active>** policy later, you can do the following:
 - Verify that the LSP was set up correctly without impacting traffic.

- Map traffic to the LSP in incremental steps.
- Map traffic to the destination loopback address (the BGP next hop) by removing the `no-install-to-address` statement once troubleshooting is complete.
- Prevent CCC connections from losing IP traffic. When an LSP determines that it does not belong to a connection, it installs the address specified with the `to` statement in the `inet.3` routing table. IP traffic is then forwarded to the CCC remote endpoint, which can cause some types of PICs to fail.

To prevent the egress router address configured using the `to` statement from being added to the `inet.3` and `inet.0` routing tables, include the `no-install-to-address` statement:

```
no-install-to-address;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls label-switched-path *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*]

Configuring Primary and Secondary LSPs

By default, an LSP routes itself hop-by-hop toward the egress router. The LSP tends to follow the shortest path as dictated by the local routing table, usually taking the same path as destination-based, best-effort traffic. These paths are “soft” in nature because they automatically reroute themselves whenever a change occurs in a routing table or in the status of a node or link.

To configure the path so that it follows a particular route, create a named path using the `path` statement, as described in “Creating Named Paths” on page 52. Then apply the named path by including the `primary` or `secondary` statement. A named path can be referenced by any number of LSPs.

To configure primary and secondary paths for an LSP, complete the steps in the following sections:

- Configuring Primary and Secondary Paths for an LSP on page 75
- Configuring the Revert Timer for LSPs on page 76
- Specifying the Conditions for Path Selection on page 77

Configuring Primary and Secondary Paths for an LSP

The `primary` statement creates the primary path, which is the LSP’s preferred path. The `secondary` statement creates an alternative path. If the primary path can no longer reach the egress router, the alternative path is used.

To configure primary and secondary paths, include the `primary` and `secondary` statements:

```
primary path-name {
  ...
```

```

    }
    secondary path-name {
    ...
    }

```

You can include these statements at the following hierarchy levels:

- [edit protocols mpls label-switched-path *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*]

When the software switches from the primary to a secondary path, it continuously attempts to revert to the primary path, switching back to it when it is again reachable, but no sooner than the retry time specified in the **retry-timer** statement. (For more information, see “Configuring the Connection Between Ingress and Egress Routers” on page 81.)

You can configure zero or one primary path. If you do not configure a primary path, the first secondary path that is established is selected as the path.

You can configure zero or more secondary paths. All secondary paths are equal, and the software tries them in the order that they are listed in the configuration. The software does not attempt to switch among secondary paths. If the current secondary path is not available, the next one is tried. To create a set of equal paths, specify secondary paths without specifying a primary path.

If you do not specify any named paths, or if the path that you specify is empty, the software makes all routing decisions necessary to reach the egress router.

Configuring the Revert Timer for LSPs

For LSPs configured with both primary and secondary paths, it is possible to configure the revert timer. If a primary path goes down and traffic is switched to the secondary path, the revert timer specifies the amount of time (in seconds) that the LSP must wait before it can revert traffic back to a primary path. If during this time, the primary path experiences any connectivity problems or stability problems, the timer is restarted.

The JUNOS Software also makes a determination as to which path is the preferred path. The preferred path is the path which has not encountered any difficulty in the last revert timer period. If both the primary and secondary paths have encountered difficulty, neither path is considered preferred. However, if one of the paths is dynamic and the other static, the dynamic path is selected as the preferred path.

The range of values you can configure for the revert timer is 0 through 65,535 seconds. The default value is 60 seconds.

If you configure a value of 0 seconds, the traffic on the LSP, once switched from the primary path to the secondary path, remains on the secondary path permanently (until the network operator intervenes or until the secondary path goes down).

You can configure the revert timer for all LSPs on the router at the [edit protocols mpls] hierarchy level or for a specific LSP at the [edit protocols mpls label-switched-path *lsp-name*] hierarchy level.

To configure the revert timer, include the **revert-timer** statement:

```
revert-timer seconds;
```

For a list of hierarchy levels at which you can include this statement, see the summary section for this statement.

Specifying the Conditions for Path Selection

When you have configured both primary and secondary paths for an LSP, you may need to ensure that only a specific path is used.

The **select** statement is optional. If you do not include it, MPLS uses an automatic path selection algorithm.

The **manual** and **unconditional** options do the following:

- **manual**—The path is immediately selected for carrying traffic as long as it is up and stable. Traffic is sent to other working paths if the current path is down or degraded (receiving errors). This parameter overrides all other path attributes except the **select unconditional** statement.
- **unconditional**—The path is selected for carrying traffic unconditionally, regardless of whether the path is currently down or degraded (receiving errors). This parameter overrides all other path attributes.

Because the **unconditional** option switches to a path without regard to its current status, be aware of the following potential consequences of specifying it:

- If a path is not currently up when you enable the **unconditional** option, traffic can be disrupted. Ensure that the path is functional before specifying the **unconditional** option.
- Once a path is selected because it has the **unconditional** option enabled, all other paths for the LSP are gradually cleared, including the primary and standby paths. No path can act as a standby to an unconditional path, so signaling those paths serves no purpose.

For a specific path, the **manual** and **unconditional** options are mutually exclusive. You can include the **select** statement with the **manual** option in the configuration of only one of an LSP's paths, and the **select** statement with the **unconditional** option in the configuration of only one other of its paths.

Enabling or disabling the **manual** and **unconditional** options for the **select** statement while LSPs and their paths are up does not disrupt traffic.

To specify that a path is selected for carrying traffic if it is up and stable for at least the revert timer window, include the **select** statement with the **manual** option:

```
select manual;
```

To specify that a path is always selected for carrying traffic, even if it is currently down or degraded, include the **select** statement with the **unconditional** option:

```
select unconditional;
```

You can include the **select** statement at the following hierarchy levels:

- [edit protocols mpls label-switched-path *lsp-name* (primary | secondary) *path-name*]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* (primary | secondary) *path-name*]

Configuring a Text Description for LSPs

You can provide a textual description for the LSP. Enclose any descriptive text that includes spaces in quotation marks (" "). Any descriptive text you include is displayed in the output of the **show mpls lsp detail** command and has no effect on the operation of the LSP.

To provide a textual description for the LSP, include the **description** statement:

```
description text;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls label-switched-path *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*]

The description text can be no more than 80 characters in length.

Configuring Fast Reroute

Fast reroute provides a mechanism for automatically rerouting traffic on an LSP if a node or link in an LSP fails, thus reducing the loss of packets traveling over the LSP.

To configure fast reroute on an LSP, include the **fast-reroute** statement on the ingress router:

```
fast-reroute {
  (bandwidth bps | bandwidth-percent percentage);
  (exclude group-names | no-exclude);
  hop-limit number;
  include-all [ group-names ];
  include-any [ group-names ];
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls label-switched-path *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*]

You do not need to configure fast reroute on the LSP's transit and egress routers. Once fast reroute is enabled, the ingress router signals all the downstream routers that fast reroute is enabled on the LSP, and each downstream router does its best to set up detours for the LSP. If a downstream router does not support fast reroute, it

ignores the request to set up detours and continues to support the LSP. A router that does not support fast reroute will cause some of the detours to fail, but otherwise has no impact on the LSP.



NOTE: To enable PFE fast reroute, configure a routing policy statement with the **load-balance per-packet** statement at the [edit policy-options policy-statement *policy-name* then] hierarchy level on each of the routers where traffic might be rerouted. See also “Configuring Load Balancing Across RSVP LSPs” on page 291.

By default, no bandwidth is reserved for the rerouted path. To allocate bandwidth for the rerouted path, include either the **bandwidth** statement or the **bandwidth-percent** statement. You can only include one of these statements at a time. If you do not include either the **bandwidth** statement or the **bandwidth-percent** statement, the default setting is to not reserve bandwidth for the detour path.

When you include the **bandwidth** statement, you can specify the specific amount of bandwidth (in bits per second [bps]) you want to reserve for the detour path. The bandwidth does not need to be identical to that allocated for the LSP.

When you specify a bandwidth percent using the **bandwidth-percent** statement, the detour path bandwidth is computed by multiplying the bandwidth percent by the bandwidth configured for the main traffic-engineered LSP. For information on how to configure the bandwidth for a traffic-engineered LSP, see “Configuring Traffic-Engineered LSPs” on page 127.

Hop-limit constraints define how many more routers a detour is allowed to traverse compared with the LSP itself. By default, the hop limit is set to 6. For example, if an LSP traverses 4 routers, any detour for the LSP can be up to 10 (that is, 4 + 6) router hops, including the ingress and egress routers.

By default, a detour inherits the same administrative (coloring) group constraints as its parent LSP when CSPF is determining the alternate path. Administrative groups, also known as link coloring or resource class, are manually assigned attributes that describe the “color” of links, such that links with the same color conceptually belong to the same class. If you specify the **include-any** statement when configuring the parent LSP, all links traversed by the alternate session must have at least one color found in the list of groups. If you specify the **include-all** statement when configuring the parent LSP, all links traversed by the alternate session must have all of the colors found in the list of groups. If you specify the **exclude** statement when configuring the parent LSP, none of the links must have a color found in the list of groups. For more information about administrative group constraints, see “Configuring Administrative Groups” on page 96.

Configuring the Optimization Interval for Fast Reroute Paths

You can enable path optimization for fast reroute by configuring the fast reroute optimize timer. The optimize timer triggers a periodic optimization process that recomputes the fast reroute detour LSPs to use network resources more efficiently.

To enable fast reroute path optimization, specify the number of seconds including the **fast-reroute optimize-timer** statement:

```
fast-reroute optimize-timer seconds;
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp]
- [edit logical-systems *logical-system-name* protocols rsvp]

Adding LSP-Related Routes to the inet.3 Routing Table

By default, a host route toward the egress router is installed in the **inet.3** routing table. (The host route address is the one you configure in the **to** statement.) Installing the host route allows BGP to perform next-hop resolution. It also prevents the host route from interfering with prefixes learned from dynamic routing protocols and stored in the **inet.0** routing table.

Unlike the routes in the **inet.0** table, routes in the **inet.3** table are not copied to the Packet Forwarding Engine, and hence they cause no changes in the system forwarding table directly. You cannot use the **ping** or **traceroute** command through these routes. The only use for **inet.3** is to permit BGP to perform next-hop resolution. To examine the **inet.3** table, use the **show route table inet.3** command.

To inject additional routes into the **inet.3** routing table, include the **install** statement:

```
install {
  destination-prefix <active>;
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls label-switched-path *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*]

The specified routes are installed as aliases into the routing table when the LSP is established. Installing additional routes allows BGP to resolve next hops within the specified prefix and to direct additional traffic for these next hops to a particular LSP.

Including the **active** option with the **install** statement installs the specified prefix into the **inet.0** routing table, which is the primary forwarding table. The result is a route that is installed in the forwarding table any time the LSP is established, which means you can ping or trace the route. Use this option with care, because this type of prefix is very similar to a static route.

You use alias routes for routers that have multiple addresses being used as BGP next hops, or for routers that are not MPLS capable. In either of these cases, the LSP can be configured to another MPLS capable system within the local domain, which then acts as a “border” router. The LSP then terminates on the border router and, from that router, Layer 3 forwarding takes the packet to the true next-hop router.

In the case of an interconnect, the domain’s border router can act as the proxy router and can advertise the prefix for the interconnect if the border router is not setting the BGP next hop to itself.

In the case of a point of presence (POP) that has routers that do not support MPLS, one router (for example, a core router) that supports MPLS can act as a proxy for the entire POP and can inject a set of prefixes that cover the POP. Thus, all routers within the POP can advertise themselves as interior BGP (IBGP) next hops, and traffic can follow the LSP to reach the core router. This means that normal IGP routing would prevail within the POP.

You cannot use the `ping` or `traceroute` commands on routes in the `inet.3` routing table.

For BGP next-hop resolution, it makes no difference whether a route is in `inet.0` or `inet.3`; the route with the best match (longest mask) is chosen. Among multiple best-match routes, the one with the highest preference value is chosen.

Configuring the Connection Between Ingress and Egress Routers

The ingress router might make many attempts to connect and reconnect to the egress router using the primary path. You can control how often the ingress router tries to establish a connection using the primary path and how long it waits between retry attempts.

The retry timer configures how long the ingress router waits before trying to connect again to the egress router using the primary path. The default retry time is 30 seconds. The time can be from 1 through 600 seconds. To modify this value, include the `retry-timer` statement:

```
retry-timer seconds;
```

You can configure this statement at the following hierarchy levels:

- [edit protocols mpls label-switched-path *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*]

By default, no limit is set to the number of times an ingress router attempts to establish or reestablish a connection to the egress router using the primary path. To limit the number of attempts, include the `retry-limit` statement:

```
retry-limit number;
```

You can configure this statement at the following hierarchy levels:

- [edit protocols mpls label-switched-path *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*]

The limit can be a value up to 10,000. When the retry limit is exceeded, no more attempts are made to establish a path connection. At this point, intervention is required to restart the primary path.

If you set a retry limit, it is reset to 1 each time a successful primary path is created.

Configuring LSP Metrics

The LSP metric is used to indicate the ease or difficulty of sending traffic over a particular LSP. Lower LSP metric values (lower cost) increase the likelihood of an LSP being used. Conversely, high LSP metric values (higher cost) decrease the likelihood of an LSP being used.

The LSP metric can be specified dynamically by the router or explicitly by the user as described in the following sections:

- Configuring Dynamic LSP Metrics on page 82
- Configuring Static LSP Metrics on page 82

Configuring Dynamic LSP Metrics

If no specific metric is configured, an LSP attempts to track the IGP metric toward the same destination (the to address of the LSP). IGP includes OSPF, IS-IS, Routing Information Protocol (RIP), and static routes. BGP and other RSVP or LDP routes are excluded.

For example, if the OSPF metric toward a router is 20, all LSPs toward that router automatically inherit metric 20. If the OSPF toward a router later changes to a different value, all LSP metrics change accordingly. If there are no IGP routes toward the router, the LSP raises its metric to 65,535.

Note that in this case, the LSP metric is completely determined by IGP; it bears no relationship to the actual path the LSP is currently traversing. If LSP reroutes (such as through reoptimization), its metric does not change, and thus it remains transparent to users. Dynamic metric is the default behavior; no configuration is required.

Configuring Static LSP Metrics

You can manually assign a fixed metric value to an LSP. Once configured with the `metric` statement, the LSP metric is fixed and cannot change:

```
metric number;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls label-switched-path *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*]

The LSP metric has several uses:

- When there are parallel LSPs with the same egress router, the metrics are compared to determine which LSP has the lowest metric value (the lowest cost) and therefore the preferred path to the destination. If the metrics are the same, the traffic is shared.

Adjusting the metric values can force traffic to prefer some LSPs over others, regardless of the underlying IGP metric.

- When an IGP shortcut is enabled (see “IGP Shortcuts” on page 34), an IGP route might be installed in the routing table with an LSP as the next hop, if the LSP is on the shortest path to the destination. In this case, the LSP metric is added to the other IGP metrics to determine the total path metric. For example, if an LSP whose ingress router is X and egress router is Y is on the shortest path to destination Z, the LSP metric is added to the metric for the IGP route from Y to Z to determine the total cost of the path. If several LSPs are potential next hops, the total metrics of the paths are compared to determine which path is preferred (that is, has the lowest total metric). Or, IGP paths and LSPs leading to the same destination could be compared by means of the metric value to determine which path is preferred.

By adjusting the LSP metric, you can force traffic to prefer LSPs, prefer the IGP path, or share the load among them.

- If router X and Y are BGP peers and if there is an LSP between them, the LSP metric represents the total cost to reach Y from X. If for any reason the LSP reroutes, the underlying path cost might change significantly, but X’s cost to reach Y remains the same (the LSP metric), which allows X to report through a BGP multiple exit discriminator (MED) a stable metric to downstream neighbors. As long as Y remains reachable through the LSP, no changes are visible to downstream BGP neighbors.

It is possible to configure IS-IS to ignore the configured LSP metric by including the `ignore-lsp-metrics` statement at the `[edit protocols isis traffic-engineering shortcuts]` hierarchy level. This statement removes the mutual dependency between IS-IS and MPLS for path computation. For more information, see the *JUNOS Routing Protocols Configuration Guide*.

Configuring CSPF Tie Breaking

When selecting a path for an LSP, CSPF uses a tie-breaking process if there are several equal-cost paths. For information about how CSPF selects a path, see “How CSPF Selects a Path” on page 31.

You can configure one of the following statements (you can only configure one of these statements at a time) to alter the behavior of CSPF tie-breaking:

- To configure a random tie-breaking rule for CSPF to use to choose among equal-cost paths, include the `random` statement:

```
random;
```
- To prefer the path with the least-utilized links, include the `least-fill` statement:

```
least-fill;
```
- To prefer the path with the most-utilized links, include the `most-fill` statement:

```
most-fill;
```

You can include each of these statements at the following hierarchy levels:

- [edit protocols mpls label-switched-path *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*]

Configuring Load Balancing for MPLS LSPs

Load balancing is used to evenly distribute traffic when the following conditions apply:

- There are multiple equal-cost next hops over different interfaces to the same destination.
- There is a single next hop over an aggregated interface.

By default, when load balancing is used to help distribute traffic, the JUNOS Software employs a hash algorithm to select a next-hop address to install into the forwarding table. Whenever the set of next hops for a destination changes in any way, the next-hop address is reselected by means of the hash algorithm.

You can configure how the hash algorithm is used to load-balance traffic across a set of equal-cost LSPs. The hash algorithm can be configured to use the first MPLS label, the first two MPLS labels, the IP payload, or the first and second MPLS labels and the IP payload.

For more information about statements configured under the [edit forwarding-options] hierarchy level, see the *JUNOS Policy Framework Configuration Guide*.

The following sections describe how to configure load balancing for MPLS LSPs:

- Using the First MPLS Label in the Hash Key on page 84
- Using the Second MPLS Label in the Hash Key on page 84
- Using the Third MPLS Label in the Hash Key on page 85
- Using the IP Payload in the Hash Key on page 85
- Using the First Two Labels and the IP Payload in the Hash Key on page 85
- Configuring Load Balancing for MPLS LSPs Without CSPF on page 86

Using the First MPLS Label in the Hash Key

To use the first MPLS label in the hash key, include the **label-1** statement at the [edit forwarding-options hash-key family mpls] hierarchy level:

```
[edit forwarding-options hash-key family mpls]
label-1;
```

Using the Second MPLS Label in the Hash Key

To use the second MPLS label in the hash key, include both the **label-1** and **label-2** statements at the [edit forwarding-options hash-key family mpls] hierarchy level:

```
[edit forwarding-options hash-key family mpls]
label-1;
```

```
label-2;
```

Using the Third MPLS Label in the Hash Key

To use the third MPLS label in the hash key, include the `label-1`, `label-2`, and `label-3` statements at the `[edit forwarding-options hash-key family mpls]` hierarchy level:

```
[edit forwarding-options hash-key family mpls]
label-1;
label-2;
label-3;
```

Using the IP Payload in the Hash Key

To use the MPLS packet's IP payload (IP version 4 [IPv4] or IP version 6 [IPv6]) in the hash key, include the `no-labels` statement and the `payload` statement with the `ip` option at the `[edit forwarding-options hash-key family mpls]` hierarchy level:

```
[edit forwarding-options hash-key family mpls]
no-labels;
payload {
  ip;
}
```



NOTE: The router determines if the MPLS payload is an IP packet by checking the byte containing the IP version number. If the IP version number is 4 (IPv4) or 6 (IPv6), the packet is assumed to be an IP packet.

Using the First Two Labels and the IP Payload in the Hash Key

To use the first and second MPLS labels and the MPLS packet's IP payload in the hash key, include the `label-1` and `label-2` statements and the `payload` statement with the `ip` option at the `[edit forwarding-options hash-key family mpls]` hierarchy level:

```
[edit forwarding-options hash-key family mpls]
label-1;
label-2;
payload {
  ip;
}
```



NOTE: You can include this combination of statements on M320 and T Series Core Routers only. If you include them on an M Series Multiservice Edge Router, only the first MPLS label and the IP payload are used in the hash key.

Configuring Load Balancing for MPLS LSPs Without CSPF

An LSP tends to load-balance its placement by randomly selecting one of the equal-cost next hops and using it exclusively. The random selection is made independently at each transit router, which compares IGP metrics alone. No consideration is given to bandwidth or congestion levels.

Disabling Normal TTL Decrementing

By default, the time to live (TTL) field value in the packet header is decremented by 1 for every hop the packet traverses in the LSP, thereby preventing loops. If the TTL field value reaches 0, packets are dropped, and an Internet Control Message Protocol (ICMP) error packet can be sent to the originating router.

If normal TTL decrement is disabled, the TTL field of IP packets entering LSPs are decremented by only 1 on transiting the LSP, making the LSP appear as a one-hop router to diagnostic tools, such as **traceroute**. Decrementing the TTL field by 1 is done by the ingress router, which pushes a label on IP packets with the TTL field in the label initialized to 255. The label's TTL field value is decremented by 1 for every hop the MPLS packet traverses in the LSP. On the penultimate hop of the LSP, the router pops the label but does not write the label's TTL field value to the IP packet's TTL field. Instead, when the IP packet reaches the egress router, the IP packet's TTL field value is decremented by 1.

When you use **traceroute** to diagnose problems with an LSP from outside that LSP, **traceroute** sees the ingress router, although the egress router performs the TTL decrement. The behavior of **traceroute** is different if it is initiated from the ingress router of the LSP. In this case, the egress router would be the first router to respond to **traceroute**.

You can disable normal TTL decrementing in an LSP so that the TTL field value does not reach 0 before the packet reaches its destination, thus preventing the packet from being dropped. You can also disable normal TTL decrementing to make the MPLS cloud appear as a single hop, thereby hiding the network topology.

There are two ways to disable TTL decrementing:

- On the ingress of the LSP, if you include the **no-decrement-ttl** statement, the ingress router negotiates with all downstream routers using a proprietary RSVP object, to ensure all routers are in agreement. If negotiation succeeds, the whole LSP behaves as one hop to transit IP traffic.

no-decrement-ttl;

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Note that the RSVP object is proprietary to the JUNOS Software and might not work with other software. This potential incompatibility applies only to RSVP-signaled LSPs. When you include the **no-decrement-ttl** statement, TTL hiding can be enforced on a per-LSP basis.

- On the router, you can include the **no-propagate-ttl** statement. This statement applies to all LSPs, regardless of whether they are RSVP-signaled or LDP-signaled. Once set, all future LSPs traversing through this router behave as a single hop to IP packets. LSPs established before you configure this statement are not affected.

```
no-propagate-ttl;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls]
- [edit logical-systems *logical-system-name* protocols mpls]

If you include the **no-propagate-ttl** statement, make sure all routers are configured consistently within an MPLS domain; failing to do so might cause the IP packet TTL to increase while in transit within LSPs. This can happen, for example, when the ingress router has **no-propagate-ttl** configured but the penultimate router does not, so the penultimate router writes the MPLS TTL value (which starts from the ingress router as 255) into the IP packet.

The operation of the **no-propagate-ttl** statement is more interoperable with other vendors' equipment. However, you must ensure that all routers are configured identically.

Configuring MPLS Soft Preemption

Soft preemption attempts to establish a new path for a preempted LSP before tearing down the original LSP. The default behavior is to tear down a preempted LSP first, signal a new path, and then reestablish the LSP over the new path. In the interval between when the path is taken down and the new LSP is established, any traffic attempting to use the LSP is lost. Soft preemption prevents this type of traffic loss. The trade-off is that during the time when an LSP is being soft preempted, two LSPs with their corresponding bandwidth requirements are used until the original path is torn down.

MPLS soft preemption is useful for network maintenance. For example, you can move all LSPs away from a particular interface, then take the interface down for maintenance without interrupting traffic. MPLS soft preemption is described in detail in Internet draft draft-ietf-mpls-soft-preemption-02.txt, *MPLS Traffic Engineering Soft Preemption*.

Soft preemption is a property of the LSP and is disabled by default. You configure it at the ingress of an LSP by including the **soft-preemption** statement:

```
soft-preemption;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls label-switched-path *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*]

You can also configure a timer for soft preemption. The timer designates the length of time the router should wait before initiating a hard preemption of the LSP. At the end of the time specified, the LSP is torn down and resigaled. The soft-preemption cleanup timer has a default value of 30 seconds; the range of permissible values is 0 through 180 seconds. A value of 0 means that soft preemption is disabled. The soft-preemption cleanup timer is global for all LSPs.

Configure the timer by including the `cleanup-timer` statement:

```
cleanup-timer seconds;
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp preemption soft-preemption]
- [edit logical-systems *logical-system-name* protocols rsvp preemption soft-preemption]



NOTE: Soft preemption cannot be configured on LSPs for which secondary paths or fast reroute has been configured. The configuration fails to commit. However, you can enable soft preemption in conjunction with node and link protection.

Configuring Automatic Bandwidth Allocation for LSPs

Automatic bandwidth allocation allows an MPLS tunnel to automatically adjust its bandwidth allocation based on the volume of traffic flowing through the tunnel. You can configure an LSP with minimal bandwidth, and this feature can dynamically adjust the LSP's bandwidth allocation based on current traffic patterns. The bandwidth adjustments do not interrupt traffic flow through the tunnel.

At the end of the automatic bandwidth allocation time interval, the current maximum average bandwidth usage is compared with the allocated bandwidth for the LSP. If the LSP needs more bandwidth, an attempt is made to set up a new path where bandwidth is equal to the current maximum average usage. If the attempt is successful, the LSP's traffic is routed through the new path and the old path is removed. If the attempt fails, the LSP continues to use its current path.



NOTE: You might not be able to use this feature to adjust the bandwidth of fast-reroute LSPs. Because the LSPs use a fixed filter (FF) reservation style, when a new path is signaled, the bandwidth might be double-counted. Double-counting can prevent a fast-reroute LSP from ever adjusting its bandwidth when automatic bandwidth allocation is enabled.

To configure automatic bandwidth allocation, complete the steps in the following sections:

- Configuring MPLS Statistics for Automatic Bandwidth Allocation on page 89
- Configuring Automatic Bandwidth Allocation on LSPs on page 89
- Requesting Automatic Bandwidth Allocation Adjustment on page 94

Configuring MPLS Statistics for Automatic Bandwidth Allocation

To enable automatic bandwidth allocation, you first need to configure MPLS statistics. Include the `auto-bandwidth` option for the `statistics` statement. You can also use the `interval` option to specify the interval for calculating the average bandwidth usage.

These settings apply to all LSPs configured on the router on which you have also configured the `auto-bandwidth` statement at the `[edit protocols mpls label-switched-path label-switched-path-name]` hierarchy level. You can also set the adjustment interval on specific LSPs.



NOTE: To prevent unnecessary resignaling of LSPs, it is best to configure an MPLS automatic bandwidth statistics interval of no more than one third the corresponding LSP adjustment interval. For example, if you configure a value of 30 seconds for the `interval` statement at the `[edit protocols mpls statistics]` hierarchy level, you should configure a value of no more than 90 seconds for the `adjust-interval` statement at the `[edit protocols mpls label-switched-path label-switched-path-name auto-bandwidth]` hierarchy level. See also “Configuring the Automatic Bandwidth Allocation Interval” on page 90.

To configure the MPLS and automatic bandwidth allocation statistics, include the `statistics` statement:

```
statistics {
  auto-bandwidth;
  file filename <files number> <size size> <world-readable | no-world-readable>;
  interval seconds;
}
```

You can include this statement at the following hierarchy levels:

- `[edit protocols mpls]`
- `[edit logical-systems logical-system-name protocols mpls]`

Configuring Automatic Bandwidth Allocation on LSPs

To enable automatic bandwidth allocation on an LSP, include the `auto-bandwidth` statement:

```
auto-bandwidth {
  adjust-interval seconds;
  adjust-threshold percent;
  adjust-threshold-overflow-limit number;
  minimum-bandwidth bps;
  maximum-bandwidth bps;
  monitor-bandwidth;
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls label-switched-path *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*]

The statements configured at the [edit protocols mpls label-switched-path *label-switched-path-name* auto-bandwidth] hierarchy level are optional and explained in the following sections:

- Configuring the Automatic Bandwidth Allocation Interval on page 90
- Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth on page 90
- Configuring the Automatic Bandwidth Adjustment Threshold on page 91
- Configuring a Limit on Bandwidth Overflow Samples on page 91
- Configuring Passive Bandwidth Utilization Monitoring on page 93

Configuring the Automatic Bandwidth Allocation Interval

At the end of the automatic bandwidth allocation interval, the automatic bandwidth computation and new path setup process is triggered.



NOTE: To prevent unnecessary resignaling of LSPs, it is best to configure an LSP adjustment interval of no more than three times the corresponding MPLS automatic bandwidth statistics interval. For example, if you configure a value of 90 seconds for the `adjust-interval` statement at the [edit protocols mpls label-switched-path *label-switched-path-name* auto-bandwidth] hierarchy level, you should configure a value of no less than 30 seconds for the `interval` statement at the [edit protocols mpls statistics] hierarchy level. See also “Configuring MPLS Statistics for Automatic Bandwidth Allocation” on page 89.

To specify the bandwidth reallocation interval in seconds for a specific LSP, include the `adjust-interval` statement:

```
adjust-interval seconds;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls label-switched-path *lsp-name* auto-bandwidth]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* auto-bandwidth]

Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth

You can maintain the LSP's bandwidth between minimum and maximum bounds by specifying values for the `minimum-bandwidth` and `maximum-bandwidth` statements.

To specify the minimum amount of bandwidth allocated for a specific LSP, include the `minimum-bandwidth` statement:

```
minimum-bandwidth bps;
```


You can include this statement at the following hierarchy levels:

- [edit protocols mpls label-switched-path *lsp-name* auto-bandwidth]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* auto-bandwidth]

To specify the maximum amount of bandwidth allocated for a specific LSP, include the `maximum-bandwidth` statement:

```
maximum-bandwidth bps;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls label-switched-path *lsp-name* auto-bandwidth]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* auto-bandwidth]

Configuring the Automatic Bandwidth Adjustment Threshold

Use the `adjust-threshold` statement to specify the sensitivity of the automatic bandwidth adjustment of an LSP to changes in bandwidth utilization. You can set the threshold for when to trigger automatic bandwidth adjustments. When configured, bandwidth demand for the current interval is determined and compared to the LSP's current bandwidth allocation. If the percentage difference in bandwidth is greater than or equal to the specified `adjust-threshold` percentage, the LSP's bandwidth is adjusted to the current bandwidth demand.

For example, assume that the current bandwidth allocation is 100 megabits per second (Mbps) and that the percentage configured for the `adjust-threshold` statement is 15 percent. If the bandwidth demand increases to 110 Mbps, the bandwidth allocation is not adjusted. However, if the bandwidth demand increases to 120 Mbps (20 percent over the current allocation) or decreases to 80 Mbps (20 percent under the current allocation), the bandwidth allocation is increased to 120 Mbps or decreased to 80 Mbps, respectively.

To configure the threshold for automatic bandwidth adjustment, include the `adjust-threshold` statement:

```
adjust-threshold percent;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls label-switched-path *lsp-name* auto-bandwidth]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* auto-bandwidth]

Configuring a Limit on Bandwidth Overflow Samples

The automatic bandwidth adjustment timer is a periodic timer which is triggered every adjust interval to determine whether any bandwidth adjustments are required on the LSP's active path. This interval is typically configured as a long period of time,

usually hours. If, at the end of adjust interval, the change in bandwidth is above a certain adjust threshold, the LSP is resigaled with the new bandwidth.

During the automatic bandwidth adjustment interval, the router might receive a steady increase in traffic (increasing bandwidth utilization) on an LSP, potentially causing congestion or packet loss. To prevent this, you can define a second trigger to prematurely expire the automatic bandwidth adjustment timer before the end of the current adjustment interval.

Every statistics interval, the router samples the average bandwidth utilization of an LSP and if this has exceeded the current maximum average bandwidth utilization, the maximum average bandwidth utilization is updated.

During each sample period, the following conditions are also checked:

- Is the current average bandwidth utilization above the current maximum average bandwidth utilization.
- Has the change in maximum average bandwidth utilization exceeded the adjust threshold (bandwidth utilization has changed significantly).

If these conditions are true, it is considered to be one bandwidth overflow sample. Using the `adjust-threshold-overflow-limit` statement, you can define a limit on the number of bandwidth overflow samples such that when the limit is reached, the current automatic bandwidth adjustment timer is expired and a bandwidth adjustment is triggered. Once this adjustment is complete, the normal automatic bandwidth adjustment timer is reset to expire after the periodic adjustment interval.

To specify a limit on the number of bandwidth overflow samples before triggering an automatic bandwidth allocation adjustment, configure the `adjust-threshold-overflow-limit` statement:

```
adjust-threshold-overflow-limit number;
```

This statement can be configured at the following hierarchy levels:

- [edit protocols mpls label-switched-path *lsp-name* auto-bandwidth]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* auto-bandwidth]

You must also configure the `adjust-threshold` and `maximum-bandwidth` statements whenever you configure the `adjust-threshold-overflow-limit` statement:

- You must configure a nonzero value for the `adjust-threshold` statement if you configure the `adjust-threshold-overflow-limit` statement. Any bandwidth increase below the value configured for the `adjust-threshold` statement does not constitute an overflow condition.
- To prevent unlimited increases in LSP bandwidth (to limit overflow beyond a certain bandwidth), you must also configure the `maximum-bandwidth` statement when you configure the `adjust-threshold-overflow-limit` statement.

Failure to configure either of these statements when you configure the `adjust-threshold-overflow-limit` statement results in a commit error.

The following describes the other aspects of the **adjust-threshold-overflow-limit** statement:

- It only applies to bandwidth overflows. If the bandwidth is decreasing, the normal automatic bandwidth adjustment interval is used.
- It does not affect manually triggered automatic bandwidth adjustment.
- It applies to single-class DiffServ-TE LSPs.
- Because the **adjust-threshold-overflow-limit** statement can trigger a bandwidth adjustment, it cannot be enabled at the same time as the **monitor-bandwidth** statement (for information about that statement, see “Configuring Passive Bandwidth Utilization Monitoring” on page 93).
- You cannot configure automatic bandwidth adjustments to occur more often than every 300 seconds. The **adjust-threshold-overflow-limit** statement is subject to the same minimum value with regard to the minimum frequency of adjustment allowed. Overflow condition based adjustments can occur no sooner than 300 seconds from the start of the overflow condition. Therefore it is required that:

sample interval x adjust-threshold-overflow-limit >= 300s

These values are checked during the commit operation. An error is returned if the value is less than 300 seconds.

- If you change the value of the **adjust-threshold-overflow-limit** statement on a working router, you can expect the following behavior:
 - If you increase the current value of the **adjust-threshold-overflow-limit** statement, the old value is replaced with the new one.
 - If you decrease the current value of the **adjust-threshold-overflow-limit** statement and the current bandwidth overflow count is less than the new value, the old value is replaced with the new one.
 - If you decrease the current value of the **adjust-threshold-overflow-limit** statement and the current bandwidth overflow count is greater than the new value, the adjustment timer is immediately expired and a bandwidth adjustment is initiated.

Configuring Passive Bandwidth Utilization Monitoring

Use the **monitor-bandwidth** statement to switch to a passive bandwidth utilization monitoring mode. In this mode, no automatic bandwidth adjustments are made, but the maximum average bandwidth utilization is continuously monitored and recorded.

To configure passive bandwidth utilization monitoring, include the **monitor-bandwidth** statement:

monitor-bandwidth;

You can include this statement at the following hierarchy levels:

- `[edit protocols mpls label-switched-path lsp-name auto-bandwidth]`

- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* auto-bandwidth]

If you have configured an LSP with primary and secondary paths, the automatic bandwidth allocation statistics are carried over to the secondary path if the primary path fails. For example, consider a primary path whose adjustment interval is half complete and whose maximum average bandwidth usage is currently calculated as 50 Mbps. If the primary path suddenly fails, the time remaining for the next adjustment and the maximum average bandwidth usage are carried over to the secondary path.

Requesting Automatic Bandwidth Allocation Adjustment

For MPLS LSP automatic bandwidth allocation adjustment, the minimum value for the adjustment interval is 5 minutes (300 seconds). You might find it necessary to trigger a bandwidth allocation adjustment manually, for example in the following circumstances:

- When you are testing automatic bandwidth allocation in a network lab.
- When the LSP is configured for automatic bandwidth allocation in monitor mode (the **monitor-bandwidth** statement is included in the configuration as described in “Configuring Passive Bandwidth Utilization Monitoring” on page 93), and want to initiate an immediate bandwidth adjustment.

To use the **request mpls lsp adjust-autobandwidth** command, the following must be true:

- Automatic bandwidth allocation must be enabled on the LSP.
- The criteria required to trigger a bandwidth adjustment have been met (the difference between the adjust bandwidth and the current LSP path bandwidth is greater than the threshold limit).

A manually triggered bandwidth adjustment operates only on the active LSP path. Also, if you have enabled periodic automatic bandwidth adjustment, the periodic automatic bandwidth adjustment parameters (the adjustment interval and the maximum average bandwidth) are not reset after a manual adjustment.

For example, suppose the periodic adjust interval is 10 hours and there are currently 5 hours remaining before an automatic bandwidth adjustment is triggered. If you initiate a manual adjustment with the **request mpls lsp adjust-autobandwidth** command, the adjust timer is not reset and still has 5 hours remaining.

To manually trigger a bandwidth allocation adjustment, you need to use the **request mpls lsp adjust-autobandwidth** command. You can trigger the command for all affected LSPs on the router, or you can specify a particular LSP:

```
user@host> request mpls lsp adjust-autobandwidth
```

Once you execute this command, the automatic bandwidth adjustment validation process is triggered. If all the criteria for adjustment are met, the LSP's active path

bandwidth is adjusted to the adjusted bandwidth value determined during the validation process.

Disabling Constrained-Path LSP Computation

If the IGP is a link-state protocol (such as IS-IS or OSPF) and supports extensions that allow the current bandwidth reservation on each router's link to be reported, constrained-path LSPs are computed by default.

The JUNOS implementations of IS-IS and OSPF include the extensions that support constrained-path LSP computation.

- IS-IS—These extensions are enabled by default. To disable this support, include the `disable` statement at the `[edit protocols isis traffic-engineering]` hierarchy level, as discussed in the *JUNOS Routing Protocols Configuration Guide*.
- OSPF—These extensions are disabled by default. To enable this support, include the `traffic-engineering` statement in the configurations of all routers running OSPF, as described in the *JUNOS Routing Protocols Configuration Guide*.

If IS-IS is enabled on a router or you enable OSPF traffic engineering extensions, MPLS performs the constrained-path LSP computation by default. For information on how constrained-path LSP computation works, see “Constrained-Path LSP Computation” on page 29.

Constrained-path LSPs have a greater chance of being established quickly and successfully for the following reasons:

- The LSP computation takes into account the current bandwidth reservation.
- Constrained-path LSPs reroute themselves away from node failures and congestion.

When constrained-path LSP computation is enabled, you can configure the LSP so that it is periodically reoptimized, as described in “Optimizing Signaled LSPs” on page 103.

When an LSP is being established or when an existing LSP fails, the constrained-path LSP computation is repeated periodically at the interval specified by the retry timer until the LSP is set up successfully. Once the LSP is set up, no recomputation is done. For more information about the retry timer, see “Configuring the Connection Between Ingress and Egress Routers” on page 81.

By default, constrained-path LSP computation is enabled. You might want to disable constrained-path LSP computation when all nodes do not support the necessary traffic engineering extensions. To disable constrained-path LSP computation, include the `no-cspf` statement:

```
no-cspf;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

If you disable constrained-path LSP computation on LSPs by configuring the `no-cspf` statement and then attempt to advertise other LSPs with lower metrics than the IGPs from this router in either IS-IS or OSPF, new LSPs cannot be established.

Configuring Administrative Groups

Administrative groups, also known as link coloring or resource class, are manually assigned attributes that describe the “color” of links, such that links with the same color conceptually belong to the same class. You can use administrative groups to implement a variety of policy-based LSP setups.

Administrative groups are meaningful only when constrained-path LSP computation is enabled.

You can assign up to 32 names and values (in the range 0 through 31), which define a series of names and their corresponding values. The administrative names and values must be identical across all routers within a single domain.



NOTE: The administrative value is distinct from the priority. You configure the priority for an LSP using the `priority` statement. See “Configuring Priority and Preemption for LSPs” on page 102.

To configure administrative groups, follow these steps:

1. Define multiple levels of service quality by including the `admin-groups` statement:

```
admin-groups {
  group-name group-value;
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls]
- [edit logical-systems *logical-system-name* protocols mpls]

The following configuration example illustrates how you might configure a set of administrative names and values for a domain:

```
[edit protocols mpls]
admin-groups {
  gold 1;
  silver 2;
  copper 3;
  best-effort 4;
}
```

2. Define the administrative groups to which an interface belongs. You can assign multiple groups to an interface. Include the `interface` statement:

```
interface interface-name {
  admin-group [ group-names ];
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls]
- [edit logical-systems *logical-system-name* protocols mpls]

If you do not include the **admin-group** statement, an interface does not belong to any group.

IGPs use the group information to build link-state packets, which are then flooded throughout the network, providing information to all nodes in the network. At any router, the IGP topology, as well as administrative groups of all the links, is available.

Changing the interface's administrative group affects only new LSPs. Existing LSPs on the interface are not preempted or recomputed to keep the network stable. If LSPs need to be removed because of a group change, issue the **clear rsvp session** command.

3. Configure an administrative group constraint for each LSP or for each primary or secondary LSP path. Include the **label-switched-path** statement:

```
label-switched-path lsp-name {
  to address;
  ...
  admin-group {
    exclude [ group-names];
    include-all [ group-names];
    include-any [ group-names];
  }
  primary path-name {
    admin-group {
      exclude [ group-names];
      include-all [ group-names];
      include-any [ group-names];
    }
  }
  secondary path-name {
    admin-group {
      exclude [ group-names];
      include-all [ group-names];
      include-any [ group-names];
    }
  }
}
```

You can include the **label-switched-path** statement at the following hierarchy levels:

- [edit protocols mpls]
- [edit logical-systems *logical-system-name* protocols mpls]

If you omit the **include-all**, **include-any**, or **exclude** statements, the path computation proceeds unchanged. The path computation is based on the constrained-path

LSP computation. For information on how the constrained-path LSP computation is calculated, see “How CSPF Selects a Path” on page 31.



NOTE: Changing the LSP’s administrative group causes an immediate recomputation of the route; therefore, the LSP might be rerouted.

Configuring Preference Values for LSPs

As an option, you can configure multiple LSPs between the same pair of ingress and egress routers. This is useful for balancing the load among the LSPs because all LSPs, by default, have the same preference level. To prefer one LSP over another, set different preference levels for individual LSPs. The LSP with the lowest preference value is used. The default preference for RSVP LSPs is 7 and for LDP LSPs is 9. These preference values are lower (more preferred) than all learned routes except direct interface routes.

To change the default preference value, include the **preference** statement:

```
preference preference;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Disabling Path Route Recording

The JUNOS implementation of RSVP supports the Record Route object, which allows an LSP to actively record the routers through which it transits. You can use this information for troubleshooting and to prevent routing loops. By default, path route information is recorded. To disable recording, include the **no-record** statement:

```
no-record;
```

For a list of hierarchy levels at which you can include the **record** and **no-record** statements, see the statement summary section for the statement.

Configuring Class of Service for MPLS LSPs

The following sections provide an overview of MPLS class of service (CoS) and describe how to configure the MPLS CoS value:

- Class of Service for MPLS Overview on page 99
- Configuring the MPLS CoS Bits on page 99
- Rewriting IEEE 802.1p Packet Headers with the MPLS CoS Value on page 100

Class of Service for MPLS Overview

When IP traffic enters an LSP tunnel, the ingress router marks all packets with a CoS value, which is used to place the traffic into a transmission priority queue. On the router, for SDH/SONET and T3 interfaces, each interface has four transmit queues. The CoS value is encoded as part of the MPLS header and remains in the packets until the MPLS header is removed when the packets exit from the egress router. The routers within the LSP utilize the CoS value set at the ingress router. The CoS value is encoded by means of the CoS bits (also known as the EXP or experimental bits). For more information, see “Label Allocation” on page 26.

MPLS class of service works in conjunction with the router’s general CoS functionality. If you do not configure any CoS features, the default general CoS settings are used. For MPLS class of service, you might want to prioritize how the transmit queues are serviced by configuring weighted round-robin, and to configure congestion avoidance using random early detection (RED). The general CoS features are described in the *JUNOS Class of Service Configuration Guide*.

Configuring the MPLS CoS Bits

When traffic enters an LSP tunnel, the CoS bits in the MPLS header are set in one of two ways:

- The number of the output queue into which the packet was buffered and the packet loss priority (PLP) bit are written into the MPLS header and are used as the packet’s CoS value. This behavior is the default, and no configuration is required. The *JUNOS Class of Service Configuration Guide* explains the IP CoS values, and summarizes how the CoS bits are treated.
- You set a fixed CoS value on all packets entering the LSP tunnel. A fixed CoS value means that all packets entering the LSP receive the same class of service.

To set a fixed CoS value on all packets entering the LSP, include the **class-of-service** statement:

```
class-of-service cos-value;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

The CoS value set using the **class-of-service** statement at the [edit protocols mpls] hierarchy level supersedes the CoS value set at the [edit class-of-service] hierarchy level for an interface. Effectively, the CoS value configured for an LSP overrides the CoS value set for an interface.

The CoS value can be a decimal number from 0 through 7. This number corresponds to a 3-bit binary number. The high-order 2 bits of the CoS value select which transmit queue to use on the outbound interface card.

The low-order bit of the CoS value is treated as the PLP bit and is used to select the RED drop profile to use on the output queue. If the low-order bit is 0, the non-PLP drop profile is used, and if the low-order bit is 1, the PLP drop profile is used. It is

generally expected that RED will more aggressively drop packets that have the PLP bit set. For more information about RED and drop profiles, see the *JUNOS Class of Service Configuration Guide*.



NOTE: Configuring the PLP drop profile to drop packets more aggressively (for example, setting the CoS value from 6 to 7) decreases the likelihood of traffic getting through.

Table 3 on page 100 summarizes how MPLS CoS values correspond to the transmit queue and PLP bit. Note that in MPLS, the mapping between the CoS bit value and the output queue is hard-coded. You cannot configure the mapping for MPLS; you can configure it only for IPv4 traffic flows, as described in the *JUNOS Class of Service Configuration Guide*.

Table 3: MPLS CoS Values

MPLS CoS Value	Bits	Transmit Queue	PLP Bit
0	000	0	Not set
1	001	0	Set
2	010	1	Not set
3	011	1	Set
4	100	2	Not set
5	101	2	Set
6	110	3	Not set
7	111	3	Set

Because the CoS value is part of the MPLS header, the value is associated with the packets only as they travel through the LSP tunnel. The value is not copied back to the IP header when the packets exit from the LSP tunnel.

Rewriting IEEE 802.1p Packet Headers with the MPLS CoS Value

For Ethernet interfaces installed on a T Series router or an M320 router with a peer connection to an M Series router or a T Series router, you can rewrite both MPLS CoS and IEEE 802.1p bits to a configured value (the MPLS CoS bits are also known as the EXP or experimental bits). Rewriting these bits allows you to pass the configured value to the Layer 2 VLAN path. To rewrite both the MPLS CoS and IEEE 802.1p bits, you must include the EXP and IEEE 802.1p rewrite rules in the class of service interface configuration. The EXP rewrite table is applied when you configure the IEEE 802.1p and EXP rewrite rules.

For information about how to configure the EXP and IEEE 802.1p rewrite rules, see the *JUNOS Class of Service Configuration Guide*.

For information on the CoS bits, see “Label Allocation” on page 26 and “Configuring Class of Service for MPLS LSPs” on page 98.

Configuring Adaptive LSPs

An LSP occasionally might need to reroute itself. Reasons include the following:

- Continuous reoptimization process is configured with the **optimize-timer** statement.
- The current path has connectivity problems.
- The LSP is preempted by another LSP configured with the **priority** statement and is forced to reroute.
- The explicit-path information for an active LSP is modified, or the LSP’s bandwidth is increased.

You can configure an LSP to be *adaptive* when it is attempting to reroute itself. When it is adaptive, the LSP holds onto existing resources until the new path is successfully established and traffic has been cut over to the new LSP. To retain its resources, an adaptive LSP does the following:

- Maintains existing paths and allocated bandwidths—This ensures that the existing path is not torn down prematurely and allows the current traffic to continue flowing while the new path is being set up.
- Avoids double-counting for links that share the new and old paths—Double-counting occurs when an intermediate router does not recognize that the new and old paths belong to the same LSP and counts them as two separate LSPs, requiring separate bandwidth allocations. If some links are close to saturation, double-counting might cause the setup of the new path to fail.

By default, adaptive behavior is disabled. You can include the **adaptive** statement in two different hierarchy levels.

If you specify the **adaptive** statement at the LSP hierarchy levels, the adaptive behavior is enabled on all primary/secondary paths of the LSP. This means both the primary and secondary paths share the same bandwidth on common links.

To configure adaptive behavior for all LSP paths, include the **adaptive** statement in the LSP configuration:

```
adaptive;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls label-switched-path *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*]

If you specify the **adaptive** statement at the [edit protocols mpls label-switched-path *lsp-name* (primary | secondary) *path-name*] hierarchy level, adaptive behavior is enabled

only on the path on which it is specified. Bandwidth double-counting occurs between different paths. However, if you also have the **adaptive** statement configured at the **[edit protocols mpls label-switched-path *lsp-name*]** hierarchy level, it overrides the adaptive behavior of each individual path.

To configure adaptive behavior for either the primary or secondary level, include the **adaptive** statement:

```
adaptive;
```

You can include this statement at the following hierarchy levels:

- **[edit protocols mpls label-switched-path *lsp-name* (primary | secondary) *path-name*]**
- **[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* (primary | secondary) *path-name*]**

Configuring Priority and Preemption for LSPs

When there is insufficient bandwidth to establish a more important LSP, you might want to tear down a less important existing LSP to free the bandwidth. You do this by preempting the existing LSP.

Whether an LSP can be preempted is determined by two properties associated with the LSP:

- **Setup priority**—Determines whether a new LSP that preempts an existing LSP can be established. For preemption to occur, the setup priority of the new LSP must be higher than that of the existing LSP. Also, the act of preempting the existing LSP must produce sufficient bandwidth to support the new LSP. That is, preemption occurs only if the new LSP can be set up successfully.
- **Reservation priority**—Determines the degree to which an LSP holds on to its session reservation after the LSP has been set up successfully. When the reservation priority is high, the existing LSP is less likely to give up its reservation and hence it is unlikely that the LSP can be preempted.

You cannot configure an LSP with a high setup priority and a low reservation priority, because permanent preemption loops might result if two LSPs are allowed to preempt each other. You must configure the reservation priority to be higher than or equal to the setup priority.

The setup priority also defines the relative importance of LSPs on the same ingress router. When the software starts, when a new LSP is established, or during fault recovery, the setup priority determines the order in which LSPs are serviced. Higher-priority LSPs tend to be established first and hence enjoy more optimal path selection.

To configure the LSP's preemption properties, include the **priority** statement:

```
priority setup-priority reservation-priority;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Both *setup-priority* and *reservation-priority* can be a value from 0 through 7. The value 0 corresponds to the highest priority, and the value 7 to the lowest. By default, an LSP has a setup priority of 7 (that is, it cannot preempt any other LSPs) and a reservation priority of 0 (that is, other LSPs cannot preempt it). These defaults are such that preemption does not happen. When you are configuring these values, the setup priority should always be less than or equal to the hold priority.

Optimizing Signaled LSPs

Once an LSP has been established, topology or resources changes might, over time, make the path suboptimal. A new path might have become available that is less congested, has a lower metric, and traverses fewer hops. You can configure the router to recompute paths periodically to determine whether a more optimal path has become available.

If reoptimization is enabled, an LSP can be rerouted through different paths by constrained-path recomputations. However, if reoptimization is disabled, the LSP has a fixed path and cannot take advantage of newly available network resources. The LSP is fixed until the next topology change breaks the LSP and forces a recomputation.

Reoptimization is not related to failover. A new path is always computed when topology failures occur that disrupt an established path.

Because of the potential system overhead involved, you need to control carefully the frequency of reoptimization. Network stability might suffer when reoptimization is enabled. By default, the **optimize-timer** statement is set to 0 (that is, it is disabled).

Configuring LSP optimization is meaningful only when constrained-path LSP computation is enabled, which is the default behavior. For more information about constrained-path LSP computation, see “Disabling Constrained-Path LSP Computation” on page 95.

To enable path reoptimization, include the **optimize-timer** statement:

```
optimize-timer seconds;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Once you have configured the **optimize-timer** statement, the reoptimization timer continues its countdown to the configured value even if you delete the **optimize-timer** statement from the configuration. The next optimization uses the new value. You can force the JUNOS Software to use a new value immediately by deleting the old value, committing the configuration, configuring the new value for the **optimize-timer** statement, and then committing the configuration again.

After reoptimization is run, the result is accepted only if it meets the following criteria:

1. The new path is not higher in IGP metric. (The metric for the old path is updated during computation, so if a recent link metric changed somewhere along the old path, it is accounted for.)
2. If the new path has the same IGP metric, it is not more hops away.

3. The new path does not cause preemption. (This is to reduce the ripple effect of preemption causing more preemption.)
4. The new path does not worsen congestion overall.

The relative congestion of the new path is determined as follows:

- a. The percentage of available bandwidth on each link traversed by the new path is compared to that for the old path, starting from the most congested links.
- b. For each current (old) path, the software stores the four smallest values for bandwidth availability for the links traversed in ascending order.
- c. The software also stores the four smallest bandwidth availability values for the new path, corresponding to the links traversed in ascending order.
- d. If any of the four new available bandwidth values are smaller than any of the corresponding old bandwidth availability values, the new path has at least one link that is more congested than the link used by the old path. Because using the link would cause more congestion, traffic is not switched to this new path.
- e. If none of the four new available bandwidth values is smaller than the corresponding old bandwidth availability values, the new path is less congested than the old path.

When all the above conditions are met, then:

5. If the new path has a lower IGP metric, it is accepted.
6. If the new path has an equal IGP metric and lower hop count, it is accepted.
7. If you choose **least-fill** as a load-balancing algorithm and if the new path reduces congestion by at least 10 percent aggregated over all links it traversed, it is accepted. For **random** or **most-fill** algorithms, this rule does not apply.
8. Otherwise, the new path is rejected.

You can disable the following reoptimization criteria (a subset of the criteria listed previously):

- If the new path has the same IGP metric, it is not more hops away.
- The new path does not cause preemption. (This is to reduce the ripple effect of preemption causing more preemption.)
- The new path does not worsen congestion overall.
- If the new path has an equal IGP metric and lower hop count, it is accepted.

To disable them, either issue the **clear mpls optimize-aggressive** command or include the **optimize-aggressive** statement:

```
optimize-aggressive;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls]
- [edit logical-systems *logical-system-name* protocols mpls]

Including the **optimize-aggressive** statement in the configuration causes the reoptimization procedure to be triggered more often. Paths are rerouted more frequently. It also limits the reoptimization algorithm to the IGP metric only.

Configuring the Smart Optimize Timer

Due to network and router resource constraints, it is typically inadvisable to configure a short interval for the optimize timer. However, under certain circumstances, it might be desirable to reoptimize a path sooner than would normally be provided by the optimize timer.

For example, an LSP is traversing a preferred path which subsequently fails. The LSP is then switched to a less desirable path to reach the same destination. Even if the original path is quickly restored, it could take an excessively long period of time for the LSP to use it again, because it has to wait for the optimize timer to reoptimize the network paths. For situations such as this, you might want to configure the smart optimize timer.

By enabling the smart optimize timer, an LSP is switched back to its original path so long as the original path has been restored within 3 minutes of going down. Also, if the original path goes down again within 60 minutes, the smart optimize timer is disabled and path optimization behaves as it normally does when the optimize timer alone is enabled. This prevents the router from using a flapping link.

To configure smart optimize timer, include the **smart-optimize-timer** statement:

```
smart-optimize-timer seconds;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls]
- [edit logical-systems *logical-system-name* protocols mpls]

Limiting the Number of Hops in LSPs

By default, each LSP can traverse a maximum of 255 hops, including the ingress and egress routers. To modify this value, include the **hop-limit** statement:

```
hop-limit number;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

The number of hops can be from 2 through 255. (A path with two hops consists of the ingress and egress routers only.)

Configuring the Bandwidth Value for LSPs

Each LSP has a bandwidth value. This value is included in the sender's Tspec field in RSVP path setup messages. To specify a bandwidth value, include the **bandwidth** statement:

```
bandwidth bps;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

You specify the bandwidth value in bits per second, with a higher value implying a greater user traffic volume. The default bandwidth is 0 bits per second.

A nonzero bandwidth requires transit and egress routers to reserve capacity along the outbound links for the path. RSVP's reservation scheme is used to reserve this capacity. Any failure in bandwidth reservation (such as failures at RSVP policy control or admission control) might cause the LSP setup to fail. If there is insufficient bandwidth on the interfaces for the transit or egress routers, the LSP is not established.

Configuring Hot Standby of Secondary Paths

By default, secondary paths are set up only as needed. To have the system maintain a secondary path in a hot-standby state indefinitely, include the **standby** statement:

```
standby;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls label-switched-path *lsp-name* secondary]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* secondary]

The hot-standby state is meaningful only on secondary paths. Maintaining a path in a hot-standby state enables swift cutover to the secondary path when downstream routers on the current active path indicate connectivity problems. Although it is possible to configure the **standby** statement at the [edit protocols mpls label-switched-path *lsp-name* primary *path-name*] hierarchy level, it has no effect on router behavior.

If you configure the **standby** statement at the following hierarchy levels, the hot-standby state is activated on all secondary paths configured beneath that hierarchy level:

- [edit protocols mpls]
- [edit protocols mpls label-switched-path *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*]

The hot-standby state has two advantages:

- It eliminates the call-setup delay during network topology changes. Call setup can suffer from significant delays when network failures trigger large numbers of LSP reroutes at the same time.
- A cutover to the secondary path can be made before RSVP learns that an LSP is down. There can be significant delays between the time the first failure is detected by protocol machinery (which can be an interface down, a neighbor becoming unreachable, a route becoming unreachable, or a transient routing loop being detected) and the time an LSP actually fails (which requires a timeout of soft state information between adjacent RSVP routers). When topology failures occur, hot-standby secondary paths can usually achieve the smallest cutover delays with minimal disruptions to user traffic.

When the primary path is considered to be stable again, traffic is automatically switched from the standby secondary path back to the primary path. The switch is performed no faster than twice the retry-timer interval and only if the primary path exhibits stability throughout the entire switch interval.

The drawback of the hot-standby state is that more state information must be maintained by all the routers along the path, which requires overhead from each of the routers.

Damping Advertisement of LSP State Changes

When an LSP changes from being up to being down, or from down to up, this transition takes effect immediately in the router software and hardware. However, when advertising LSPs into IS-IS and OSPF, you may want to damp LSP transitions, thereby not advertising the transition until a certain period of time has transpired (known as the hold time). In this case, if the LSP goes from up to down, the LSP is not advertised as being down until it has remained down for the hold-time period. Transitions from down to up are advertised into IS-IS and OSPF immediately. Note that LSP damping affects only the IS-IS and OSPF advertisements of the LSP; other routing software and hardware react immediately to LSP transitions.

To damp LSP transitions, include the `advertisement-hold-time` statement:

```
advertisement-hold-time seconds;
```

seconds can be a value from 0 through 65,535 seconds. The default is 5 seconds.

You can include this statement at the following hierarchy levels:

- `[edit protocols mpls]`
- `[edit logical-systems logical-system-name protocols mpls]`

Chapter 6

DiffServ-Aware Traffic Engineering Configuration Guidelines

This chapter describes how to configure DiffServ-aware traffic engineering for LSPs and multiclass LSPs:

- DiffServ-Aware Traffic Engineering Introduction on page 110
- DiffServ-Aware Traffic Engineering Standards on page 110
- DiffServ-Aware Traffic Engineering Terminology on page 110
- DiffServ-Aware Traffic Engineering Overview on page 112
- DiffServ-Aware Traffic Engineering Features on page 112
- DiffServ-Aware Traffic Engineered LSPs on page 113
- DiffServ-Aware Traffic Engineered LSPs Overview on page 113
- DiffServ-Aware Traffic Engineered LSPs Operation on page 113
- Multiclass LSPs on page 114
- Multiclass LSP Overview on page 114
- Establishing a Multiclass LSP on the Differentiated Services Domain on page 115
- Configuring Routers for DiffServ-Aware Traffic Engineering on page 115
- Bandwidth Oversubscription Overview on page 119
- LSP Size Oversubscription on page 120
- Link Size Oversubscription on page 120
- Class Type Oversubscription and Local Oversubscription Multipliers on page 121
- Class Type Bandwidth and the LOM on page 121
- LOM Calculation for the MAM and Extended MAM Bandwidth Models on page 122
- LOM Calculation for the Russian Dolls Bandwidth Model on page 122
- Example: LOM Calculation on page 122
- Configuring the Bandwidth Subscription Percentage for LSPs on page 124
- Configuring LSPs for DiffServ-Aware Traffic Engineering on page 125
- Configuring Multiclass LSPs on page 128

DiffServ-Aware Traffic Engineering Introduction

Differentiated Services (DiffServ)-aware traffic engineering provides a way to guarantee a specified level of service over an MPLS network. The routers providing DiffServ-aware traffic engineering are part of a differentiated services network domain. All routers participating in a differentiated services domain must have DiffServ-aware traffic engineering enabled.

To help ensure that the specified service level is provided, it is necessary to ensure that no more than the amount of traffic specified is sent over the differentiated services domain. You can accomplish this goal by configuring a policer to police or rate limit the volume of traffic transiting the differentiated service domain. For more information about how to configure policers for label-switched paths (LSPs), see “Configuring Policers for LSPs” on page 160.

This feature can help to improve the quality of Internet services such as voice over IP (VoIP). It also makes it possible to better emulate an Asynchronous Transfer Mode (ATM) circuit over an MPLS network.

DiffServ-Aware Traffic Engineering Standards

The following RFCs provide information on DiffServ-aware traffic engineering and multiclass LSPs:

- RFC 3270, *Multi-Protocol Label Switching (MPLS) Support of Differentiated Services*
- RFC 3564, *Requirements for Support of Differentiated Services-aware MPLS Traffic Engineering*
- RFC 4124, *Protocol Extensions for Support of Differentiated-Service-Aware MPLS Traffic Engineering*
- RFC 4125, *Maximum Allocation Bandwidth Constraints Model for Diff-Serv-aware MPLS Traffic Engineering*
- RFC 4127, *Russian Dolls Bandwidth Constraints Model for Diff-Serv-aware MPLS*

These RFCs are available on the IETF website at <http://www.ietf.org/>.

DiffServ-Aware Traffic Engineering Terminology

B

Bandwidth model

The bandwidth model determines the values of the available bandwidth advertised by the interior gateway protocols (IGPs).

C

CAC Call admission control (CAC) checks to ensure there is adequate bandwidth on the path before the LSP is established. If the bandwidth is insufficient, the LSP is not established and an error is reported.

Class type A collection of traffic flows that is treated equivalently in a differentiated services domain. A class type maps to a queue and is much like a class-of-service (CoS) forwarding class in concept. It is also known as a traffic class.

D

Differentiated Services Differentiated Services make it possible to give different treatment to traffic based on the experimental (EXP) bits in the MPLS header. Traffic must be marked appropriately and CoS must be configured.

Differentiated Services domain The routers in a network that have Differentiated Services enabled.

DiffServ-aware traffic engineering A type of constraint-based routing. It can enforce different bandwidth constraints for different classes of traffic. It can also do CAC on each traffic engineering class when an LSP is established.

M

MAM The maximum allocation bandwidth constraint model divides the available bandwidth between the different classes. Sharing of bandwidth between the class types is not allowed.

Multiclass LSP A multiclass LSP functions like a standard LSP, but it also allows you to reserve bandwidth from multiple class types. The EXP bits of the MPLS header are used to distinguish between class types.

R

RDM The Russian dolls bandwidth constraint model makes efficient use of bandwidth by allowing the class types to share bandwidth.

T

Traffic engineering class A paired class type and priority.

Traffic engineering class map A map between the class types, priorities, and traffic engineering classes. The traffic engineering class mapping must be consistent across the Differentiated Services domain.

DiffServ-Aware Traffic Engineering Overview

Differentiated Services give different treatment to traffic based on the EXP bits in the MPLS label header and allow you to provide multiple classes of service.

The following sections describe DiffServ-aware traffic engineering:

- DiffServ-Aware Traffic Engineering Features on page 112
- DiffServ-Aware Traffic Engineered LSPs on page 113
- Multiclass LSPs on page 114

DiffServ-Aware Traffic Engineering Features

DiffServ-aware traffic engineering provides the following features:

- Traffic engineering at a per-class level rather than at an aggregate level
- Different bandwidth constraints for different class types (traffic classes)
- Different queuing behaviors per class, allowing the router to forward traffic based on the class type

In comparison, standard traffic engineering does not consider CoS, and it completes its work on an aggregate basis across all Differentiated Service classes.

DiffServ-aware traffic engineering provides the following advantages:

- Traffic engineering can be performed on a specific class type instead of at the aggregate level.
- Bandwidth constraints can be enforced on each specific class type.
- It forwards traffic based on the EXP bits.

This makes it possible to guarantee service and bandwidth across an MPLS network. With DiffServ-aware traffic engineering, among other services, you can provide ATM circuit emulation, VoIP, and a guaranteed bandwidth service.

The following describes how the IGP, Constrained Shortest Path First (CSPF), and RSVP participate in DiffServ-aware traffic engineering:

- The IGP can advertise the unreserved bandwidth for each traffic engineering class to the other members of the differentiated services domain. The traffic engineering database stores this information.
- A CSPF calculation is performed considering the bandwidth constraints for each class type. If all the constraints are met, the CSPF calculation is considered successful.
- When RSVP signals an LSP, it requests bandwidth for specified class types.

DiffServ-Aware Traffic Engineered LSPs

A DiffServ-aware traffic engineered LSP is an LSP configured to reserve bandwidth for one of the supported class types and to carry traffic for that class type. The following sections discuss this type of LSPs:

- DiffServ-Aware Traffic Engineered LSPs Overview on page 113
- DiffServ-Aware Traffic Engineered LSPs Operation on page 113

DiffServ-Aware Traffic Engineered LSPs Overview

A DiffServ-aware traffic engineered LSP is an LSP configured with a bandwidth reservation for a specific class type. This LSP can carry traffic for a single class type. On the packets, the class type is specified by the EXP bits (also known as the class-of-service bits) and the per-hop behavior (PHB) associated with the EXP bits. The mapping between the EXP bits and the PHB is static, rather than being signaled in RSVP.

The class type must be configured consistently across the Differentiated Services domain, meaning the class type configuration must be consistent from router to router in the network. You can unambiguously map a class type to a queue. On each node router, the class-of-service queue configuration for an interface translates to the available bandwidth for a particular class type on that link.

For more information about topics related to LSPs and DiffServ-aware traffic engineering, see the following:

- For forwarding classes and class of service, see the *JUNOS Class of Service Configuration Guide*.
- For EXP bits, see “Label Allocation” on page 26.
- For differentiated services, see RFC 3270, *Multi-Protocol Label Switching (MPLS) Support of Differentiated Services*.
- For information about how the IGPs and RSVP have been modified to support Differentiated Services-aware MPLS traffic engineering, see RFC 4124, *Protocol Extensions for Support of Differentiated-Service-Aware MPLS Traffic Engineering*.

DiffServ-Aware Traffic Engineered LSPs Operation

When configuring a DiffServ-aware traffic engineered LSP, you specify the class type and the bandwidth associated with it. The following occurs when an LSP is established with bandwidth reservation from a specific class type:

1. The IGPs advertise how much unreserved bandwidth is available for the traffic engineering classes.
2. When calculating the path for an LSP, CSPF is used to ensure that the bandwidth constraints are met for the class type carried by the LSP at the specified priority level.

CSPF also checks to ensure that the bandwidth model is configured consistently on each router participating in the LSP. If the bandwidth model is inconsistent, CSPF does not compute the path (except for LSPs from class type `ct0`).

3. Once a path is found, RSVP signals the LSP using the `Classtype` object in the path message. At each node in the path, the available bandwidth for the class types is adjusted as the path is set up.

An LSP that requires bandwidth from a particular class (except class type `ct0`) cannot be established through routers that do not understand the `Classtype` object. Preventing the use of routers that do not understand the `Classtype` object helps to ensure consistency throughout the Differentiated Services domain by preventing the LSP from using a router that cannot support Differentiated Services.

By default, LSPs are signaled with setup priority 7 and holding priority 0. An LSP configured with these values cannot preempt another LSP at setup time and cannot be preempted.

It is possible to have both LSPs configured for DiffServ-aware traffic engineering and regular LSPs configured at the same time on the same physical interfaces. For this type of heterogeneous environment, regular LSPs carry best-effort traffic by default. Traffic carried in the regular LSPs must have the correct EXP settings (either by remarking the EXP settings or by assuming that the traffic arrived with the correct EXP settings from the upstream router).

Multiclass LSPs

Multiclass LSPs function like standard LSPs, but they also allow you to configure multiple class types with guaranteed bandwidth. The EXP bits of the MPLS header are used to distinguish between class types. Multiclass LSPs can be configured for a variety of purposes. For example, you can configure a multiclass LSP to emulate the behavior of an ATM circuit. An ATM circuit can provide service-level guarantees to a class type. A multiclass LSP can provide a similar guaranteed level of service.

The following sections discuss multiclass LSPs:

- Multiclass LSP Overview on page 114
- Establishing a Multiclass LSP on the Differentiated Services Domain on page 115

Multiclass LSP Overview

A multiclass LSP is an LSP that can carry several class types. One multiclass LSP can be used to support up to four class types. On the packets, the class type is specified by the EXP bits (also known as the class-of-service bits) and the per-hop behavior (PHB) associated with the EXP bits. The mapping between the EXP bits and the PHB is static, rather than being signaled in RSVP.

Once a multiclass LSP is configured, traffic from all of the class types can:

- Follow the same path
- Be rerouted along the same path

- Be taken down at the same time

Class types must be configured consistently across the Differentiated Services domain, meaning the class type configuration must be consistent from router to router in the network.

You can unambiguously map a class type to a queue. On each node router, the CoS queue configuration for an interface translates to the available bandwidth for a particular class type on that link.

The combination of a class type and a priority level forms a traffic engineering class. The IGP can advertise up to eight traffic engineering classes for each link.

For more information about the EXP bits, see “Label Allocation” on page 26.

For more information about forwarding classes, see the *JUNOS Class of Service Configuration Guide*.

Establishing a Multiclass LSP on the Differentiated Services Domain

The following occurs when a multiclass LSP is established on the differentiated services domain:

1. The IGPs advertise how much unreserved bandwidth is available for the traffic engineering classes.
2. When calculating the path for a multiclass LSP, CSPF is used to ensure that the constraints are met for all the class types carried by the multiclass LSP (a set of constraints instead of a single constraint).
3. Once a path is found, RSVP signals the LSP using an RSVP object in the path message. At each node in the path, the available bandwidth for the class types is adjusted as the path is set up. The RSVP object is a hop-by-hop object. Multiclass LSPs cannot be established through routers that do not understand this object. Preventing routers that do not understand the RSVP object from carrying traffic helps to ensure consistency throughout the differentiated services domain by preventing the multiclass LSP from using a router that is incapable of supporting differentiated services.

By default, multiclass LSPs are signaled with setup priority 7 and holding priority 0. A multiclass LSP configured with these values cannot preempt another LSP at setup time and cannot be preempted.

It is possible to have both multiclass LSPs and regular LSPs configured at the same time on the same physical interfaces. For this type of heterogeneous environment, regular LSPs carry best-effort traffic by default. Traffic carried in the regular LSPs must have the correct EXP settings.

Configuring Routers for DiffServ-Aware Traffic Engineering

To configure DiffServ-aware traffic engineering, include the `diffserv-te` statement:

```
diffserv-te {
```

```

bandwidth-model {
    extended-mam;
    mam;
    rdm;
}
te-class-matrix {
    traffic-class {
        tnumber {
            priority priority;
            traffic-class cnumber priority priority;
        }
    }
}
}

```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls]
- [edit logical-systems *logical-system-name* protocols mpls]

You must include the **diffserv-te** statement in the configuration on all routers participating in the Differentiated Services domain. However, you are not required to configure the traffic engineering class matrix (by including the **te-class-matrix** statement at the [edit protocols mpls **diffserv-te**] or [edit logical-systems *logical-system-name* protocols mpls **diffserv-te**] hierarchy level).

To configure DiffServ-aware traffic engineering, complete the procedures in the following sections:

- Configuring the Bandwidth Model on page 116
- Configuring Traffic Engineering Classes on page 117
- Configuring Class of Service for Diffserv-Aware Traffic Engineering on page 119

Configuring the Bandwidth Model

You must configure a bandwidth model on all routers participating in the Differentiated Services domain. The bandwidth models available are MAM, extended MAM, and RDM:

- Maximum allocation bandwidth constraints model (MAM)—Defined in RFC 4125, *Maximum Allocation Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering*.
- Extended MAM—A proprietary bandwidth model that behaves much like standard MAM. If you configure multiclass LSPs, you must configure the extended MAM bandwidth model.
- Russian-dolls bandwidth allocation model (RDM)—Makes efficient use of bandwidth by allowing the class types to share bandwidth. RDM is defined in RFC 4127, *Russian Dolls Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering*.

To configure a bandwidth model, include the `bandwidth-model` statement and specify one of the bandwidth model options:

```
bandwidth-model {
  extended-mam;
  mam;
  rdm;
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls diffserv-te]
- [edit logical-systems *logical-system-name* protocols mpls diffserv-te]



NOTE: If you change the bandwidth model on an ingress router, all the LSPs enabled on the router are taken down and resigaled.

Configuring Traffic Engineering Classes

Configuring traffic engineering classes is optional. Table 4 on page 117 shows the default values for everything in the traffic engineering class matrix. The default mapping is expressed in terms of the default forwarding classes defined in the CoS configuration.

Table 4: Default Values for the Traffic Engineering Class Matrix

Traffic Engineering Class	Class Type	Queue	Priority
te0	ct0	0	7
te1	ct1	1	7
te2	ct2	2	7
te3	ct3	3	7
te4	ct0	0	0
te5	ct1	1	0
te6	ct2	2	0
te7	ct3	3	0

If you want to override the default mappings, you can configure traffic engineering classes 0 through 7. For each traffic engineering class, you configure a class type (or queue) from 0 through 3. For each class type, you configure a priority from 0 through 7.

To configure traffic engineering classes explicitly, include the `te-class-matrix` statement:

```

te-class-matrix {
  tenumber {
    priority priority;
    traffic-class {
      ctnumber priority priority;
    }
  }
}

```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls diffserv-te]
- [edit logical-systems *logical-system-name* protocols mpls diffserv-te]

The following example shows how to configure traffic engineering class **te0** with class type **ct1** and a priority of 4:

```

[edit protocols mpls diffserv-te]
te-class-matrix {
  te0 traffic-class ct1 priority 4;
}

```



NOTE: If you explicitly configure a value for one of the traffic engineering classes, all the default values in the traffic engineering class matrix are dropped.

When you explicitly configure traffic engineering classes, you must also configure a bandwidth model; otherwise, the configuration commit operation fails. See “Configuring the Bandwidth Model” on page 116.

Requirements and Limitations for the Traffic Engineering Class Matrix

When you configure a traffic engineering class matrix, be aware of the following requirements and limitations:

- A mapping configuration is local and affects only the router on which it is configured. It does not affect other systems participating in the differentiated services domain. However, for a Differentiated Services domain to function properly, you need to configure the same traffic engineering class matrix on all the routers participating in the same domain.
- When explicitly configuring traffic engineering classes, you must configure the classes in sequence (**te0**, **te1**, **te2**, **te3**, and so on); otherwise, the configuration commit operation fails.

The first traffic engineering class you configure must be **te0**; otherwise, the configuration commit operation fails.

Configuring Class of Service for Diffserv-Aware Traffic Engineering

To configure Diffserv-aware traffic engineering, you must also configure class of service. The following example illustrates a class-of-service configuration that would allocate 25 percent of the link bandwidth to each class:

```
class-of-service {
  interfaces {
    all {
      scheduler-map simple-map;
    }
  }
  scheduler-maps {
    simple-map {
      forwarding-class assured-forwarding scheduler simple_sched;
      forwarding-class best-effort scheduler simple_sched;
      forwarding-class network-control scheduler simple_sched;
      forwarding-class expedited-forwarding scheduler simple_sched;
    }
  }
  schedulers {
    simple_sched {
      transmit-rate percent 25;
      buffer-size percent 25;
    }
  }
}
```

For more information on how to configure class of service, see the *JUNOS Class of Service Configuration Guide*.

Bandwidth Oversubscription Overview

LSPs are established with bandwidth reservations configured for the maximum amount of traffic you expect to traverse the LSP. Not all LSPs carry the maximum amount of traffic over their links at all times. For example, even if the bandwidth for link A has been completely reserved, actual bandwidth might still be available but not currently in use. This excess bandwidth can be used by allowing other LSPs to also use link A, oversubscribing the link. You can oversubscribe the bandwidth configured for individual class types or specify a single value for all of the class types using an interface.

You can use oversubscription to take advantage of the statistical nature of traffic patterns and to permit higher utilization of links.

The following examples describe how you might use bandwidth oversubscription and undersubscription:

- Use oversubscription on class types where peak periods of traffic do not coincide in time.

- Use oversubscription of class types carrying best-effort traffic. You take the risk of temporarily delaying or dropping traffic in exchange for making better utilization of network resources.
- Give different degrees of oversubscription or undersubscription of traffic for the different class types. For instance, you configure the subscription for classes of traffic as follows:
 - Best effort—`ct0 1000`
 - Voice—`ct3 1`

When you undersubscribe a class type for a multiclass LSP, the total demand of all RSVP sessions is always less than the actual capacity of the class type. You can use undersubscription to limit the utilization of a class type.

The bandwidth oversubscription calculation occurs on the local router only. Because no signaling or other interaction is required from other routers in the network, the feature can be enabled on individual routers without being enabled or available on other routers which might not support this feature. Neighboring routers do not need to know about the oversubscription calculation, they rely on the IGP.

The following sections describe the types of bandwidth oversubscription available in the JUNOS Software:

- LSP Size Oversubscription on page 120
- Link Size Oversubscription on page 120
- Class Type Oversubscription and Local Oversubscription Multipliers on page 121

LSP Size Oversubscription

For LSP size oversubscription, you simply configure less bandwidth than the peak rate expected for the LSP. You also might need to adjust the configuration for automatic policers. Automatic policers manage the traffic assigned to an LSP, ensuring that it does not exceed the configured bandwidth values. LSP size oversubscription requires that the LSP can exceed its configured bandwidth allocation.

Policing is still possible. However, the policer must be manually configured to account for the maximum bandwidth planned for the LSP, rather than for the configured value.

Link Size Oversubscription

You can increase the maximum reservable bandwidth on the link and use the inflated values for bandwidth accounting. Use the **subscription** statement to oversubscribe the link. The configured value is applied to all class type bandwidth allocations on the link. For more information about link size oversubscription, see “Configuring the Bandwidth Subscription Percentage for LSPs” on page 124.

Class Type Oversubscription and Local Oversubscription Multipliers

Local oversubscription multipliers (LOMs) allow different oversubscription values for different class types. LOMs are useful for networks where the oversubscription ratio needs to be configured differently on different links and where oversubscription values are required for different classes. You might use this feature to oversubscribe class types handling best-effort traffic, but use no oversubscription for class types handling voice traffic. An LOM is calculated locally on the router. No information related to an LOM is signaled to other routers in the network.

An LOM is configurable on each link and for each class type. The per-class type LOM allows you to increase or decrease the oversubscription ratio. The per-class-type LOM is factored into all local bandwidth accounting for admission control and IGP advertisement of unreserved bandwidths.

The LOM calculation is tied to the bandwidth model (MAM, extended MAM, and Russian dolls) used, because the effect of oversubscription across class types must be accounted for accurately.



NOTE: All LOM calculations are performed by the JUNOS Software and require no user intervention.

The formulas related to the oversubscription of class types are described in the following sections:

- Class Type Bandwidth and the LOM on page 121
 - LOM Calculation for the MAM and Extended MAM Bandwidth Models on page 122
 - LOM Calculation for the Russian Dolls Bandwidth Model on page 122
 - Example: LOM Calculation on page 122
-

Class Type Bandwidth and the LOM

The following formula expresses the relationship between the bandwidth of the class type and the LOM. The normalized bandwidth of the class type (N_B) is equal to the reserved bandwidth of the class type (R_B) divided by the LOM of the class type (L_C):

$$N_B = R_B / L_C$$

When calculating available bandwidth, you need to subtract the normalized bandwidth from the relevant bandwidth constraint.



NOTE: When using an LOM, values advertised for the available bandwidth might be larger than the bandwidth constraint values. However, the values advertised in the maximum link bandwidth advertisement are not affected by local oversubscription.

LOM Calculation for the MAM and Extended MAM Bandwidth Models

The following formulas show how the LOM is calculated for the MAM and extended MAM bandwidth models.

$$\text{Unreserved TE-Class}(i) = \text{LOMc} \times [\text{BCc} - \text{SUM} (\text{Normalized} (\text{CTc}, q))] \text{ for } q \leq p$$

Or

$$\text{Unreserved TE-Class}(i) = (\text{LOMc} \times \text{BCc}) - \text{SUM} (\text{Reserved} (\text{CTc}, q)) \text{ for } q \leq p$$

where:

- LOMc—LOM for class type *c*.
- BCc—Bandwidth constraint for class type *c*.
- CTc—Class type *c*.
- TE-Class(*i*) <--> (CTc , preemption *p*) in the configured TE-Class mapping.

LOM Calculation for the Russian Dolls Bandwidth Model

The following formulas show how the LOM is calculated for the Russian dolls bandwidth model:

$$\begin{aligned} \text{Unreserved TE-Class} (i) = & \text{LOMc} \times \text{MIN} [\\ & [\text{BCc} - \text{SUM} (\text{Normalized} (\text{CTb}, q))] \text{ for } q \leq p \text{ and } c \leq b \leq 7, \\ & \dots \\ & [\text{BC0} - \text{SUM} (\text{Normalized} (\text{CTb}, q))] \text{ for } q \leq p \text{ and } 0 \leq b \leq 7, \\ &] \end{aligned}$$

where:

- LOMc—LOM for class type *c*.
- BCc—Bandwidth constraint for class type *c*.
- TE-Class(*i*) <--> (CTc , preemption *p*) in the configured TE-Class mapping.

Note that the impact of an LSP on the unreserved bandwidth of a class type does not depend only on the LOM for that class type—it also depends on the LOM for the class type of the LSP.

Example: LOM Calculation

The following example illustrates how an LOM calculation is made for four classes of traffic: **ct0**, **ct1**, **ct2**, and **ct3**.

The class types have been assigned the following values:

$ct0 = 40$
 $ct1 = 30$
 $ct2 = 20$
 $ct3 = 10$

These class type values yield the following bandwidth constraints:

$BC0 = (ct3 + ct2 + ct1 + ct0) = 100$
 $BC1 = (ct3 + ct2 + ct1) = 60$
 $BC2 = (ct3 + ct2) = 30$
 $BC3 = (ct3) = 10$

LSPs from class type **ct0** can take up to 100 percent of bandwidth on the link. LSPs from class type **ct1** can take up to 60 percent of the bandwidth on the link, and so on.

If you assume for this example that the class types have the following LOM values:

$LOM(ct0) = 8$
 $LOM(ct1) = 4$
 $LOM(ct2) = 2$
 $LOM(ct3) = 1$

In the absence of any other reservation, LSPs from class type **ct0** can take up to 800 percent of the available bandwidth ($8 \times 100 = 800$). In the absence of any other reservation, LSPs from class type **ct1** can take up to 240 percent of the available bandwidth ($4 \times 60 = 240$). and so on.

The maximum amount of bandwidth that can be reserved is:

$ct0 = LOM(ct0) \times BC0 = 800$
 $ct1 = LOM(ct1) \times BC1 = 240$
 $ct2 = LOM(ct2) \times BC2 = 60$
 $ct3 = LOM(ct3) \times BC3 = 10$

For the undersubscribed class type **ct3**, the maximum reservable bandwidth is the same as the bandwidth constraint. For the overbooked class types, these values are not the values of the bandwidth constraint-taking into account the oversubscription for each class type separately. The oversubscription per class type in the sum is not taken into account because ultimately the entire bandwidth constraint can be filled with the bandwidth reservation of just one class type, so you have to account for that class type's bandwidth oversubscription only.

When calculating the available bandwidth for **CTc**, you need to express reservations from other classes as if they were from **CTc**. The reservation from class **ctx** is normalized with the LOM of **ctx**, but it is then multiplied by the LOM of **CTc**.

For the previous example, assume that **LSP1** has class type **ct3** configured with bandwidth of 10 and a priority of 0.

The values for the reservable bandwidth will be:

$ct0 = 8 \times (100 - 10) = 720$
 $ct1 = 4 \times \min((100-10), (60-10)) = 200$
 $ct2 = 2 \times \min((100-10), (60-10), (30-10)) = 40$

$$ct3 = 1 \times \min((100-10), (60-10), (30-10), (10-10)) = 0$$

These numbers can be rationalized as follows: the normalized reservation is 10 percent. If this bandwidth came from class type **ct0**, it would be equivalent to an overbooked reservation of 80 percent. You can see that 720 percent ($800 - 80 = 720$) of the bandwidth remains available for other LSPs.

Configuring the Bandwidth Subscription Percentage for LSPs

By default, RSVP allows all of a class type's bandwidth (100 percent) to be used for RSVP reservations. When you oversubscribe a class type for a multiclass LSP, the aggregate demand of all RSVP sessions is allowed to exceed the actual capacity of the class type.

If you want to oversubscribe or undersubscribe all of the class types on an interface using the same percentage bandwidth, configure the percentage using the **subscription** statement:

```
subscription percentage;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section.

To undersubscribe or oversubscribe the bandwidth for each class type, configure a percentage for each class type (**ct0**, **ct1**, **ct2**, and **ct3**) option for the **subscription** statement. When you oversubscribe a class type, an LOM is applied to calculate the actual bandwidth reserved. See "Class Type Oversubscription and Local Oversubscription Multipliers" on page 121 for more information.

```
subscription {
  ct0 percentage;
  ct1 percentage;
  ct2 percentage;
  ct3 percentage;
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section.

percentage is the percentage of class type bandwidth that RSVP allows to be used for reservations. It can be a value from 0 through 65,000 percent. If you specify a value greater than 100, you are oversubscribing the interface or class type.

The value you configure when you oversubscribe a class type is a percentage of the class type bandwidth that can actually be used. The default subscription value is 100 percent.

You can use the **subscription** statement to disable new RSVP sessions for one or more class types. If you configure a percentage of 0, no new sessions (including those with zero bandwidth requirements) are permitted for the class type.

Existing RSVP sessions are not affected by changing the subscription factor. To clear an existing session, issue the **clear rsvp session** command. For more information on

the clear `rsvp session` command, see the *JUNOS Routing Protocols and Policies Command Reference*.

Constraints on Configuring Bandwidth Subscription

Be aware of the following issues when configuring bandwidth subscription:

- If you configure bandwidth constraints at the `[edit class-of-service interface interface-name]` hierarchy level, they override any bandwidth configuration you specify at the `[edit protocols rsvp interface interface-name bandwidth]` hierarchy level for Diffserv-TE. Also note that either of the CoS or RSVP bandwidth constraints can override the interface hardware bandwidth constraints.
- If you configure a bandwidth subscription value for a specific interface that differs from the value configured for all interfaces (by including different values for the subscription statement at the `[edit protocols rsvp interface interface-name]` and `[edit protocols rsvp interface all]` hierarchy levels), the interface-specific value is used for that interface.
- You can configure subscription for each class type only if you also configure a bandwidth model. If no bandwidth model is configured, the commit operation fails with the following error message:

```
user@host# commit check

[edit protocols rsvp interface all]
'subscription'
RSVP: Must have a diffserv-te bandwidth model configured when configuring
subscription per traffic class.
error: configuration check-out failed
```

- You cannot include the `subscription` statement both in the configuration for a specific class type and the configuration for the entire interface. The commit operation fails with the following error message:

```
user@host# commit check

[edit protocols rsvp interface all]
'subscription'
RSVP: Cannot configure both link subscription and per traffic class
subscription.
error: configuration check-out failed
```

Configuring LSPs for DiffServ-Aware Traffic Engineering

You must configure the Differentiated Services domain (see “Configuring Routers for DiffServ-Aware Traffic Engineering” on page 115) before you can enable DiffServ-aware traffic engineering for LSPs. The Differentiated Services domain provides the underlying class types and corresponding traffic engineering classes that you reference in the LSP configuration. The traffic engineering classes must be configured consistently on each router participating in the Differentiated Services domain for the LSP to function properly.



NOTE: You must configure either MAM or RDM as the bandwidth model when you configure DiffServ-aware traffic engineering for LSPs. See “Configuring the Bandwidth Model” on page 116.

The actual data transmitted over this Differentiated Services domain is carried by an LSP. Each LSP relies on the EXP bits of the MPLS packets to enable DiffServ-aware traffic engineering. Each LSP can carry traffic for a single class type.

All the routers participating in the LSP must be Juniper Networks routers running JUNOS Release 6.3 or later. The network can include routers from other vendors and Juniper Networks routers running earlier versions of the JUNOS Software. However, the DiffServ-aware traffic engineering LSP cannot traverse these routers.



NOTE: You cannot simultaneously configure multiclass LSPs and DiffServ-aware traffic engineering LSPs on the same router.

To enable DiffServ-aware traffic engineering for LSPs, you need to configure the following:

- Configuring Class of Service for the Interfaces on page 126
- Configuring IGP on page 126
- Configuring Traffic-Engineered LSPs on page 127
- Configuring Policing for LSPs on page 127
- Configuring Fast Reroute for Traffic-Engineered LSPs on page 128

Configuring Class of Service for the Interfaces

The existing class-of-service (CoS) infrastructure ensures that traffic that is consistently marked receives the scheduling guarantees for its class. The classification, marking, and scheduling necessary to accomplish this are configured using the existing JUNOS Software CoS features.



NOTE: ATM interfaces are not supported.

For information on how to configure CoS, see the *JUNOS Class of Service Configuration Guide*.

Configuring IGP

You can configure either IS-IS or OSPF as the IGP. The IS-IS and OSPF configurations for routers supporting LSPs are standard. For information on how to configure these protocols, see the *JUNOS Routing Protocols Configuration Guide*.

Configuring Traffic-Engineered LSPs

You configure an LSP by using the standard LSP configuration statements and procedures. To configure DiffServ-aware traffic engineering for the LSP, specify a class type bandwidth constraint by including the **bandwidth** statement:

```
label-switched-path lsp-name {
    bandwidth {
        ctnumber bps;
    }
}
```

For a list of hierarchy levels at which you can include the **bandwidth** statement, see the statement summary sections for this statement.

If you do not specify a bandwidth for a class type, **ct0** is automatically specified as the queue for the LSP. You can configure only one class type for each LSP, unlike multiclass LSPs.

The class type statements specify bandwidth (in bits per second) for the following classes:

- **ct0**—Bandwidth reserved for class 0
- **ct1**—Bandwidth reserved for class 1
- **ct2**—Bandwidth reserved for class 2
- **ct3**—Bandwidth reserved for class 3

You can configure setup and holding priorities for an LSP, but the following restrictions apply:

- The combination of class and priority must be one of the configured traffic engineering classes. The default setup priority is 7 and the default holding priority is 0.
- Configuring an invalid combination of class type and priority causes the commit operation to fail.
- Automatic bandwidth allocation is not supported. If you configure automatic bandwidth allocation, the commit operation fails.
- LSPs configured with the **bandwidth** statement but without specifying a class type use the default class type **ct0**.
- For migration issues, see Internet draft [draft-ietf-tewg-diff-te-proto-07.txt](#).

Configuring Policing for LSPs

Policing allows you to control the amount of traffic forwarded through a particular LSP. Policing helps to ensure that the amount of traffic forwarded through an LSP never exceeds the requested bandwidth allocation. You can configure multiple policers for each LSP.

For information on how to configure a policer for an LSP, see “Configuring Policers for LSPs” on page 160.

Configuring Fast Reroute for Traffic-Engineered LSPs

You can configure fast reroute for traffic engineered LSPs (LSPs carrying a single class of traffic). It is also possible to reserve bandwidth on the detour path for the class of traffic when fast reroute is enabled. The same class type number is used for both the traffic engineered LSP and its detour.

If you configure the router to reserve bandwidth for the detour path, a check is made to ensure that the link is capable of handling DiffServ-aware traffic engineering and for CoS capability before accepting it as a potential detour path. Unsupported links are not used.

You can configure the amount of bandwidth to reserve for detours using either the **bandwidth** statement or the **bandwidth-percent** statement. You can only configure one these statements at a time. If you do not configure either the **bandwidth** statement or the **bandwidth-percent** statement, the default setting is to not reserve bandwidth for the detour path (the bandwidth guarantee will be lost if traffic is switched to the detour).

When you configure the **bandwidth** statement, you can specify the specific amount of bandwidth (in bits per second [bps]) you want to reserve for the detour path. For information, see “Configuring Fast Reroute” on page 78.

When you configure the **bandwidth-percent** statement, the detour path bandwidth is computed by multiplying it to the bandwidth configured for the main traffic engineered LSP. For information on how to configure the bandwidth for a traffic engineered LSP, see “Configuring Traffic-Engineered LSPs” on page 127.

To configure the percent of bandwidth used by the detour path based on the bandwidth of the protected path, include the **bandwidth-percent** statement:

```
bandwidth-percent percentage;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls label-switched-path *lsp-name* fast-reroute]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* fast-reroute]

Configuring Multiclass LSPs

A multiclass LSP is an LSP configured to reserve bandwidth for multiple class types and also carries the traffic for these class types. The differentiated service behavior is determined by the EXP bits.

You must configure the Differentiated Services domain (see “Configuring Routers for DiffServ-Aware Traffic Engineering” on page 115) before you can enable a multiclass LSP. The Differentiated Services domain provides the underlying class types and corresponding traffic engineering classes that you reference in a multiclass LSP

configuration. The traffic engineering classes must be configured consistently on each router participating in the Differentiated Services domain for the multiclass LSP to function properly.



NOTE: You must configure extended MAM as the bandwidth model when you configure multiclass LSPs. See “Configuring the Bandwidth Model” on page 116.

All the routers participating in a multiclass LSP must be Juniper Networks routers running JUNOS Release 6.2 or later. The network can include routers from other vendors and Juniper Networks routers running earlier versions of the JUNOS Software. However, the multiclass LSP cannot traverse these routers.

To enable multiclass LSPs, you need to configure the following:

- Configuring Class of Service for the Interfaces on page 129
- Configuring the IGP on page 129
- Configuring Class-Type Bandwidth Constraints for Multiclass LSPs on page 129
- Configuring Policing for Multiclass LSPs on page 131
- Configuring Fast Reroute for Multiclass LSPs on page 131

Configuring Class of Service for the Interfaces

The existing class-of-service infrastructure ensures that traffic that is consistently marked receives the scheduling guarantees for its class. The classification, marking, and scheduling necessary to consistently mark traffic are configured with the existing JUNOS Software CoS features.



NOTE: ATM interfaces are not supported.

For information on how to configure CoS, see the *JUNOS Class of Service Configuration Guide*.

Configuring the IGP

You can configure either IS-IS or OSPF. The IS-IS and OSPF configurations for routers supporting multiclass LSPs are standard. For information about how to configure these protocols, see the *JUNOS Routing Protocols Configuration Guide*.

Configuring Class-Type Bandwidth Constraints for Multiclass LSPs

You configure a multiclass LSP by using the standard LSP configuration statements and procedures. To configure an LSP as a multiclass LSP, specify the class type bandwidth constraints by including the **bandwidth** statement:

```
bandwidth {
  ct0 bps;
  ct1 bps;
```

```

    ct2 bps;
    ct3 bps;
}

```

For a list of hierarchy levels at which you can include the **bandwidth** statement, see the statement summary sections for these statements.

The class type statements specify bandwidth (in bits per second) for the following classes:

- **ct0**—Bandwidth reserved for class 0
- **ct1**—Bandwidth reserved for class 1
- **ct2**—Bandwidth reserved for class 2
- **ct3**—Bandwidth reserved for class 3

For example, to configure 50 megabytes of bandwidth for class type 1 and 30 megabytes of bandwidth for class type 2, include the **bandwidth** statement as follows:

```

[edit protocols mpls]
label-switched-path traffic-class {
    bandwidth {
        ct1 50M;
        ct2 30M;
    }
}

```

You cannot configure a bandwidth for a class type and also configure a bandwidth at the `[edit protocols mpls label-switched-path lsp-name bandwidth]` hierarchy level. For example, the following configuration cannot be committed:

```

[edit protocols mpls]
label-switched-path traffic-class {
    bandwidth {
        20M;
        ct1 10M;
    }
}

```

You can configure setup and holding priorities for a multiclass LSP, but the following restrictions apply:

- The setup and holding priorities apply to all classes for which bandwidth is requested.
- The combination of class and priority must be one of the configured traffic engineering classes. The default traffic engineering class configuration results in multiclass LSPs that cannot preempt and cannot be preempted. The default setup priority is 7 and the default holding priority is 0.
- Configuring an invalid combination of class type and priority causes the commit operation to fail.

- Automatic bandwidth allocation is not supported for multiclass LSPs. If you configure automatic bandwidth allocation, the commit operation fails.
- LSPs configured with the **bandwidth** statement but without specifying a class type use the default class type **ct0**.

Configuring Policing for Multiclass LSPs

Policing allows you to control the amount of traffic forwarded through a particular multiclass LSP. Policing helps to ensure that the amount of traffic forwarded through an LSP never exceeds the requested bandwidth allocation. You can configure multiple policers for each multiclass LSP. You can also enable automatic policing for multiclass LSPs.

For information on how to configure a policer for a multiclass LSP, see “Configuring Policers for LSPs” on page 160 and “Configuring Automatic Policers” on page 162.

Configuring Fast Reroute for Multiclass LSPs

You can enable fast reroute for multiclass LSPs. The bandwidth guarantees for the class types can be carried over to the detour path in case the primary path of the multiclass LSP fails. The same traffic class types configured for the primary multiclass LSP are also signaled for the detour LSP.

The bandwidth guarantee for the detour path is a percentage of the bandwidth configured for the class types of the primary path. For example, you configure a value of 50 percent for the detour path and the protected LSP carries traffic for class types CT0 through CT3. The detour path is signaled with the same class types (CT0 through CT3) but with 50 percent of the bandwidth configured for the protected LSP.

If you configure the router to reserve bandwidth for the detour path, a check is made to ensure that the link is capable of handling DiffServ-aware traffic engineering, that all of the traffic class types needed are available, and for CoS capability before accepting it as a potential detour path. Unsupported links are not used.

The bandwidth percentage for fast reroute is signaled from the ingress router to the egress router. All of the intermediate devices must complete their own CSPF computations and signaling.

When you configure the **bandwidth-percent** statement, the detour path bandwidth is computed by multiplying it to the bandwidth configured for the primary multiclass LSP. For information on how to configure the bandwidth for the multiclass LSP, see “Configuring Traffic-Engineered LSPs” on page 127.

To configure the percent of bandwidth used by the detour path based on the bandwidth of the protected path, include the **bandwidth-percent** statement:

```
bandwidth-percent percentage;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls label-switched-path *lsp-name* fast-reroute]

- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* fast-reroute]

Chapter 7

Static and Explicit-Path LSP Configuration Guidelines

The following sections describe how to configure static and explicit-path label-switched paths (LSPs):

- Configuring Static LSPs on page 133
- Configuring Explicit-Path LSPs on page 140

Configuring Static LSPs

To configure static LSPs, configure the ingress router and each router along the path up to and including the egress router.

For the ingress router, configure which packets to tag (based on the packet's IP destination address), the next router in the LSP, and the tag to apply to the packet. Manually assigned labels can have values from 1,000,000 through 1,048,575. Optionally, you can apply preference and class-of-service (CoS) values to the packets.

For the intermediate routers in the path, configure the next router in the path and the tag to apply to the packet. Again, you can optionally apply preference and CoS values to the packets.

For the egress router, you generally just remove the label and continue forwarding the packet to the next hop. However, if the previous router removed the label, the egress router examines the packet's IP header and forwards the packet toward its IP destination.

To configure static MPLS, perform the following tasks:

- Configuring the Ingress Router for Static LSPs on page 133
- Configuring the Intermediate and Egress Routers for Static LSPs on page 136
- Configuring Static Unicast Routes for Point-to-Multipoint LSPs on page 138

Configuring the Ingress Router for Static LSPs

The ingress router checks the IP address in the incoming packet's destination address field and, if it finds a match in the routing table, applies the label associated with that address to the packets. The label has forwarding information associated with it, including the address of the next-hop router, and the route preference and CoS values.

To configure static LSPs on the ingress router, include the **static-path** statement:

```
static-path inet {
  prefix {
    class-of-service value;
    double-push bottom-label top-label;
    next-hop (address | interface-name | address/interface-name);
    preference preference;
    push out-label;
    triple-push bottom-label middle-label top-label;
  }
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls]
- [edit logical-systems *logical-system-name* protocols mpls]

The **next-hop** and **push** statements are required; the other statements are optional.

Each **static-path** statement consists of the following parts:

- Criteria to use to analyze an incoming packet:
 - The **inet** option creates an LSP that handles IPv4 packets. All static MPLS routes created using the **inet** option are installed in the default IPv4 routing table (**inet.0**), and the creating protocol is identified as **static**. This process is no different from creating static IPv4 routes at the [edit routing-options static] hierarchy level.
 - In the **prefix** option, you configure the IP destination address to check when incoming packets are analyzed. If the address matches, the specified label, **out-label**, is assigned to the packet, and the packet enters an LSP. Each prefix that you specify is installed as a static route in the routing table. You can specify one or more **prefix** statements at the [edit protocols mpls static-path] hierarchy level.
- The **next-hop** statement, which supplies the IP address of the next hop to the destination. You can specify this as the IP address of the next hop, the interface name (for point-to-point interfaces only), or as **address/interface-name** to specify an IP address on an operational interface. When the next hop is on a directly attached interface, the route is installed in the routing table. You cannot configure a LAN or nonbroadcast multiaccess (NBMA) interface as a next-hop interface.
- Label properties applied to the packet in the LSP, which are defined by the following statements:
 - **push out-label**—Push one more label on top of the stack.
 - **double-push bottom-label top-label**—Push two more labels on top of the stack.
 - **triple-push bottom-label middle-label top-label**—Push three more labels on top of the stack.

A label is a 20-bit integer, so it can be a number from 0 through 1,048,575 (2²⁰– 1). Labels 0 through 999,999 are for internal use. Labels 1,000,000

through 1,048,575 are unassigned by the JUNOS Software and are available for static LSPs. When you configure static LSPs, you can use only this range of labels.

- Preference of this route (defined by the `preference preference` statement).
- CoS value to apply to the packet (defined by the `class-of-service cos-value` statement).


To determine whether a static ingress route is installed, use the command `show route table inet.0 protocol static`. The following is sample output. The `push` keyword identifies that a label is to be added in front of an IP packet.

```
10.0.0.0          *[Static/5] 00:01:48
> to 11.1.1.1 via so-0/0/0, push 1000123
```

Example: Configuring the Ingress Router

Configure the ingress router for a static LSP that consists of three routers (see Figure 20 on page 135).

Figure 20: Static MPLS Configuration



For packets addressed to 10.0.0.0, assign label 1000123 and transmit them to the next-hop router at 11.1.1.1:

```
[edit]
interfaces {
  so-0/0/0 {
    unit 0 {
      family mpls;
    }
  }
}
protocols {
  mpls {
    static-path inet {
      10.0.0.0 {
        next-hop 11.1.1.1;
        push 1000123;
      }
    }
  }
}
```

```

    }
    interface so-0/0/0;
  }
}

```

To determine whether the static ingress route is installed, use the following command:

```
user@host> show route table inet.0 protocol static
```

The following is a sample of the output. The **push 1000123** keyword identifies the route.

```

10.0.0.0/8          *[Static/5] 00:01:48
> to 11.1.1.1 via so-0/0/0, push 1000123

```

Configuring the Intermediate and Egress Routers for Static LSPs

Intermediate and egress routers perform similar functions—they modify the label that has been applied to a packet. An intermediate router can change the label. An egress router removes the label (if the packet still contains a label) and continues forwarding the packet to its destination.

To configure static LSPs on intermediate and egress routers, include the **interface** statement:

```

interface (interface-name | all) {
  disable;
  admin-group [ group-names ];
  label-map in-label {
    class-of-service cos-value;
    next-hop (address | interface-name | address/interface-name) | (discard | reject);
    (pop | swap <out-label>);
    preference preference;
    swap-push swap-label push-label;
  }
}

```

You can include these statements at the following hierarchy levels:

- [edit protocols mpls]
- [edit logical-systems *logical-system-name* protocols mpls]

For the **label-map** statement configuration, the **next-hop** | (**reject** | **discard**) and **pop** | **swap** statements are required. The remaining statements are optional.

Each statement within the **interface** statement consists of the following parts:

- Criteria to use to analyze the labeled packet. Two criteria are used: the interface on which the packet was received (specified in the opening **interface** statement itself) and the packet's label (specified in the **label-map** statement).
- The **next-hop** statement, which supplies the IP address of the next hop to the destination. The address is specified as the IP address of the next hop, or the interface name (for point-to-point interfaces only), or **address** and **interface-name**

to specify an IP address on an operational interface. When the specified next hop is on a directly attached interface, this route is installed in the routing table. You cannot configure a LAN or NBMA interface as a next-hop interface.

- Operation to perform on the labeled packet:
 - For egress routers, remove the packet's label altogether (**pop**).
 - For intermediate routers only, exchange the label for another label (**swap out-label**).
 - Discard the packet, sending an ICMP unreachable message to the packet's originator (**reject**).
 - Discard the packet without sending an ICMP unreachable message to the packet's originator (**discard**).
- Label properties to apply to the packet (all are optional):
 - Preference value for this route (**preference preference**).
 - For intermediate routers only, the CoS value to apply to the packet (**class-of-service cos-value**).

You can specify any number of **label-map** statements at the **[edit protocols mpls interface interface-name]** hierarchy level.

The static routes are installed in the default MPLS routing table, **mpls.0**, and the creating protocol is identified as **static**. To verify that a static route is properly installed, use the command **show route table mpls.0 protocol static**. The following is an example of the output:

```
1000123          *[Static/5] 00:00:38
> to 12.2.2.2 via so-5/0/0.0, swap 1000456
```

Example: Configuring an Intermediate Router

For packets labeled 1000123 arriving on interface **so-0/0/0**, assign the label 1000456, and transmit them to the next-hop router at 12.2.2.2:

```
[edit]
interfaces {
  so-0/0/0 {
    unit 0 {
      family mpls;
    }
  }
}
protocols {
  mpls {
    interface so-0/0/0 {
      label-map 1000123 {
        next-hop 12.2.2.2;
        swap 1000456;
      }
    }
  }
}
```

```
}
}
```

To determine whether the static intermediate route is installed, use the following command:

```
user@host> show route table mpls.0 protocol static
```

The following is a sample of the output. The `swap 1000456` keyword identifies the route.

```
1000123      *[Static/5] 00:01:48
> to 12.2.2.2 via so-0/0/0, swap 1000456
```

Example: Configuring an Egress Router

For packets labeled `1000456` arriving on interface `so-0/0/0`, remove the label and transmit the packets to the next-hop router at `13.3.3.3`:

```
[edit]
interfaces {
  so-0/0/0 {
    unit 0 {
      family mpls;
    }
  }
}
protocols {
  mpls {
    interface so-0/0/0 {
      label-map 1000456 {
        next-hop 13.3.3.3;
        pop;
      }
    }
  }
}
```

To determine whether the static egress route is installed, use the following command:

```
user@host> show route table mpls.0 protocol static
```

The following is a sample of the output. The `pop` keyword identifies the egress route.

```
1000456      *[Static/5] 00:01:48
> to 13.3.3.3 via so-0/0/0, pop
```

Configuring Static Unicast Routes for Point-to-Multipoint LSPs

You can configure a static unicast IP route with a point-to-multipoint LSP as the next hop. For more information on point-to-multipoint LSPs, see “Point-to-Multipoint LSPs” on page 48, “Point-to-Multipoint LSP Configuration Guidelines” on page 141, and “Configuring CCC Switching for Point-to-Multipoint LSPs” on page 424.

To configure a static unicast route for a point-to-multipoint LSP, complete the following steps:

1. On the ingress PE router, configure a static IP unicast route with the point-to-multipoint LSP name as the next hop by including the `p2mp-lsp-next-hop` statement:

```
p2mp-lsp-next-hop point-to-multipoint-lsp-next-hop;
```

You can include this statement at the following hierarchy levels:

- `[edit routing-options static route route-name]`
- `[edit logical-systems logical-system-name routing-options static route route-name]`

2. On the egress PE router, configure a static IP unicast route with the same destination address configured in Step 1 (the address configured at the `[edit routing-options static route]` hierarchy level) by including the `next-hop` statement:

```
next-hop address;
```

You can include this statement at the following hierarchy levels:

- `[edit routing-options static route route-name]`
- `[edit logical-systems logical-system-name routing-options static route route-name]`



NOTE: CCC and static routes cannot use the same point-to-multipoint LSP.

For more information on static routes, see the *JUNOS Routing Protocols Configuration Guide*.

The following `show route` command output displays a unicast static route pointing to a point-to-multipoint LSP on the ingress PE router where the LSP has two branch next hops:

```
user@host> show route 5.5.5.5 detail
inet.0: 29 destinations, 30 routes (28 active, 0 holddown, 1 hidden)
5.5.5.5/32 (1 entry, 1 announced)
  *Static Preference: 5
    Next hop type: Flood
    Next hop: via so-0/3/2.0 weight 1
    Label operation: Push 100000
    Next hop: via t1-0/1/1.0 weight 1
    Label operation: Push 100064
    State: <Active Int Ext>
    Local AS: 10458
    Age: 2:41:15
    Task: RT
    Announcement bits (2): 0-KRT 3-BGP.0.0.0.0+179
    AS path: I
```

Configuring Explicit-Path LSPs

If you disable constrained-path label-switched path (LSP) computation, as described in “Disabling Constrained-Path LSP Computation” on page 95, you can configure LSPs manually or allow the LSPs to follow the IGP path.

When explicit-path LSPs are configured, the LSP is established along the path you specified. If the path is topologically not feasible, either because the network is partitioned or insufficient resources are available along some parts of the path, the LSP will fail. No alternative paths can be used. If the setup succeeds, the LSP stays on the defined path indefinitely.

To configure an explicit-path LSP, follow these steps:

1. Configure the path information in a named path, as described in “Creating Named Paths” on page 52. To configure complete path information, specify every router hop between the ingress and egress routers, preferably using the **strict** attribute. To configure incomplete path information, specify only a subset of router hops, using the **loose** attribute in places where the path is incomplete.

For incomplete paths, the MPLS routers complete the path by querying the local routing table. This query is done on a hop-by-hop basis, and each router can figure out only enough information to reach the next explicit hop. It might be necessary to traverse a number of routers to reach the next (loose) explicit hop.

Configuring incomplete path information creates portions of the path that depend on the current routing table, and this portion of the path can reroute itself as the topology changes. Therefore, an explicit-path LSP that contains incomplete path information is not completely fixed. These types of LSPs have only a limited ability to repair themselves, and they tend to create loops or flaps depending on the contents of the local routing table.

2. To configure the LSP and point it to the named path, use either the **primary** or **secondary** statement, as described in “Configuring Primary and Secondary LSPs” on page 75.
3. Disable constrained-path LSP computation by including the **no-cspf** statement either as part of the LSP or as part of a **primary** or **secondary** statement. For more information, see “Disabling Constrained-Path LSP Computation” on page 95.
4. Configure any other LSP properties.

Using explicit-path LSPs has the following drawbacks:

- More configuration effort is required.
- Configured path information cannot take into account dynamic network bandwidth reservation, so the LSPs tend to fail when resources become depleted.
- When an explicit-path LSP fails, you might need to manually repair it.

Because of these limitations, we recommend that you use explicit-path LSPs only in controlled situations, such as to enforce an optimized LSP placement strategy resulting from computations with an offline simulation software package.

Chapter 8

Point-to-Multipoint LSP Configuration Guidelines

This chapter discusses the following topics:

- Configuring Primary and Branch LSPs for Point-to-Multipoint LSPs on page 141
- Example: Configuring a Point-to-Multipoint LSP on page 143
- Configuring Link Protection for Point-to-Multipoint LSPs on page 144
- Configuring Graceful Restart for Point-to-Multipoint LSPs on page 144
- Configuring a Multicast RPF Check Policy for Point-to-Multipoint LSPs on page 145
- Configuring Ingress PE Router Redundancy for Point-to-Multipoint LSPs on page 146
- Enabling Point-to-Point LSPs to Monitor Egress PE Routers on page 147
- Preserving Point-to-Multipoint LSP Functioning with Different JUNOS Software Releases on page 147

Configuring Primary and Branch LSPs for Point-to-Multipoint LSPs

A point-to-multipoint MPLS label-switched path (LSP) is an RSVP LSP with multiple destinations. By taking advantage of the MPLS packet replication capability of the network, point-to-multipoint LSPs avoid unnecessary packet replication at the ingress router. For more information about point-to-multipoint LSPs, see “Point-to-Multipoint LSPs” on page 48.

To configure a point-to-multipoint LSP, you need to configure the primary LSP from the ingress router and the branch LSPs that carry traffic to the egress routers, as described in the following sections:

- Configuring the Primary Point-to-Multipoint LSP on page 141
- Configuring a Branch LSP for Point-to-Multipoint LSPs on page 142

Configuring the Primary Point-to-Multipoint LSP

A point-to-multipoint LSP must have a configured primary point-to-multipoint LSP to carry traffic from the ingress router. The configuration of the primary point-to-multipoint LSP is similar to a signaled LSP. See “Configuring the Ingress Router for MPLS-Signaled LSPs” on page 52 for more information. In addition to the

conventional LSP configuration, you need to specify a path name for the primary point-to-multipoint LSP by including the **p2mp** statement:

```
p2mp p2mp-lsp-name;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls label-switched-path *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*]

You can enable the optimization timer for point-to-multipoint LSPs. See “Optimizing Signaled LSPs” on page 103 for more information.

Configuring a Branch LSP for Point-to-Multipoint LSPs

The primary point-to-multipoint LSP sends traffic to two or more branch LSPs carrying traffic to each of the egress provider edge (PE) routers. In the configuration for each of these branch LSPs, the point-to-multipoint LSP path name you specify must be identical to the path name configured for the primary point-to-multipoint LSP. See “Configuring the Primary Point-to-Multipoint LSP” on page 141 for more information.

To associate a branch LSP with the primary point-to-multipoint LSP, specify the point-to-multipoint LSP name by including the **p2mp** statement:

```
p2mp p2mp-lsp-name;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls label-switched-path *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*]



NOTE: Any change in any of the branch LSPs of a point-to-multipoint LSP, either due to a user action or an automatic adjustment made by the router, causes the primary and branch LSPs to be resigaled. The new point-to-multipoint LSP is signaled first before the old path is taken down.

The following sections describe how you can configure the branch LSP as a dynamically signaled path using Constrained Shortest Path First (CSPF), as a static path, or as a combination of dynamic and static paths:

- Configuring the Branch LSP as a Dynamic Path on page 142
- Configuring the Branch LSP as a Static Path on page 143

Configuring the Branch LSP as a Dynamic Path

By default, the branch LSP for a point-to-multipoint LSP is signaled dynamically using CSPF and requires no configuration.

When a point-to-multipoint LSP is changed, either by the addition or deletion of new destinations or by the recalculation of the path to existing destinations, certain nodes in the tree might receive data from more than one incoming interface. This can happen under the following conditions:

- Some of the branch LSPs to destinations are statically configured and might intersect with statically or dynamically calculated paths to other destinations.
- When a dynamically calculated path for a branch LSP results in a change of incoming interface for one of the nodes in the network, the older path is not immediately torn down after the new one has been signaled. This ensures that any data in transit relying on the older path can reach its destination. However, network traffic can potentially use either path to reach the destination.
- A faulty router at the ingress calculates the paths to two different branch destinations such that a different incoming interface is chosen for these branch LSPs on a router node common to these branch LSPs.

Configuring the Branch LSP as a Static Path

You can configure the branch LSP for a point-to-multipoint LSP as a static path. See “Configuring Static LSPs” on page 133 for more information.

Example: Configuring a Point-to-Multipoint LSP

Configure a point-to-multipoint LSP based on the topology shown in Figure 18 on page 49. There are four branch LSPs, each belonging to a single point-to-multipoint LSP called `p2mp-lsp-sample`.

```
[edit protocols mpls]
label-switched-path branch-LSP-to-PE2 {
  to 10.255.235.25;
  p2mp p2mp-lsp-sample;
  primary path1;
}
label-switched-path branch-LSP-to-PE2 {
  to 10.255.235.25;
  p2mp p2mp-lsp-sample;
  primary path2;
}
label-switched-path branch-LSP-to-PE3 {
  to 10.255.241.34;
  p2mp p2mp-lsp-sample;
  primary path3;
}
label-switched-path branch-LSP-to-CE4 {
  to 10.255.244.125;
  p2mp p2mp-lsp-sample;
  primary path4;
}
```

Configuring Link Protection for Point-to-Multipoint LSPs

Link protection helps to ensure that traffic going over a specific interface to a neighboring router can continue to reach this router if that interface fails. When link protection is configured for an interface and a point-to-multipoint LSP that traverses this interface, a bypass LSP is created that handles this traffic if the interface fails. The bypass LSP uses a different interface and path to reach the same destination.

To extend link protection to all of the paths used by a point-to-multipoint LSP, link protection must be configured on each router that each branch LSP traverses. If you enable link protection on a point-to-multipoint LSP, you must enable link protection on all of the branch LSPs.

The Internet draft *draft-ietf-mpls-rsvp-te-p2mp-01.txt*, *Extensions to RSVP-TE for Point to Multipoint TE LSPs*, describes link protection for point-to-multipoint LSPs.

To enable link protection on point-to-multipoint LSPs, complete the following steps:

1. Configure link protection on each branch LSP. To configure link protection, include the **link-protection (MPLS)** statement:

```
link-protection;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls label-switched-path *branch-lsp-name*]
 - [edit logical-systems *logical-system-name* protocols mpls label-switched-path *branch-lsp-name*]
2. Configure link protection for each RSVP interface on each router that the branch LSP traverses. For information on how to configure link protection on RSVP interfaces, see “Configuring Link Protection on Interfaces Used by LSPs” on page 282.

For more information on how to configure link protection, see “Configuring Node Protection or Link Protection for LSPs” on page 281.

Configuring Graceful Restart for Point-to-Multipoint LSPs

You can configure graceful restart on point-to-multipoint LSPs. Graceful restart allows a router undergoing a restart to inform its adjacent neighbors of its condition. The restarting router requests a grace period from the neighbor or peer, which can then cooperate with the restarting router. The restarting router can still forward MPLS traffic during the restart period; convergence in the network is not disrupted. The restart is not apparent to the rest of the network, and the restarting router is not removed from the network topology. RSVP graceful restart can be enabled on both transit routers and ingress routers.

To enable graceful restart on a router handling point-to-multipoint LSP traffic, include the **graceful-restart** statement:

`graceful-restart;`

You can include this statement at the following hierarchy levels:

- `[edit routing-options]`
- `[edit logical-systems logical-system-name routing-options]`

The graceful restart configuration for point-to-multipoint LSPs is identical to that of point-to-point LSPs. For more information on how to configure graceful restart, see “Configuring RSVP Graceful Restart” on page 289.

Configuring a Multicast RPF Check Policy for Point-to-Multipoint LSPs

You can control whether a reverse path forwarding (RPF) check is performed for a source and group entry before installing a route in the multicast forwarding cache. This makes it possible to use point-to-multipoint LSPs to distribute multicast traffic to PIM islands situated downstream from the egress routers of the point-to-multipoint LSPs.

By configuring the `rpf-check-policy` statement, you can disable RPF checks for a source and group pair. You would typically configure this statement on the egress routers of a point-to-multipoint LSP, because the interface receiving the multicast traffic on a point-to-multipoint LSP egress router might not always be the RPF interface.

You can also configure a routing policy to act upon a source and group pair. This policy behaves like an import policy, so if no policy term matches the input data, the default policy action is “acceptance.” An accept policy action enables RPF checks. A reject policy action (applied to all source and group pairs that are not accepted) disables RPF checks for the pair.

To configure a multicast RPF check policy for a point-to-multipoint LSP, specify the RPF check policy using the `rpf-check-policy` statement:

`rpf-check-policy policy;`

You can include this statement at the following hierarchy levels:

- `[edit routing-options multicast]`
- `[edit logical-systems logical-system-name routing-options multicast]`

You also must configure a policy for the multicast RPF check. You configure policies at the `[edit policy-options]` hierarchy level. For more information, see the *JUNOS Policy Framework Configuration Guide*.



NOTE: When you configure the `rpf-check-policy` statement, the JUNOS Software cannot perform RPF checks on incoming traffic and therefore cannot detect traffic arriving on the wrong interface. This might cause routing loops to form.

Example: Configuring Multicast RPF Check Policy for a Point-to-Multipoint LSP

Configure a policy to ensure that an RPF check is not performed for sources with prefix 128.83/16 or longer that belong to groups having a prefix of 228/8 or longer:

```
[edit]
policy-options {
  policy-statement rpf-sg-policy {
    from {
      route-filter 228.0.0.0/8 orlonger;
      source-address-filter 128.83.0.0/16 orlonger;
    }
    then {
      reject;
    }
  }
}
```

Configuring Ingress PE Router Redundancy for Point-to-Multipoint LSPs

You can configure one or more PE routers as part of a backup PE router group to enable ingress PE router redundancy. You accomplish this by configuring the IP addresses of the backup PE routers (at least one backup PE router is required) and the local IP address used by the local PE router.

You must also configure a full mesh of point-to-point LSPs between the primary and backup PE routers. You also need to configure BFD on these LSPs. See “Configuring BFD for RSVP-Signaled LSPs” on page 167 and “Configuring BFD for LDP LSPs” on page 351 for more information.

To configure ingress PE router redundancy for point-to-multipoint LSPs, include the **backup-pe-group** statement:

```
backup-pe-group pe-group-name {
  backups [addresses];
  local-address address;
}
```

For a list of hierarchy levels at which you can include these statements, see the statement summary sections for these statements.

After you configure the ingress PE router redundancy backup group, you must also apply the group to a static route on the PE router. This ensures that the static route is active (installed in the forwarding table) when the local PE router is the designated forwarder for the backup PE group. You can only associate a backup PE router group with a static route that also has the **p2mp-lsp-next-hop** statement configured. For more information, see “Configuring Static Unicast Routes for Point-to-Multipoint LSPs” on page 138.

Enabling Point-to-Point LSPs to Monitor Egress PE Routers

Configuring an LSP with the `associate-backup-pe-groups` statement enables it to monitor the status of the PE router to which it is configured. You can configure multiple backup PE router groups using the same router's address. A failure of this LSP indicates to all of the backup PE router groups that the destination PE router is down. The `associate-backup-pe-groups` statement is not tied to a specific backup PE router group. It applies to all groups that are interested in the status of the LSP to that address.

To allow an LSP to monitor the status of the egress PE router, include the `associate-backup-pe-groups` statement:

```
associate-backup-pe-groups;
```

This statement can be configured at the following hierarchy levels:

- [edit protocols mpls label-switched-path *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*]

If you configure the `associate-backup-pe-groups` statement, you must configure BFD for the point-to-point LSP. For information on how to configure BFD for an LSP, see “Configuring BFD for MPLS IPv4 LSPs” on page 167 and “Configuring BFD for LDP LSPs” on page 351.

You also must configure a full mesh of point-to-point LSPs between the PE routers in the backup PE router group. A full mesh is required so that each PE router within the group can independently determine the status of the other PE routers, allowing each router to independently determine which PE router is currently the designated forwarder for the backup PE router group.

If you configure multiple LSPs with the `associate-backup-pe-groups` statement to the same destination PE router, the first LSP configured is used to monitor the forwarding state to that PE router. If you configure multiple LSPs to the same destination, make sure to configure similar parameters for the LSPs. With this configuration scenario, a failure notification might be triggered even though the remote PE router is still up.

Preserving Point-to-Multipoint LSP Functioning with Different JUNOS Software Releases

In JUNOS Release 9.1 and earlier, Resv messages that include the `S2L_SUB_LSP` object are rejected by default. In JUNOS Release 9.2 and later, such messages are accepted by default. To ensure proper functioning of point-to-multipoint LSPs in a network that includes both devices running JUNOS Release 9.1 and earlier and devices running JUNOS 9.2 and later, you must include the `no-p2mp-sublsp` statement in the configuration of the devices running JUNOS 9.2 and later:

```
no-p2mp-sublsp;
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp]
- [edit logical-systems *logical-system-name* protocols rsvp]

Chapter 9

Miscellaneous MPLS Properties Configuration Guidelines

This chapter discusses the following topics:

- Configuring MPLS to Pop the Label on the Ultimate-Hop Router on page 149
- Configuring Traffic Engineering for LSPs on page 150
- Enabling Interarea Traffic Engineering on page 152
- Enabling Inter-AS Traffic Engineering for LSPs on page 153
- Configuring MPLS to Gather Statistics on page 156
- Configuring System Log Messages and SNMP Traps for LSPs on page 157
- Configuring MPLS Firewall Filters and Policers on page 158
- Configuring MPLS Rewrite Rules on page 165
- Configuring BFD for MPLS IPv4 LSPs on page 167
- Pinging LSPs on page 170
- Tracing MPLS and LSP Packets and Operations on page 171

Configuring MPLS to Pop the Label on the Ultimate-Hop Router

You can control the label value advertised on the egress router of a label-switched path (LSP). The default advertised label is label 3 (Implicit Null Label). If label 3 is advertised, the penultimate-hop router removes the label and sends the packet to the egress router. By enabling ultimate-hop popping, label 0 (IPv4 Explicit Null Label) is advertised. Ultimate-hop popping ensures that any packets traversing an MPLS network include a label.

To configure MPLS to pop the label on the ultimate-hop router, include the **explicit-null** statement:

```
explicit-null;
```

You can configure this statement at the following hierarchy levels:

- [edit protocols mpls]
- [edit logical-systems *logical-system-name* protocols mpls]



NOTE: Juniper Networks routers queue packets based on the incoming label. Routers from other vendors might queue packets differently. Keep this in mind when working with networks containing routers from multiple vendors.

For more information about labels, see “Label Description” on page 25 and “Label Allocation” on page 26.

Configuring Traffic Engineering for LSPs

When you configure an LSP, a host route (a 32-bit mask) is installed in the ingress router toward the egress router; the address of the host route is the destination address of the LSP. Typically, you configure the BGP option (`traffic-engineering bgp`), allowing only BGP to use LSPs in its route calculations. The other `traffic-engineering` statement options, allow you to alter this behavior in the master instance. This functionality is not available for specific routing instances. Also, you can enable only one of the `traffic-engineering` statement options (`bgp`, `bgp-igp`, `bgp-igp-both-ribs`, or `mpls-forwarding`) at a time.



NOTE: Enabling or disabling any of the `traffic-engineering` statement options causes all the MPLS routes to be removed and then reinserted into the routing tables.

You can configure OSPF and traffic engineering to advertise the LSP metric in summary link-state advertisements (LSAs) as described in the section “Advertising the LSP Metric in Summary LSAs” on page 152.

The following sections describe how to configure traffic engineering for LSPs:

- Using RSVP and LDP Routes for Traffic Forwarding on page 150
- Using RSVP and LDP Routes for Forwarding in Virtual Private Networks on page 151
- Using RSVP and LDP Routes for Forwarding But Not Route Selection on page 151
- Advertising the LSP Metric in Summary LSAs on page 152

Using RSVP and LDP Routes for Traffic Forwarding

Configure the `bgp-igp` option of the `traffic-engineering` statement to cause BGP and the interior gateway protocols (IGPs) to use LSPs for forwarding traffic destined for egress routers. The `bgp-igp` option causes all `inet.3` routes to be moved to the `inet.0` routing table.

On the ingress router, include the `traffic-engineering bgp-igp` statement:

```
traffic-engineering bgp-igp;
```

You can include this statement at the following hierarchy levels:

- `[edit protocols mpls]`

- [edit logical-systems *logical-system-name* protocols mpls]



NOTE: The `bgp-igp` option of the `traffic-engineering` statement cannot be configured for virtual private networks (VPNs). VPN routing instances require that routes be in the `inet.3` routing table.

Using RSVP and LDP Routes for Forwarding in Virtual Private Networks

VPNs rely on the routes in the `inet.3` routing table to function properly. For VPNs, configure the `bgp-igp-both-ribs` option of the `traffic-engineering` statement to cause BGP and the IGPs to use LSPs for forwarding traffic destined for egress routers. The `bgp-igp-both-ribs` option installs the ingress routes in both the `inet.0` routing table (for IPv4 unicast routes) and the `inet.3` routing table (for MPLS path information).

On the ingress router, include the `traffic-engineering bgp-igp-both-ribs` statement:

```
traffic-engineering bgp-igp-both-ribs;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls]
- [edit logical-systems *logical-system-name* protocols mpls]

Using RSVP and LDP Routes for Forwarding But Not Route Selection

If you configure the `traffic-engineering bgp-igp` statement or the `traffic-engineering bgp-igp-both-ribs` statement, high-priority RSVP and LDP routes can supersede IGP routes in the `inet.0` routing table. IGP routes might no longer be redistributed since they are no longer the active routes.

When you configure the `mpls-forwarding` option at either the [edit logical-systems *logical-system-name* protocols mpls `traffic-engineering`] hierarchy level or the [edit protocols mpls `traffic-engineering`] hierarchy level, RSVP and LDP routes are used for forwarding but are excluded from route selection. These routes are added to both the `inet.0` and `inet.3` routing tables. RSVP and LDP routes in the `inet.0` routing table are given a low preference when the active route is selected. However, RSVP and LDP routes in the `inet.3` routing table are given a normal preference and are therefore used for selecting forwarding next hops.

When you activate the `mpls-forwarding` option, routes whose state is `ForwardingOnly` are preferred for forwarding even if their preference is lower than that of the currently active route. To examine the state of a route, execute a `show route detail` command.

To configure, include the `traffic-engineering mpls-forwarding` statement:

```
traffic-engineering mpls-forwarding;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls]
- [edit logical-systems *logical-system-name* protocols mpls]

When you configure the `mpls-forwarding` option, IGP shortcut routes are copied to the `inet.0` routing table only.

Advertising the LSP Metric in Summary LSAs

You can configure MPLS and OSPF to treat an LSP as a link. This configuration allows other routers in the network to use this LSP. To accomplish this goal, you need to configure MPLS and OSPF traffic engineering to advertise the LSP metric in summary LSAs.

For MPLS, include the `traffic-engineering bgp-igp` and `label-switched-path` statements:

```
traffic-engineering bgp-igp;
label-switched-path lsp-name {
  to address;
}
```

You can include these statements at the following hierarchy levels:

- [edit protocols mpls]
- [edit logical-systems *logical-system-name* protocols mpls]

For OSPF, include the `lsp-metric-into-summary` statement:

```
lsp-metric-into-summary;
```

You can include this statement at the following hierarchy levels:

- [edit protocols ospf traffic-engineering shortcuts]
- [edit logical-systems *logical-system-name* protocols ospf traffic-engineering shortcuts]

For more information about MPLS traffic engineering, see “Configuring Traffic Engineering for LSPs” on page 150. For more information about OSPF traffic engineering, see the *JUNOS Routing Protocols Configuration Guide*.

Enabling Interarea Traffic Engineering

The JUNOS Software can signal a contiguous traffic-engineered LSP across multiple OSPF areas. The LSP signaling must be done using either nesting or contiguous signaling, as described in RFC 4206, *Label-Switched Paths (LSP) Hierarchy with Generalized Multi-Protocol Label Switching (GMPLS) Traffic Engineering (TE)*. However, contiguous signaling support is limited to just basic signaling. Reoptimization is not supported with contiguous signaling.

The following describes some of the interarea traffic engineering features:

- Interarea traffic engineering can be enabled when the loose-hop area border routers (ABRs) are configured on the ingress router using CSPF for the Explicit

Route Object (ERO) calculation within an OSPF area. ERO expansion is completed on the ABRs.

- Interarea traffic engineering can be enabled when CSPF is enabled, but without ABRs specified in the LSP configuration on the ingress router (ABRs can be automatically designated).
- Differentiated Services (DiffServ) traffic engineering is supported as long as the class type mappings are uniform across multiple areas.

To enable interarea traffic engineering, include the **expand-loose-hop** statement in the configuration for each LSP transit router:

```
expand-loose-hop;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls]
- [edit logical-systems *logical-system-name* protocols mpls]

Enabling Inter-AS Traffic Engineering for LSPs

Generally, traffic engineering is possible for LSPs that meet the following conditions:

- Both ends of the LSP are in the same OSPF area or at the same IS-IS level.
- The two ends of the LSP are in different OSPF areas within the same autonomous system (AS). LSPs that end in different IS-IS levels are not supported.
- The two ends of an explicit-path LSP are in different OSPF ASs and the autonomous system border routers (ASBRs) are configured statically as the loose hops supported on the explicit-path LSP. For more information about explicit-path LSPs, see “Static and Explicit-Path LSP Configuration Guidelines” on page 133.

Without statically defined ASBRs on LSPs, traffic engineering is not possible between one routing domain, or AS, and another. However, when the ASs are under the control of single service provider, it is possible in some cases to have traffic engineered LSPs span the ASs and dynamically discover the OSPF ASBRs linking them (IS-IS is not supported with this feature).

Inter-AS traffic engineered LSPs are possible as long as certain network requirements are met, none of the limiting conditions apply, and OSPF passive mode is configured with EGBP. Details are provided in the following sections:

- Inter-AS Traffic Engineering Requirements on page 153
- Inter-AS Traffic Engineering Limitations on page 154
- Configuring OSPF Passive TE Mode on page 155

Inter-AS Traffic Engineering Requirements

The proper establishment and functioning of inter-AS traffic engineered LSPs depend on the following network requirements, all of which must be met:

- All ASs are under control of a single service provider.
- OSPF is used as the routing protocol within each AS, and EBGp is used as the routing protocol between the ASs.
- ASBR information is available inside each AS.
- EBGp routing information is distributed by OSPF, and an IBGP full mesh is in place within each AS.
- Transit LSPs are *not* configured on the inter-AS links, but *are* configured between entry and exit point ASBRs on each AS.
- The EBGp link between ASBRs in different ASs is a direct link and must be configured as a passive traffic engineering link under OSPF. The remote link address itself, not the loopback or any other link address, is used as the remote node identifier for this passive link. For more information about OSPF passive traffic engineering mode configuration, see “Configuring OSPF Passive TE Mode” on page 155.

In addition, the address used for the remote node of the OSPF passive traffic engineering link must be the same as the address used for the EBGp link. For more information about OSPF and BGP in general, see the *JUNOS Routing Protocols Configuration Guide*.

Inter-AS Traffic Engineering Limitations

Only LSP hierarchical, or nested, signaling is supported for inter-AS traffic engineered LSPs. Only point-to-point LSPs are supported (there is no point-to-multipoint support).

In addition, the following limitations apply. Any one of these conditions is sufficient to render inter-AS traffic engineered LSPs impossible, even if the above requirements are met.

- The use of multihop BGP is not supported.
- The use of policers or topologies that prevent BGP routes from being known inside the AS is not supported.
- Multiple ASBRs on a LAN between EBGp peers are not supported. Only one ASBR on a LAN between EBGp peers is supported (others ASBRs can exist on the LAN, but cannot be advertised).
- Route reflectors or policies that hide ASBR information or prevent ASBR information from being distributed inside the ASs are not supported.
- Bidirectional LSPs are not supported (LSPs are unidirectional from the traffic engineering perspective).
- Topologies with both inter-AS and intra-AS paths to the same destination are not supported.

In addition, several features that are routine with all LSPs are not supported with inter-AS traffic engineering:

- Admin group link colors are not supported.
- Secondary standby is not supported.

- Reoptimization is not supported.
- Crankback on transit routers is not supported.
- Diverse path calculation is not supported.
- Graceful restart is not supported.

These lists of limitations or unsupported features with inter-AS traffic engineered LSPs are not exhaustive.

Configuring OSPF Passive TE Mode

Ordinarily, interior routing protocols such as OSPF are not run on links between ASs. However, for inter-AS traffic engineering to function properly, information about the inter-AS link, in particular, the address on the remote interface, must be made available inside the AS. This information is not normally included either in EBGp reachability messages or in OSPF routing advertisements.

To flood this link address information within the AS and make it available for traffic engineering calculations, you must configure OSPF passive mode for traffic engineering on each inter-AS interface. You must also supply the remote address for OSPF to distribute and include in the traffic engineering database.

To configure OSPF passive mode for traffic engineering on an inter-AS interface, include the **passive** statement for the link at the **[edit protocols ospf area *area-id* interface *interface-name*]** hierarchy level:

```
passive {
  traffic-engineering {
    remote-node-id ip-address; /* IP address at far end of inter-AS link */
  }
}
```

OSPF must be properly configured on the router. The following example configures the inter-AS link **so-1/1/0** to distribute traffic engineering information with OSPF within the AS. The local IP address on the link is **192.168.207.1** and the remote address is **192.168.207.2**.

```
[edit protocols ospf area 0.0.0.0]
interface so-1/1/0 {
  unit 0 {
    passive {
      traffic-engineering {
        remote-node-id 192.168.207.2;
      }
    }
    family inet {
      address 192.168.207.1;
    }
  }
}
```

Configuring MPLS to Gather Statistics

You can configure MPLS so that it periodically gathers traffic statistics about all MPLS sessions, including transit sessions, by configuring the **statistics** statement. You must configure the **statistics** statement if you want to collect MPLS traffic statistics using SNMP polling of MPLS Management Information Bases (MIBs).

To enable MPLS statistics collection, include the **statistics** statement:

```
statistics {
  auto-bandwidth;
  file filename <files number> <size size> <world-readable | no-world-readable>;
  interval seconds;
}
```

You can configure these statements at the following hierarchy levels:

- [edit protocols mpls]
- [edit logical-systems *logical-system-name* protocols mpls]

The default interval is 300 seconds.

The statistics are placed in a file, with one entry per LSP. During the specified interval, the following information is recorded in this file:

- The number of packets, number of bytes, packets per second, and bytes per second transmitted by each LSP.
- The percent of bandwidth transmitted over a given LSP in relation to the bandwidth percentage configured for that LSP. If no bandwidth is configured for an LSP, 0 percent is recorded in the percentage column.

At the end of each periodic report, a summary shows the current time, total number of sessions, number of sessions read, number of sessions ignored, and read errors, if any. Ignored sessions are typically those not in the up state or those with a reserved (0 through 15) incoming label (typically the egress point of an LSP). The reason for a read error appears on the same line as the entry for the LSP on which the error occurred. Gathering statistics is an unreliable process; occasional read errors might affect their accuracy. The following is a sample of the information included in the output file:

lsp6	0 pkt	0 Byte	0 pps	0 Bps	0
lsp5	0 pkt	0 Byte	0 pps	0 Bps	0
lsp6.1	34845 pkt	2926980 Byte	1049 pps	88179 Bps	132
lsp5.1	0 pkt	0 Byte	0 pps	0 Bps	0
lsp4	0 pkt	0 Byte	0 pps	0 Bps	0
Dec 7 17:28:38 Total 6 sessions: 5 success, 0 fail, 1 ignored					

Configuring System Log Messages and SNMP Traps for LSPs

Whenever an LSP makes a transition from up to down, or down to up, and whenever an LSP switches from one active path to another, the ingress router generates a system log message and sends an SNMP trap. The following shows a sample system log message:

```
MPLS lsp sheep1 up on primary(any) Route 192.168.1.1 192.168.1.2 192.168.1.3
MPLS lsp sheep1 change on primary(any) Route 192.168.1.1 192.168.1.2 192.168.1.3
MPLS lsp sheep1 down on primary(any)
MPLS lsp sheep1 up on secondary(any) Route 192.168.1.1 192.168.1.2 192.168.1.3
MPLS lsp sheep1 change on secondary(any) to primary(any), Route 192.168.1.1
192.168.1.2 192.168.1.3
```

For information about the MPLS SNMP traps and the proprietary MPLS MIBs, see the *JUNOS Network Management Configuration Guide*.

To generate system log messages for LSPs, include the **syslog** option to the **log-updown** statement:

```
log-updown {
    syslog;
}
```

To generate SNMP traps for LSPs, include the **trap** option to the **log-updown** statement:

```
log-updown {
    trap;
}
```

To generate SNMP traps whenever an LSP path goes down, include the **trap-path-down** option to the **log-updown** statement:

```
log-updown {
    trap-path-down;
}
```

To generate SNMP traps whenever an LSP path comes up, include the **trap-path-up** option to the **log-updown** statement:

```
log-updown {
    trap-path-up;
}
```

To disable the generation of system log messages, include the **no-syslog** option to the **log-updown** statement:

```
log-updown {
    no-syslog;
}
```

To disable the generation of SNMP traps, include the **no-trap** statement:

```
no-trap {
  mpls-lsp-traps;
  rfc3812-traps;
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls log-updown]
- [edit logical-systems *logical-system-name* protocols mpls log-updown]

For scalability reasons, only the ingress router generates SNMP traps. By default, MPLS issues traps for all configured LSPs. If you have many LSPs, the number of traps can become quite large. To disable the generation of SNMP traps, configure the **no-trap** statement.

The **no-trap** statement also includes the following options which allow you to block certain categories of MPLS SNMP traps:

- **mpls-lsp-traps**—Blocks the MPLS LSP traps defined in the `jnx-mpls.mib`, but allows the `rfc3812.mib` traps.
- **rfc-3812-traps**—Blocks the traps defined in the `rfc3812.mib`, but allows the MPLS LSP traps defined in the `jnx-mpls.mib`.

Configuring MPLS Firewall Filters and Policers

You can configure an MPLS firewall filter to count packets based on the EXP bits for the top-level MPLS label in a packet. You can also configure policers for MPLS LSPs.

The following sections discuss MPLS firewall filters and policers:

- Configuring MPLS Firewall Filters on page 158
- Examples: Configuring MPLS Firewall Filters on page 159
- Configuring Policers for LSPs on page 160
- Example: Configuring an LSP Policer on page 161
- Configuring Automatic Policers on page 162
- Writing Different DSCP and EXP Values in MPLS-Tagged IP Packets on page 165

Configuring MPLS Firewall Filters

You can configure an MPLS firewall filter to count packets based on the experimental (EXP) bits for the top-level MPLS label in a packet. You can then apply this filter to a specific interface. You can also configure a policer for the MPLS filter to police (that is, rate-limit) the traffic on the interface to which the filter is attached. You cannot apply MPLS firewall filters to Ethernet (fxp0) or loopback (lo0) interfaces.

You can configure an MPLS firewall filter on the M Series Multiservice Edge Routers and the T Series Core Routers.

You can configure the following match criteria attributes for MPLS filters at the [edit firewall family mpls filter *filter-name* term *term-name* from] hierarchy level:

- exp
- exp-except

These attributes can accept EXP bits in the range 0 through 7. You can configure the following choices:

- A single EXP bit—for example, **exp 3**;
- Several EXP bits—for example, **exp 0, 4**;
- A range of EXP bits—for example, **exp [0-5]**;

If you do not specify a match criterion (that is, you do not configure the **from** statement and use only the **then** statement with the **count** action keyword), all the MPLS packets passing through the interface on which the filter is applied will be counted.

You also can configure any of the following action keywords at the [edit firewall family mpls filter *filter-name* term *term-name* then] hierarchy level:

- count
- accept
- discard
- next
- policer

For more information about how to configure firewall filters, see the *JUNOS Policy Framework Configuration Guide*. For more information about how to configure interfaces, see the *JUNOS Network Interfaces Configuration Guide* and the *JUNOS Services Interfaces Configuration Guide*.

Examples: Configuring MPLS Firewall Filters

The following examples illustrate how you might configure an MPLS firewall filter and then apply the filter to an interface. This filter is configured to count MPLS packets with EXP bits set to either 0 or 4.

The following shows a configuration for an MPLS firewall filter:

```
[edit firewall]
family mpls {
  filter expf {
    term expt0 {
      from {
        exp 0,4;
      }
      then {
        count counter0;
        accept;
      }
    }
  }
}
```

```

    }
  }
}

```

The following shows how to apply the MPLS firewall filter to an interface:

```

[edit interfaces]
so-0/0/0 {
  mtu 4474;
  encapsulation ppp;
  sonet-options {
    fcs 32;
  }
  unit 0 {
    point-to-point;
    family mpls {
      filter {
        input expf;
        output expf;
      }
    }
  }
}

```

The MPLS firewall filter is applied to the input and output of an interface (see the input and output statements in the preceding example).

Configuring Policers for LSPs

MPLS LSP policing allows you to control the amount of traffic forwarded through a particular LSP. Policing helps to ensure that the amount of traffic forwarded through an LSP never exceeds the requested bandwidth allocation. LSP policing is supported on regular LSPs, LSPs configured with DiffServ-aware traffic engineering, and multiclass LSPs. You can configure multiple policers for each multiclass LSP. For regular LSPs, each LSP policer is applied to all of the traffic traversing the LSP. The policer's bandwidth limitations become effective as soon as the total sum of traffic traversing the LSP exceeds the configured limit.

You configure the multiclass LSP and DiffServ-aware traffic engineering LSP policers in a filter. The filter can be configured to distinguish between the different class types and apply the relevant policer to each class type. The policers distinguish between class types based on the EXP bits.

You configure LSP policers under the **family any** filter. The **family any** filter is used because the policer is applied to traffic entering the LSP. This traffic might be from different families: IPv6, MPLS, and so on. You do not need to know what sort of traffic is entering the LSP, as long as the match conditions apply to all types of traffic.

You can configure only those match conditions that apply across all types of traffic. The following are the supported match conditions for LSP policers:

- forwarding-class
- packet-length

- interface
- interface-set

To enable a policer on an LSP, first you need to configure a policing filter and then include it in the LSP configuration. For information on how to configure policers, see the *JUNOS Policy Framework Configuration Guide*.

To configure a policer for an LSP, specify a filter by including the **filter** option to the **policing** statement:

```
policing {
    filter filter-name;
}
```

You can include the **policing** statement at the following hierarchy levels:

- [edit protocols mpls label-switched-path *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*]

LSP Policer Limitations

When configuring MPLS LSP policers, be aware of the following limitations:

- LSP policers are supported for packet LSPs only.
- LSP policers are supported for unicast next hops only. Multicast next hops are not supported.
- LSP policers are not supported on aggregated interfaces.
- The LSP policer runs before any output filters.
- Traffic sourced from the Routing Engine (for example, ping traffic) does not take the same forwarding path as transit traffic. This type of traffic cannot be policed.
- LSP policers work on all T Series routers and on M Series routers that have the Internet Processor II application-specific integrated circuit (ASIC).

Example: Configuring an LSP Policer

The following example shows how you can configure a policing filter for an LSP:

```
[edit firewall]
policer police-ct1 {
    if-exceeding {
        bandwidth-limit 50m;
        burst-size-limit 1500;
    }
    then {
        discard;
    }
}
policer police-ct0 {
    if-exceeding {
```

```

        bandwidth-limit 200m;
        burst-size-limit 1500;
    }
    then {
        discard;
    }
}
family any {
    filter bar {
        term discard-ct0 {
            then {
                policer police-ct0;
                accept;
            }
        }
    }
    term discard-ct1 {
        then {
            policer police-ct1;
            accept;
        }
    }
}
}

```

Configuring Automatic Policers

Automatic policing of LSPs allows you to provide strict service guarantees for network traffic. Such guarantees are especially useful in the context of Differentiated Services for traffic engineered LSPs, providing better emulation for ATM wires over an MPLS network. For more information about Differentiated Services for LSPs, see “DiffServ-Aware Traffic Engineering Configuration Guidelines” on page 109.

Differentiated Services for traffic engineered LSPs allow you to provide differential treatment to MPLS traffic based on the EXP bits. To ensure these traffic guarantees, it is insufficient to simply mark the traffic appropriately. If traffic follows a congested path, the requirements might not be met.

LSPs are guaranteed to be established along paths where enough resources are available to meet the requirements. However, even if the LSPs are established along such paths and are marked properly, these requirements cannot be guaranteed unless you ensure that no more traffic is sent to an LSP than there is bandwidth available.

It is possible to police LSP traffic by manually configuring an appropriate filter and applying it to the LSP in the configuration. However, for large deployments it is cumbersome to configure thousands of different filters. Configuration groups cannot solve this problem either, since different LSPs might have different bandwidth requirements, requiring different filters. To police traffic for numerous LSPs, it is best to configure automatic policers.

When you configure automatic policers for LSPs, a policer is applied to all of the LSPs configured on the router. However, you can disable automatic policing on specific LSPs.



NOTE: You cannot configure automatic policing for LSPs carrying CCC traffic.

The following sections describe how to configure automatic policers for LSPs:

- Configuring Automatic Policers for LSPs on page 163
- Configuring Automatic Policers for DiffServ-Aware Traffic Engineering LSPs on page 164
- Configuring Automatic Policers for Point-to-Multipoint LSPs on page 164
- Disabling Automatic Policing on an LSP on page 164
- Example: Configuring Automatic Policing for an LSP on page 165

Configuring Automatic Policers for LSPs

To configure automatic policers for standard LSPs (neither DiffServ-aware traffic engineered LSPs nor multiclass LSPs), include the **auto-policing** statement with either the **class all** *policer-action* option or the **class ct0** *policer-action* option:

```
auto-policing {
  class all policer-action;
  class ct0 policer-action;
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls]
- [edit logical-systems *logical-system-name* protocols mpls]

You can configure the following policer actions for automatic policers:

- **drop**—Drop all packets.
- **loss-priority-high**—Set the packet loss priority (PLP) to high.
- **loss-priority-low**—Set the PLP to low.

These policer actions are applicable to all types of LSPs. The default policer action is to do nothing.

Automatic policers for LSPs police traffic based on the amount of bandwidth configured for the LSPs. You configure the bandwidth for an LSP using the **bandwidth** statement at the [edit protocols mpls label-switched-path *lsp-path-name*] hierarchy level. If you have enabled automatic policers on a router, change the bandwidth configured for an LSP, and commit the revised configuration, the change does not take effect on the active LSPs. To force the LSPs to use the new bandwidth allocation, issue a **clear mpls lsp** command.



NOTE: You cannot configure automatic policers for LSPs that traverse aggregated interfaces or Multilink Point-to-Point Protocol (MLPPP) interfaces.

Configuring Automatic Policers for DiffServ-Aware Traffic Engineering LSPs

To configure automatic policers for DiffServ-aware traffic engineering LSPs and for multiclass LSPs, include the `auto-policing` statement:

```
auto-policing {
  class all policer-action;
  class ctnumber policer-action;
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls]
- [edit logical-systems *logical-system-name* protocols mpls]

You include either the `class all policer-action` statement or a `class ctnumber policer-action` statement for each of one or more classes (you can configure a different policer action for each class). For a list of the actions that you can substitute for the *policer-action* variable, see “Configuring Automatic Policers for LSPs” on page 163. The default policer action is to do nothing.



NOTE: You cannot configure automatic policers for LSPs that traverse aggregated interfaces or MLPPP interfaces.

Configuring Automatic Policers for Point-to-Multipoint LSPs

You can configure automatic policers for point-to-multipoint (P2MP) LSPs by including the `auto-policing` statement with either the `class all policer-action` option or the `class ct0 policer-action` option. You only need to configure the `auto-policing` statement on the primary P2MP LSP (for more information on primary P2MP LSPs, see “Configuring the Primary Point-to-Multipoint LSP” on page 141). No additional configuration is required on the sub LSPs for the P2MP LSP.

The automatic policer configuration for P2MP LSPs is identical to the automatic policer configuration for standard LSPs. For more information, see “Configuring Automatic Policers for LSPs” on page 163.

Disabling Automatic Policing on an LSP

When you enable automatic policing, all of the LSPs on the router or logical system are affected. To disable automatic policing on a specific LSP on a router where you have enabled automatic policing, include the `policing` statement with the `no-automatic-policing` option:

```
policing no-automatic-policing;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls label-switched-path *lsp-name*]

- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*]

Example: Configuring Automatic Policing for an LSP

Configure automatic policing for a multiclass LSP, specifying different actions for class types ct0, ct1, ct2, and ct3.

```
[edit protocols mpls]
diffserv-te {
  bandwidth-model extended-mam;
}
auto-policing {
  class ct1 loss-priority-low;
  class ct0 loss-priority-high;
  class ct2 drop;
  class ct3 loss-priority-low;
}
traffic-engineering bgp-igp;
label-switched-path sample-lsp {
  to 3.3.3.3;
  bandwidth {
    ct0 11;
    ct1 1;
    ct2 1;
    ct3 1;
  }
}
interface fxp0.0 {
  disable;
}
interface t1-0/5/3.0;
interface t1-0/5/4.0;
```

Writing Different DSCP and EXP Values in MPLS-Tagged IP Packets

You can selectively set the DiffServ code point (DSCP) field of MPLS-tagged IPv4 and IPv6 packets to 0 without affecting output queue assignment, and continue to set the MPLS EXP field according to the configured rewrite table, which is based on forwarding classes. You can accomplish this by configuring a firewall filter for the MPLS-tagged packets.

For instructions on how to write different DSCP and EXP values in MPLS-tagged IP packets, see the *JUNOS Class of Service Configuration Guide*. For instructions on how to configure firewall filters, see the *JUNOS Policy Framework Configuration Guide*.

Configuring MPLS Rewrite Rules

You can apply a number of different rewrite rules to MPLS packets.

For more information about how to configure statements at the [edit class-of-service] hierarchy level, see the *JUNOS Class of Service Configuration Guide*.

The following sections describe how you can apply rewrite rules to MPLS packets:

- Rewriting the EXP Bits of All Three Labels of an Outgoing Packet on page 166
- Rewriting MPLS and IPv4 Packet Headers on page 166

Rewriting the EXP Bits of All Three Labels of an Outgoing Packet

In interprovider, carrier-of-carrier, and complex traffic engineering scenarios, it is sometimes necessary to push three labels on the next hop, using a swap-push-push or triple-push operation.

By default, on M Series routers except the M320, the top MPLS EXP label of an outgoing packet is not rewritten when you configure swap-push-push and triple-push operations. You can rewrite the EXP bits of all three labels of an outgoing packet, thereby maintaining the class of service (CoS) of an incoming MPLS or non-MPLS packet.

To push three labels on incoming MPLS packets, include the `exp-swap-push-push default` statement at the [edit class-of-service interfaces *interface-name* unit *logical-unit-number* rewrite-rules] hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules]
  exp-swap-push-push default;
```

To push three labels on incoming non-MPLS packets, include the `exp-push-push-push default` statement at the [edit class-of-service interfaces *interface-name* unit *logical-unit-number* rewrite-rules] hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules]
  exp-push-push-push default;
```

For more information about how to configure statements at the [edit class-of-service] hierarchy level, see the *JUNOS Class of Service Configuration Guide*.

Rewriting MPLS and IPv4 Packet Headers

You can apply a rewrite rule to MPLS and IPv4 packet headers simultaneously. This allows you to initialize MPLS EXP and IP precedence bits at LSP ingress. You can configure different rewrite rules depending on whether the traffic is VPN or non-VPN.

To rewrite MPLS and IPv4 packet headers, include the `protocol` statement at the [edit class-of-service interfaces *interface-name* unit *logical-unit-number* rewrite-rules *exp rewrite-rule-name*] hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules
  exp rewrite-rule-name]
  protocol types;
```

Use the `protocol` statement to specify the types of MPLS packets and packet headers to which to apply the rewrite rule. The MPLS packet can be a standard MPLS packet or an MPLS packet with an IPv4 payload. Specify the type of MPLS packet by using the following options:

- **mpls-any**—Applies the rewrite rule to MPLS packets and writes the code point value to MPLS headers.
- **mpls-inet-both**—Applies the rewrite rule to VPN MPLS packets with IPv4 payloads. Writes the code point value to the MPLS and IPv4 headers in T Series and M320 routers. On M Series routers, except the M320, the **mpls-inet-both** option causes all ingress MPLS LSP packets with IPv4 payloads to be initialized with 000 code points for IP precedence and MPLS EXP values.
- **mpls-inet-both-non-vpn**—Applies the rewrite rule to any non-VPN MPLS packets with IPv4 payloads. Writes the code point value to the MPLS and IPv4 headers in T Series and M320 routers. On M Series routers, except the M320, the **mpls-inet-both-non-vpn** option causes all ingress MPLS LSP packets with IPv4 payloads to be initialized with 000 code points for IP precedence and MPLS EXP values.

For a detailed example on how to configure rewrite rules for MPLS and IPv4 packets and for more information about how to configure class of service, see the *JUNOS Class of Service Configuration Guide*.

Configuring BFD for MPLS IPv4 LSPs

You can configure Bidirectional Forwarding Detection (BFD) protocol on MPLS IPv4 LSPs as outlined in the Internet draft draft-ietf-bfd-mpls-02.txt, *BFD for MPLS LSPs*. BFD is used as a periodic Operation, Administration, and Maintenance (OAM) feature for LSPs to detect LSP data plane faults. You can configure BFD for LSPs that use either LDP or RSVP as the signaling protocol.

You can also use the LSP **ping** commands to detect LSP data plane faults. However, BFD has a couple of benefits: it requires less computer processing than LSP **ping** commands and can quickly detect faults in large numbers of LSPs (LSP **ping** commands must be issued for each LSP individually). On the other hand, BFD cannot be used to verify the control plane against the data plane at the egress LSR, which is possible when an LSP **ping** echo request is associated with a forwarding equivalence class (FEC).

For configuration instructions for LDP-signaled LSPs, see “Configuring BFD for LDP LSPs” on page 351. For configuration instructions for RSVP-signaled LSPs, see the following section.

Configuring BFD for RSVP-Signaled LSPs

BFD for RSVP supports unicast IPv4 LSPs. When BFD is configured for an RSVP LSP on the ingress router, it is enabled on the primary path and on all standby secondary paths for that LSP. You can enable BFD for all LSPs on a router or for specific LSPs. If you configure BFD for a specific LSP, whatever values configured globally for BFD are overridden. The BFD sessions originate only at the ingress router and terminate at the egress router.

An error is logged whenever a BFD session for a path fails. The following shows how BFD for RSVP LSP log messages might appear:

RPD_MPLS_PATH_BFD_UP: MPLS BFD session for path path1 up on LSP R0_to_R3
 RPD_MPLS_PATH_BFD_DOWN: MPLS BFD session for path path1 down on LSP R0_to_R3

You can configure BFD for all of the RSVP LSPs on the router, a specific LSP, or the primary path of a specific LSP. To configure BFD for RSVP LSPs, include the `oam` and `bfd-liveness-detection` statements.

```
oam {
  bfd-liveness-detection {
    failure-action {
      make-before-break teardown-timeout seconds;
      teardown;
    }
    failure-action teardown;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    minimum-transmit-interval milliseconds;
    multiplier detection-time-multiplier;
  }
  lsp-ping-interval seconds;
}
```

You can configure this statement at the following hierarchy levels:

- [edit protocols mpls]
- [edit protocols mpls label-switched-path *lsp-name*]
- [edit protocols mpls label-switched-path *lsp-name* primary *path-name*]

The `bfd-liveness-detection` statement includes the following options:

- `minimum-interval`—Specifies the minimum transmit and receive interval.
- `minimum-receive-interval`—Specifies the minimum receive interval. The range is from 1 through 255,000 milliseconds.
- `minimum-transmit-interval`—Specifies the minimum transmit interval. The range is from 1 through 255,000 milliseconds.
- `multiplier`—Specifies the detection time multiplier. The range is from 1 through 255.



NOTE: To avoid triggering false negatives, configure a BFD fault detection time that is longer than the fast reroute time.

You can also configure the `lsp-ping-interval` option to adjust the time interval between LSP pings. The LSP ping command for RSVP-signaled LSPs is `ping mpls rsvp`. For more information on the `ping mpls rsvp` command, see the *JUNOS System Basics and Services Command Reference*.

Configuring a Failure Action for the BFD Session on an RSVP LSP

When the BFD session for an RSVP LSP goes down, the LSP is torn down and resigaled. Traffic can be switched to a standby LSP or to simply teardown the LSP path. Any actions performed are logged.

When a BFD session for an RSVP LSP path goes down, you can configure the JUNOS Software to resignal the LSP path or to simply disable the LSP path. A standby LSP path could be configured to handle traffic while the primary LSP path is unavailable. The router can automatically recover from LSP failures that can be detected by BFD. By default, if a BFD session fails, the event is simply logged.

To enable the JUNOS Software to tear down an RSVP LSP path in the event of a BFD event, include the **failure-action** statement:

```
failure-action {
    make-before-break teardown-timeout seconds;
    teardown;
}
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

You can configure either the **teardown** or **make-before-break** options:

- **teardown**—Causes the LSP path to be taken down and resigaled immediately.
- **make-before-break**—Causes the JUNOS Software to attempt to signal a new LSP path before tearing down the old LSP path. You can also configure the **teardown-timeout** option to automatically tear down the LSP after the time period specified if the attempt to resignal the LSP fails within the **teardown-timeout** interval. If you specify a value of 0 for the **teardown-timeout** interval, the LSP is taking down and resigaled immediately (the same behavior as when you configure the **teardown** option).

To configure a failure action for all of the RSVP LSPs include the **failure-action** statement at the [edit protocols mpls oam bfd-liveness-detection] hierarchy level. To configure a failure action for a specific RSVP LSP, include the **failure-action** statement at the [edit protocols mpls label-switched-path *lsp-name* oam bfd-liveness-detection] hierarchy level.

To configure a failure action for a specific primary path, include the **failure-action** statement at the [edit protocols mpls label-switched path *lsp-name* primary *path-name* oam bfd-liveness-detection] hierarchy level. To configure a failure action for a specific secondary LSP path, include the **failure-action** statement at the [edit protocols mpls label-switched-path *lsp-name* secondary *path-name* oam bfd-liveness-detection] hierarchy level.

Pinging LSPs

The following sections describe how to use the `ping mpls` command to confirm LSP functioning.

- Pinging MPLS LSPs on page 170
- Pinging Point-to-Multipoint LSPs on page 170
- Pinging the Endpoint Address of MPLS LSPs on page 171
- Pinging CCC LSPs on page 171
- Pinging Layer 3 VPNs on page 171
- Support for LSP Ping and Traceroute Commands Based on RFC 4379 on page 171

Pinging MPLS LSPs

You can ping a specific LSP. Echo requests are sent over the LSP as MPLS packets. The payload is a User Datagram Protocol (UDP) packet forwarded to an address in the 127/8 range (127.0.0.1 by default, this address is configurable) and port 3503. The label and interface information for building and sending this information as an MPLS packet is the same as for standard LSP traffic.

When the echo request arrives at the egress node, the receiver checks the contents of the packet and sends a reply containing the correct return value, by using UDP. The router sending the echo request waits to receive an echo reply after a timeout of 2 seconds (you cannot configure this value).

You must configure MPLS at the `[edit protocols mpls]` hierarchy level on the remote router to be able to ping an LSP terminating there. You must configure MPLS even if you intend to ping only LDP forwarding equivalence classes (FECs).

To ping an MPLS LSP use the `ping mpls <count count> <ldp <fec>> <rsvp <exp forwarding-class> <lsp-name>>` command. To ping a secondary MPLS LSP, use the `ping mpls <count count> <rsvp <lsp-name>> standby path-name` command. For a detailed description of this command, see the *JUNOS Routing Protocols and Policies Command Reference*.



NOTE: The `ping mpls` command is not supported within routing instances.

Pinging Point-to-Multipoint LSPs

To ping a point-to-multipoint LSP, use the `ping mpls rsvp lsp-name multipoint` or `ping mpls rsvp egress address` commands. The `ping mpls rsvp lsp-name multipoint` command returns a list of all of the egress router identifiers and the current status of the point-to-multipoint LSP egress routers. The `ping mpls rsvp lsp-name multipoint egress address` command returns the current status of the specified egress router.

Pinging the Endpoint Address of MPLS LSPs

To determine whether an LSP between two provider edge (PE) routers is up and running, you can ping the endpoint address of the LSP. To ping an MPLS LSP endpoint, use the `ping mpls lsp-end-point address` command. This command tells you what type of LSP (RSVP or LDP) terminates at the address specified and whether that LSP is up or down.

For a detailed description of this command, see the *JUNOS Routing Protocols and Policies Command Reference*.

Pinging CCC LSPs

You can ping a specific CCC LSP. The CCC LSP ping command is identical to the one used for MPLS LSPs. The command you use is `ping mpls <count count> <rsvp <lsp-name>>`. You can also ping a secondary standby CCC LSP by using the `ping mpls <count count> <rsvp <lsp-name>> standby path-name` command.

For a detailed description of this command, see the *JUNOS Routing Protocols and Policies Command Reference*.

Pinging Layer 3 VPNs

You can use a similar command, `ping mpls l3vpn vpn-name prefix prefix <count count>`, to ping a Layer 3 VPN. For more information about this command, see the *JUNOS VPNs Configuration Guide* and the *JUNOS Routing Protocols and Policies Command Reference*.

Support for LSP Ping and Traceroute Commands Based on RFC 4379

The JUNOS Software partially supports LSP `ping` and `traceroute` commands based on RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*. However, the JUNOS Software only supports this functionality on LSP transit routers. If a `ping` or `traceroute` command is issued from a router that fully supports RFC 4379, it can propagate correctly on routers running the JUNOS Software.

LSP `ping` and `traceroute` commands based on RFC 4379 attempt to trace the path taken by an LSP by relying on MPLS TTL expiration. An LSP can take multiple paths from ingress to egress. This occurs in particular with Equal Cost Multipath (ECMP). The LSP `traceroute` command can trace all possible paths to an LSP egress node.

Tracing MPLS and LSP Packets and Operations

To trace MPLS and LSP packets and operations, include the `traceoptions` statement:

```
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  flag flag;
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

You can specify the following MPLS-specific flags in the MPLS **traceoptions** statement:

- **all**—Trace all operations.
- **connection**—Trace all circuit cross-connect (CCC) activity.
- **connection-detail**—Trace detailed CCC activity.
- **cspf**—Trace CSPF computations.
- **cspf-link**—Trace links visited during CSPF computations.
- **cspf-node**—Trace nodes visited during CSPF computations.
- **error**—Trace MPLS error conditions.
- **graceful-restart**—Trace MPLS graceful restart events.
- **lsping**—Trace LSP ping packets and return codes.
- **state**—Trace all LSP state transitions.

When you configure trace options to track an MPLS LSP using the **cspf** option, the CSPF log displays information about the MPLS LSP using the term “generalized MPLS” (GMPLS). For example, a message in the CSPF log might state that the “link passes GMPLS constraints”. Generalized MPLS (GMPLS) is a superset of MPLS, so this message is normal and does not affect proper MPLS LSP operation.

For general information about tracing and global tracing options, see the *JUNOS Routing Protocols Configuration Guide*.

Chapter 10

Summary of MPLS Configuration Statements

This chapter shows the complete MPLS configuration statements. The statements are organized alphabetically.

adaptive

Syntax	adaptive;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	During reroute, do not double-count bandwidth on links shared by the old and new paths. Including this statement causes RSVP to use shared explicit (SE) reservation styles and assists in smooth transition during rerouting.
Default	The configured object is disabled.
Usage Guidelines	See “Configuring Adaptive LSPs” on page 101.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

adjust-interval

Syntax	adjust-interval <i>seconds</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth], [edit protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the bandwidth reallocation interval.
Options	<i>seconds</i> —Bandwidth reallocation interval, in seconds. Range: 300 through 4,294,967,295 seconds Default: 86,400 seconds
Usage Guidelines	See “Configuring the Automatic Bandwidth Allocation Interval” on page 90.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

adjust-threshold

Syntax	adjust-threshold <i>percent</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth], [edit protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify how sensitive the automatic bandwidth adjustment for a label-switched path (LSP) is to changes in bandwidth utilization.
Options	<i>percent</i> —Bandwidth demand for the current bandwidth adjustment interval is determined and compared to the LSP’s current bandwidth allocation. If the percentage difference in bandwidth is greater than or equal to the percentage specified by this statement, the LSP’s bandwidth is adjusted to the current bandwidth demand.
Usage Guidelines	See “Configuring the Automatic Bandwidth Adjustment Threshold” on page 91.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

adjust-threshold-overflow-limit

Syntax	adjust-threshold-overflow-limit <i>number</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth], [edit protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth]
Release Information	Statement introduced in JUNOS Release 7.5.
Description	Specify the number of consecutive bandwidth overflow samples before triggering a bandwidth adjustment.
Options	<i>number</i> —Number of consecutive bandwidth overflow samples. Range: 1 through 65,535 Default: This feature is disabled by default.
Usage Guidelines	See “Configuring a Limit on Bandwidth Overflow Samples” on page 91.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

admin-down

Syntax	admin-down;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced in JUNOS Release 8.2.
Description	Set the A-bit in the Admin Status object. When set, this bit indicates the administrative down status for an LSP. This feature is used specifically by non-packet GMPLS LSPs. It does not affect control path setup or data forwarding for packet LSPs.
Usage Guidelines	See “Allowing Non-Packet GMPLS LSPs to Establish Paths Through Routers Running the JUNOS Software” on page 462.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

admin-group

See the following sections:

- admin-group (for Interfaces) on page 176
- admin-group (for LSPs) on page 177

admin-group (for Interfaces)

Syntax	admin-group [<i>group-names</i>];
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i>], [edit protocols mpls interface <i>interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Define administrative groups for an interface.
Options	<i>group-names</i> —One or more names of groups defined with the admin-groups statement at the [edit protocols mpls] hierarchy level.
Usage Guidelines	See “Configuring Administrative Groups” on page 96.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Topics	admin-groups

admin-group (for LSPs)

Syntax	admin-group { exclude [<i>group-names</i>]; include-all [<i>group-names</i>]; include-any [<i>group-names</i>]; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>], [edit protocols mpls], [edit protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Define the administrative groups to include or exclude for an LSP and for a path's primary and secondary paths.
Options	The statements are explained separately.
Usage Guidelines	See "Configuring Administrative Groups" on page 96.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

admin-groups

Syntax	admin-groups { <i>group-name group-value</i> ; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure administrative groups to implement link coloring of resource classes.
Options	<p><i>group-name</i>—Name of the group. You can assign up to 32 names. The names and their corresponding values must be identical across all routers within a single domain.</p> <p><i>group-value</i>—Value assigned to the group. The names and their corresponding values must be identical across all routers within a single domain.</p> <p>Range: 0 through 31</p>
Usage Guidelines	See “Configuring Administrative Groups” on page 96.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Topics	admin-group

advertisement-hold-time

Syntax	advertisement-hold-time <i>seconds</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Do not advertise when the LSP goes from up to down, for a certain period of time known as hold time.
Options	<p><i>seconds</i>—Hold time, in seconds.</p> <p>Range: 0 through 65,535 seconds</p> <p>Default: 5 seconds</p>
Usage Guidelines	See “Damping Advertisement of LSP State Changes” on page 107.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

allow-fragmentation

Syntax	allow-fragmentation;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls path-mtu], [edit protocols mpls path-mtu]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Allow IP packets to be fragmented before they are encapsulated in MPLS.
Usage Guidelines	See “Enabling Packet Fragmentation” on page 294.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

associate-backup-pe-groups

Syntax	associate-backup-pe-groups;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced in JUNOS Release 9.0.
Description	Enable an LSP to monitor the status of its destination PE router. You can configure multiple backup PE router groups using the same router's address. Backup PE router groups provide ingress PE router redundancy when point-to-multipoint (P2MP) LSPs are configured for multicast distribution. A failure of this LSP indicates to all of the backup PE router groups that the destination PE router is down. This statement is not tied to a specific backup PE router group. It applies to all groups that are interested in the status of the LSP to the destination address.
Usage Guidelines	See “Enabling Point-to-Point LSPs to Monitor Egress PE Routers” on page 147.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

auto-bandwidth

Syntax auto-bandwidth {
 adjust-interval *seconds*;
 adjust-threshold *percent*;
 adjust-threshold-overflow-limit *number*;
 maximum-bandwidth *bps*;
 minimum-bandwidth *bps*;
 monitor-bandwidth;
 }

Hierarchy Level [edit protocols mpls label-switched-path *lsp-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Allow an MPLS tunnel to automatically adjust its bandwidth allocation based on the volume of traffic flowing through the tunnel.

Options The statements are explained separately.

Usage Guidelines See “Configuring Automatic Bandwidth Allocation for LSPs” on page 88.

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

auto-policing

Syntax	<pre> auto-policing { class all (drop loss-priority-high loss-priority-low); class ctnumber (drop loss-priority-high loss-priority-low); } </pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Enable the automatic policing of all the MPLS LSPs on the router or logical system.
Options	<p>class all—Apply the same policer action to all the class types (ct0, ct1, ct2, and ct3).</p> <p>class <i>ctnumber</i>—Specific class type (ct0, ct1, ct2, or ct3) to which to apply a policer action.</p> <p>Policer actions—You can specify the following policer actions:</p> <p>Default: no action</p> <ul style="list-style-type: none"> ■ drop—Drop all packets. ■ loss-priority-high—Set the packet loss priority (PLP) to high. ■ loss-priority-low—Set the PLP to low.
Usage Guidelines	See “Configuring Automatic Policers” on page 162.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Topics	policing

backup-pe-group

Syntax	<pre> backup-pe-group <i>pe-group-name</i> { backups [<i>addresses</i>]; local-address <i>address</i>; }</pre>
Hierarchy Level	<pre> [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast], [edit logical-systems <i>logical-system-name</i> routing-options multicast], [edit routing-instances <i>routing-instance-name</i> routing-options multicast], [edit routing-options multicast]</pre>
Release Information	Statement introduced in JUNOS Release 9.0.
Description	Configure a backup provider edge (PE) group for ingress PE router redundancy when point-to-multipoint (P2MP) label-switched paths (LSPs) are used for multicast distribution.
Options	<p>backups <i>addresses</i>—Specify the address of backup PE routers for ingress PE redundancy when point-to-multipoint LSPs are used for multicast distribution.</p> <p>local-address <i>address</i>—Specify the address of the local PE router for ingress PE redundancy when point-to-multipoint LSPs are used for multicast distribution.</p> <p><i>pe-group-name</i>—Specify the name for the group of PE routers that provide ingress PE router redundancy for P2MP LSPs.</p>
Usage Guidelines	See “Configuring Ingress PE Router Redundancy for Point-to-Multipoint LSPs” on page 146.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

bandwidth

Syntax	<pre>bandwidth bps { ct0 bps; ct1 bps; ct2 bps; ct3 bps; }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> fast-reroute], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>], [edit protocols mpls], [edit protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i> fast-reroute], [edit protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>]</pre>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>When configuring an LSP, specify the traffic rate associated with the LSP.</p> <p>When configuring fast reroute, allocate bandwidth for the reroute path. By default, no bandwidth is reserved for the rerouted path. The fast reroute bandwidth does not need to be identical to that allocated for the LSP itself.</p> <p>When configuring a multiclass LSP, use the <i>ctnumber bandwidth</i> statements to specify the bandwidth to be allocated for each class type.</p>
Options	<p><i>bps</i>—Bandwidth, in bits per second. You can specify this as an integer value. You can also use the abbreviations k (for a thousand), m (for a million), or g (for a billion).</p> <p>Range: Any positive integer Default: 0 (no bandwidth is reserved)</p> <p><i>ctnumber bps</i>—Bandwidth for the specified class type, in bits per second. You can specify this as an integer value. If you do so, count your zeros carefully, or you can use the abbreviations k (for a thousand), m (for a million), or g (for a billion [also called a thousand million]).</p> <p>Range: Any positive integer Default: 0 (no bandwidth is reserved)</p>
Usage Guidelines	See “Configuring Fast Reroute” on page 78, “Configuring the Bandwidth Value for LSPs” on page 106, “Configuring Traffic-Engineered LSPs” on page 127, and “Configuring Class-Type Bandwidth Constraints for Multiclass LSPs” on page 129.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

bandwidth-model

Syntax	bandwidth-model { extended-mam; mam; rdm; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls diffserv-te], [edit protocols mpls diffserv-te]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure the bandwidth model for differentiated services. Note that you cannot configure both bandwidth models at the same time.
Options	<p>extended-mam—The extended maximum allocation model (MAM) is a bandwidth model based on MAM.</p> <p>mam—The MAM is defined in RFC 4125, <i>Maximum Allocation Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering</i>.</p> <p>rdm—The Russian dolls bandwidth allocation model (RDM) is defined in RFC 4127, <i>Russian Dolls Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering</i>. RDM makes efficient use of bandwidth by allowing the class types to share bandwidth.</p>
Usage Guidelines	See “Configuring the Bandwidth Model” on page 116.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

bandwidth-percent

Syntax	<code>bandwidth-percent percentage;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> fast-reroute], [edit protocols mpls label-switched-path <i>lsp-name</i> fast-reroute]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure the percentage of bandwidth to reserve for the detour path in case the primary path for a traffic engineered LSP or a multiclass LSP fails. The percentage configured indicates the percentage of the protected path's bandwidth that is reserved for the detour path.
Options	<i>percentage</i> —The percentage of the protected path's bandwidth that is reserved for the detour path.
Usage Guidelines	See “Configuring Fast Reroute” on page 78, “Configuring Fast Reroute for Traffic-Engineered LSPs” on page 128 and “Configuring Fast Reroute for Multiclass LSPs” on page 131.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

bfd-liveness-detection

Syntax	<pre> bfd-liveness-detection { failure-action { make-before-break teardown-timeout <i>seconds</i>; teardown; } minimum-interval <i>milliseconds</i>; minimum-receive-interval <i>milliseconds</i>; minimum-transmit-interval <i>milliseconds</i>; multiplier <i>detection-time-multiplier</i>; }</pre>
Hierarchy Level	[edit protocols mpls label-switched-path <i>lsp-name</i> oam], [edit protocols mpls oam]
Release Information	Statement introduced in JUNOS Release 7.6. The failure-action statement option was added in JUNOS Release 8.5.
Description	Enable Bidirectional Forwarding Detection (BFD) for all of the MPLS LSPs or for just a specific LSP.
Options	<p>minimum-interval—Minimum transmit and receive interval. Range: 50 through 255,000 milliseconds Default: 50</p> <p>minimum-receive-interval—Minimum receive interval. Range: 50 through 255,000 milliseconds Default: 50</p> <p>minimum-transmit-interval—Minimum transmit interval. Range: 50 through 255,000 milliseconds Default: 50</p> <p>multiplier—Detection time multiplier. Range: 1 through 255 Default: 3</p> <p>The failure-action statement is explained separately.</p>
Usage Guidelines	See “Configuring BFD for MPLS IPv4 LSPs” on page 167.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

class-of-service

Syntax	<code>class-of-service cos-value;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i> label-map (default-route <i>in-label</i>)], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls static-path inet <i>address</i>], [edit protocols mpls], [edit protocols mpls interface <i>interface-name</i> label-map (default-route <i>in-label</i>)], [edit protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>], [edit protocols mpls static-path inet <i>address</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>Class-of-service (CoS) value given to all packets in the LSP.</p> <p>The CoS value might affect the scheduling or queuing algorithm of traffic traveling along an LSP.</p>
Options	<i>cos-value</i> —CoS value. A higher value typically corresponds to a higher level of service. Range: 0 through 7 Default: If you do not specify a CoS value, the IP precedence bits from the packet's IP header are used as the packet's CoS value.
Usage Guidelines	See “Configuring Class of Service for MPLS LSPs” on page 98, “Configuring the Ingress Router for Static LSPs” on page 133, and “Configuring the Intermediate and Egress Routers for Static LSPs” on page 136.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

default-route

Syntax	default-route { class-of-service <i>cos-value</i> ; (next-hop (<i>address</i> <i>interface-name</i> <i>address/interface-name</i>) (discard reject); (pop swap <i>out-label</i>); preference <i>preference</i> ; swap-push <i>swap-label push-label</i> ; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls interface (<i>interface-name</i> all) label-map], [edit protocols mpls interface (<i>interface-name</i> all) label-map]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Process MPLS packets that have not been assigned label values and have no corresponding entry in the mpls.0 table. The remaining statements are explained separately.
Usage Guidelines	See “Configuring a Default Route for Unlabeled MPLS Packets” on page 60.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

description

Syntax	description <i>text</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Provides a textual description of the LSP. Enclose any descriptive text that includes spaces in quotation marks (" "). Any descriptive text you include is displayed in the output of the <code>show mpls lsp detail</code> command and has no effect on the operation of the LSP.
Options	<i>text</i> —Provide a textual description of the LSP.
Usage Guidelines	See “Configuring a Text Description for LSPs” on page 78.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

diffserv-te

Syntax

```
diffserv-te {
  bandwidth-model {
    extended-mam;
    mam;
    rdm;
  }
  te-class-matrix {
    tnumber {
      priority priority;
      traffic-class {
        ctnumber priority priority;
      }
    }
  }
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls],
[edit protocols mpls]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify properties for differentiated services in traffic engineering.

Options The statements are explained separately.

Usage Guidelines See “Configuring Routers for DiffServ-Aware Traffic Engineering” on page 115.

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

disable

Syntax	disable;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth], [edit protocols mpls], [edit protocols mpls interface <i>interface-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Disable the functionality of the configured object.
Default	The configured object is enabled (operational) unless explicitly disabled.
Usage Guidelines	See “LSP Configuration Overview” on page 70.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

discard

Syntax	discard;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i> label-map (default-route <i>in-label</i>)], [edit protocols mpls interface <i>interface-name</i> label-map (default-route <i>in-label</i>)],
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Do not forward packets that match the incoming label. Instead, drop the packets and do not send an Internet Control Message Protocol (ICMP) unreachable message.
Usage Guidelines	See “Configuring the Intermediate and Egress Routers for Static LSPs” on page 136.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

double-push

Syntax	double-push <i>bottom-label top-label</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Push two more labels on top of the stack.
Options	<i>bottom-label</i> —(Optional) Label value. <i>top-label</i> —(Optional) Label value. Range: Labels 0 through 999,999 are for internal use. Labels 1,000,000 through 1,048,575 are unassigned by the JUNOS Software and are available for static LSPs. When you configure static LSPs, you can use only this range of labels.
Usage Guidelines	See “Configuring the Ingress Router for Static LSPs” on page 133.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

encoding-type

Syntax	encoding-type (ethernet packet pdh sonet-sdh);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> lsp-attributes], [edit protocols mpls label-switched-path <i>lsp-name</i> lsp-attributes]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the encoding type of payload carried by the LSP. It can be any of the following: <ul style="list-style-type: none"> ■ ethernet—Ethernet ■ packet—Packet ■ pdh—Plesiochronous digital hierarchy (PDH) ■ sonet-sdh—SONET/SDH
Default	packet
Usage Guidelines	See “Configuring the Encoding Type” on page 460.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

exclude

See the following sections:

- [exclude \(for Administrative Groups\)](#) on page 192
- [exclude \(for Fast Reroute\)](#) on page 193

exclude (for Administrative Groups)

Syntax `exclude [group-names];`

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* admin-group],
 [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* (primary | secondary) *path-name* admin-group],
 [edit protocols mpls label-switched-path *lsp-name* admin-group],
 [edit protocols mpls label-switched-path *lsp-name* (primary | secondary) *path-name* admin-group]

Release Information Statement introduced before JUNOS Release 7.4.

Description Define the administrative groups to exclude for an LSP or for a path's primary and secondary paths.

Options *group-names*—Names of one or more groups defined with the `admin-groups` statement.

Usage Guidelines See “Configuring Administrative Groups” on page 96.

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

exclude (for Fast Reroute)

Syntax	(exclude [<i>group-names</i>] no-exclude);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> fast-reroute], [edit protocols mpls label-switched-path <i>lsp-name</i> fast-reroute]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Control exclusion of administrative groups: <ul style="list-style-type: none"> ■ exclude—Define the administrative groups to exclude for fast reroute. ■ no-exclude—Disable administrative group exclusion.
Options	<i>group-names</i> —Names of one or more groups defined with the admin-groups statement.
Usage Guidelines	See “Configuring Fast Reroute” on page 78.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Topics	admin-groups

expand-loose-hop

Syntax	expand-loose-hop;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
Release Information	Statement introduced in JUNOS Release 7.6.
Description	Allow an LSP to traverse multiple OSPF areas within a service provider’s network.
Usage Guidelines	See “Enabling Interarea Traffic Engineering” on page 152.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

explicit-null

Syntax	explicit-null;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Advertise label 0 to the egress router of an LSP.
Default	If you do not include the explicit-null statement in the MPLS configuration, label 3 (implicit null) is advertised.
Usage Guidelines	See “Configuring RSVP to Pop the Label on the Ultimate-Hop Router” on page 295.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

failure-action

Syntax	failure-action { make-before-break teardown-timeout <i>seconds</i> ; teardown; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls oam bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> oam bfd-liveness-detection], [edit protocols mpls label-switched-path <i>lsp-name</i> oam bfd-liveness-detection], [edit protocols mpls oam bfd-liveness-detection]
Release Information	Statement introduced in JUNOS Release 9.4.
Description	Configure route and next-hop properties in the event of a Bidirectional Forwarding Detection (BFD) protocol session failure event on an RSVP label-switched path (LSP). The failure event could be an existing BFD session that has gone down or a BFD session that never came up. RSVP adds back the route or next hop when the relevant BFD session comes back up.
Options	<p>make-before-break—When a BFD session fails for an RSVP LSP, an attempt is made to signal a new LSP path before tearing down the old LSP path.</p> <p>teardown—When a BFD session fails for an RSVP LSP, the associated LSP path is taken down and resigaled immediately.</p> <p>teardown-timeout <i>seconds</i>—When you configure the make-before-break option, you can specify a time in seconds for the teardown-timeout option. At the end of the time specified, the associated RSVP LSP is automatically torn down and resigaled.</p>
Usage Guidelines	See “Configuring a Failure Action for the BFD Session on an RSVP LSP” on page 169.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

fast-reroute

Syntax	<pre>fast-reroute { (bandwidth <i>bps</i> bandwidth-percent <i>percentage</i>); (exclude [<i>group-names</i>] no-exclude); hop-limit <i>number</i>; (include-all [<i>group-names</i>] no-include-all); (include-any [<i>group-names</i>] no-include-any); }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Establish detours for the LSP so that if a node or link in the LSP fails, the traffic on the LSP can be rerouted with minimal packet loss.
Options	The statements are explained separately.
Usage Guidelines	See “Configuring Fast Reroute” on page 78.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

fate-sharing

Syntax	<pre>fate-sharing { group <i>group-name</i> { cost <i>value</i>; from <i>address</i> <to <i>address</i>>; } }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options], [edit routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify groups of objects that share characteristics resulting in backup paths to be used if primary paths fail. All objects are treated as /32 host addresses. You specify one or more objects within a group. The objects can be LAN interfaces, router IDs, or point-to-point links. The sequence is insignificant.
Options	<p>cost <i>value</i>—Cost assigned to the group. Range: 1 through 65,535 Default: 1</p> <p>from <i>address</i>—Address of the router or address of the LAN/NBMA interface. For example, an Ethernet network with four hosts in the same fate-sharing group would require you to list all four of the separate from addresses in the group.</p> <p>group <i>group-name</i>—Each fate-sharing group must have a name, which can have a maximum of 32 characters, including letters, numbers, periods (.), and hyphens (-). You can define up to 512 groups.</p> <p>to <i>address</i>—(Optional) Address of egress router. For point-to-point link objects, you must specify both a from and a to address.</p>
Usage Guidelines	See “Configuring Alternate Backup Paths Using Fate Sharing” on page 54.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

from

Syntax	<code>from address;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the source address to use for the LSP. The address you specify does not affect the outgoing interface used by the LSP.
Default	If you do not include this statement, the software automatically selects the loopback interface as the address.
Options	<i>address</i> —IP address.
Usage Guidelines	See “Configuring the Ingress Router Address for LSPs” on page 73.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

gpid

Syntax	<code>gpid (ethernet hdlc ipv4 pos-scrambling-crc-16 pos-no-scrambling-crc-16 pos-scrambling-crc-32 pos-no-scrambling-crc-32 ppp);</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> lsp-attributes], [edit protocols mpls label-switched-path <i>lsp-name</i> lsp-attributes]
Release Information	Statement introduced before JUNOS Release 7.4. pos-scrambling-crc-16, pos-no-scrambling-crc-16, pos-scrambling-crc-32, and pos-no-scrambling-crc-32 options added in JUNOS Release 8.0.
Description	Specify the type of payload carried by the LSP. It can be any of the following: <ul style="list-style-type: none"> ■ <code>ethernet</code>—Ethernet (GPID value: 33) ■ <code>hdlc</code>—High-level Data Link Control (HDLC) (GPID value: 44) ■ <code>ipv4</code>—IP version 4 (GPID value: 0x0800) ■ <code>pos-no-scrambling-crc-16</code>—for interoperability with other vendors' equipment (GPID value: 29) ■ <code>pos-no-scrambling-crc-32</code>—for interoperability with other vendors' equipment (GPID value: 30) ■ <code>pos-scrambling-crc-16</code>—for interoperability with other vendors' equipment (GPID value: 31) ■ <code>pos-scrambling-crc-32</code>—for interoperability with other vendors' equipment (GPID value: 32) ■ <code>ppp</code>—Point-to-Point Protocol (PPP) (GPID value: 50)
Default	<code>ipv4</code>
Usage Guidelines	See “Configuring the GPID” on page 461.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

hop-limit

Syntax	hop-limit <i>number</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> fast-reroute], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>], [edit protocols mpls], [edit protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i> fast-reroute], [edit protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For an LSP, specify the maximum number of routers that the LSP can traverse, including the ingress and egress routers. For fast reroute, how many more routers a detour is allowed to traverse compared with the LSP itself. For example, if an LSP traverses 4 routers, any detour for the LSP can be no more than 10 router hops, including the ingress and egress routers.
Options	<i>number</i> —Maximum number of hops. Range: 2 through 255 (for an LSP); 0 through 255 (for fast reroute) Default: 255 (for an LSP); 6 (for fast reroute)
Usage Guidelines	See “Configuring Fast Reroute” on page 78 and “Limiting the Number of Hops in LSPs” on page 105.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

icmp-tunneling

Syntax	icmp-tunneling;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Enable ICMP tunneling, which can be used for debugging and tracing purposes.
Usage Guidelines	See “Configuring ICMP Message Tunneling” on page 67.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

include-all

See the following sections:

- include-all (for Administrative Groups) on page 201
- include-all (for Fast Reroute) on page 202

include-all (for Administrative Groups)

Syntax include-all [*group-names*];

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* admin-group],
 [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* (primary | secondary) *path-name* admin-group],
 [edit protocols mpls label-switched-path *lsp-name* admin-group],
 [edit protocols mpls label-switched-path *lsp-name* (primary | secondary) *path-name* admin-group]

Release Information Statement introduced before JUNOS Release 7.4.

Description Require the LSP to traverse links that include all of the defined administrative groups.

Options *group-names*—One or more names of groups defined with the admin-groups statement.

Usage Guidelines See “Configuring Administrative Groups” on page 96.

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

Related Topics admin-groups

include-all (for Fast Reroute)

Syntax	(include-all [<i>group-names</i>] no-include-all);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> fast-reroute], [edit protocols mpls label-switched-path <i>lsp-name</i> fast-reroute]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Control inclusion of administrative groups: <ul style="list-style-type: none"> ■ include-all—Define the administrative groups that must all be included for fast reroute. ■ no-include-all—Disable administrative group inclusion.
Options	<i>group-names</i> —One or more names of groups defined with the admin-groups statement.
Usage Guidelines	See “Configuring Fast Reroute” on page 78.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

include-any

See the following sections:

- include-any (for Administrative Groups) on page 203
- include-any (for Fast Reroute) on page 204

include-any (for Administrative Groups)

Syntax include-any [*group-names*];

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* admin-group],
 [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* (primary | secondary) *path-name* admin-group],
 [edit protocols mpls label-switched-path *lsp-name* admin-group],
 [edit protocols mpls label-switched-path *lsp-name* (primary | secondary) *path-name* admin-group]

Release Information Statement introduced before JUNOS Release 7.4.

Description Define the administrative groups to include for an LSP or for a path's primary and secondary paths.

Options *group-names*—One or more names of groups defined with the admin-groups statement.

Usage Guidelines See “Configuring Administrative Groups” on page 96.

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

include-any (for Fast Reroute)

Syntax	(include-any [<i>group-names</i>] no-include-any);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> fast-reroute], [edit protocols mpls label-switched-path <i>lsp-name</i> fast-reroute]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Control inclusion of administrative groups: <ul style="list-style-type: none"> ■ include-any—Define the administrative groups to include for fast reroute. ■ no-include-any—Disable administrative group inclusion.
Options	<i>group-names</i> —One or more names of groups defined with the admin-groups statement.
Usage Guidelines	See “Configuring Fast Reroute” on page 78.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

install

Syntax	install { <i>destination-prefix</i> <active>; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Associate one or more prefixes with an LSP. When the LSP is up, all the prefixes are installed as entries into the inet.3 routing table.
Options	active —(Optional) Install the route into the inet.0 routing table. This allows you to issue a ping or traceroute command on this address. <i>destination-prefix</i> —Address to associate with the LSP.
Usage Guidelines	See “Adding LSP-Related Routes to the inet.3 Routing Table” on page 80.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

interface

Syntax	<pre> interface (<i>interface-name</i> all) { disable; admin-group [<i>group-names</i>]; label-map(<i>label-map</i> default-route) { class-of-service <i>cos-value</i>; discard; next-hop (<i>address</i> <i>interface-name</i> <i>address/interface-name</i>) (discard reject); (pop swap <i>out-label</i>); preference <i>preference</i>; reject; swap <i>swap-label</i>; swap-push <i>swap-label push-label</i>; } } </pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Enable MPLS on one or more interfaces.
Options	<p><i>interface-name</i>—Name of the interface on which to configure MPLS. To configure all interfaces, specify all. For details about specifying interfaces, see the <i>JUNOS Network Interfaces Configuration Guide</i>.</p> <p>The remaining options are explained separately.</p>
Usage Guidelines	See “Configuring the Intermediate and Egress Routers for Static LSPs” on page 136.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

ipv6-tunneling

Syntax	ipv6-tunneling;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Allow IPv6 routes to be resolved over an MPLS network by converting all routes stored in the <i>inet.3</i> routing table to IPv4-compatible IPv6 addresses and then copying them into the <i>inet6.3</i> routing table. This routing table can be used to resolve next hops for both <i>inet6</i> and <i>inet6-vpn</i> routes.
Usage Guidelines	See “Enabling IPv6 Tunneling on PE Routers” on page 67.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

label-map

Syntax	label-map (<i>in-label</i> default-route) { class-of-service <i>cos-value</i> ; next-hop (<i>address</i> <i>interface-name</i> <i>address/interface-name</i>) (discard reject); (pop swap <i>out-label</i>); preference <i>preference</i> ; swap-push <i>swap-label push-label</i> ; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i>], [edit protocols mpls interface <i>interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For static MPLS only, specify the label to match.
Options	<i>in-label</i> —Incoming label value. Range: 0 through 1,048,575. Labels 0 through 999,999 are for internal use. Labels 1,000,000 through 1,048,575 are unassigned by the JUNOS Software and are available for static LSPs. When you configure static LSPs, you can use only this range of labels. The remaining statements are explained separately.
Usage Guidelines	See “Configuring the Intermediate and Egress Routers for Static LSPs” on page 136.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

label-switched-path

Syntax `label-switched-path lsp-name {`
 `disable;`
 `adaptive;`
 `admin-down;`
 `admin-group {`
 `exclude [group-names];`
 `include-all [group-names];`
 `include-any [group-names];`
 `}`
 `auto-bandwidth {`
 `adjust-interval seconds;`
 `adjust-threshold percentage;`
 `maximum-bandwidth bps;`
 `minimum-bandwidth bps;`
 `monitor-bandwidth;`
 `}`
 `bandwidth bps {`
 `ct0 bps;`
 `ct1 bps;`
 `ct2 bps;`
 `ct3 bps;`
 `}`
 `class-of-service cos-value;`
 `description text;`
 `fast-reroute {`
 `(bandwidth bps | bandwidth-percent percentage);`
 `(exclude [group-names] | no-exclude);`
 `hop-limit number;`
 `(include-all [group-names] | no-include-all);`
 `(include-any [group-names] | no-include-any);`
 `}`
 `from address;`
 `hop-limit number;`
 `install {`
 `destination-prefix/prefix-length <active>;`
 `}`
 `ldp-tunneling;`
 `link-protection;`
 `lsp-attributes {`
 `encoding-type (ethernet | packet | pdh | sonet-sdh);`
 `gpipid (ethernet | hdlc | ipv4 | ppp);`
 `signal-bandwidth type;`
 `switching-type type;`
 `}`
 `metric metric;`
 `no-cspf;`
 `no-decrement-ttl;`
 `node-link-protection;`
 `optimize-timer seconds;`
 `p2mp path-name;`
 `policing {`

```

    filter filter-name;
    no-automatic-policing;
}
preference preference;
primary path-name {
    adaptive;
    admin-group {
        exclude [ group-names ];
        include-all [ group-names ];
        include-any [ group-names ];
    }
    bandwidth bps {
        ct0 bps;
        ct1 bps;
        ct2 bps;
        ct3 bps;
    }
    class-of-service cos-value;
    hop-limit number;
    no-cspf;
    no-decrement-ttl;
    optimize-timer seconds;
    preference preference;
    priority setup-priority reservation-priority;
    (record | no-record);
    select (manual | unconditional);
    standby;
}
priority setup-priority reservation-priority;
(random | least-fill | most-fill);
(record | no-record);
retry-limit number;
retry-timer seconds;
revert-timer seconds;
secondary path-name {
    adaptive;
    admin-group {
        exclude[ group-names ];
        include-all [ group-names ];
        include-any [ group-names ];
    }
    bandwidth bps {
        ct0 bps;
        ct1 bps;
        ct2 bps;
        ct3 bps;
    }
}
class-of-service cos-value;
hop-limit number;
no-cspf;
no-decrement-ttl;
optimize-timer seconds;
preference preference;
priority setup-priority reservation-priority;
(record | no-record);
select (manual | unconditional);

```

```
        standby;
    }
    soft-preemption;
    standby;
    to address;
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
}
```

Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure an LSP to use in dynamic MPLS. When configuring an LSP, you must specify the address of the egress router in the to statement. All remaining statements are optional.
Options	<i>lsp-name</i> —Name that identifies the LSP. The name can be up to 32 characters and can contain letters, digits, periods, and hyphens. To include other characters, enclose the name in quotation marks. The name must be unique within the ingress router. The remaining statements are explained separately.
Usage Guidelines	See “LSP Configuration Overview” on page 70.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

ldp-tunneling

Syntax	ldp-tunneling;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Enable the LSP to be used for LDP tunneling.
Usage Guidelines	See “Enabling LDP over RSVP-Established LSPs” on page 362.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

least-fill

See random

link-protection (MPLS)

Syntax	link-protection;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>Enable link protection on the specified LSP, which helps to ensure that traffic sent over a specific interface to a neighboring router can continue to reach the router if that interface fails. For point-to-multipoint LSPs, including this statement extends link protection to all of the paths used by the LSP.</p> <p>To fully enable link protection, you must also include the link-protection statement at the [edit protocols rsvp interface <i>interface-name</i>] or [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i>] hierarchy level.</p>
Default	Link protection is disabled.
Usage Guidelines	See “Configuring Link Protection for Point-to-Multipoint LSPs” on page 144 and “Configuring Node Protection or Link Protection for LSPs” on page 281.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Topics	link-protection (RSVP)

log-updown

Syntax	<pre> log-updown { no-trap { mpls-lsp-traps; rfc3812-traps; } (syslog no-syslog); trap; trap-path-down; trap-path-up; } </pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
Release Information	Statement introduced before JUNOS Release 7.4. The <code>mpls-lsp-traps</code> and <code>rfc-3812-traps</code> options added in JUNOS Release 9.0.
Description	Log a message or send an SNMP trap whenever an LSP makes a transition from up to down, or vice versa, and whenever an LSP switches from one active path to another. Only the ingress router performs these operations.
Default	There is no default behavior for this statement. If you do not specify the options, the configuration cannot be committed.
Options	<p><code>no-syslog</code>—Do not log a message to the system log file.</p> <p><code>no-trap</code>—Do not send an SNMP trap.</p> <p><code>syslog</code>—Log a message to the system log file.</p> <p><code>trap</code>—Send an SNMP trap.</p> <p><code>trap-path-down</code>—Send an SNMP trap when an LSP path goes down.</p> <p><code>trap-path-up</code>—Send an SNMP trap when an LSP path comes up.</p> <p>The <code>no-trap</code> statement is explained separately.</p>
Usage Guidelines	See “Configuring System Log Messages and SNMP Traps for LSPs” on page 157 and the <i>JUNOS Network Management Configuration Guide</i> .
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Topics	<code>no-trap</code> , <code>traceoptions</code>

lsp-attributes

Syntax	lsp-attributes { encoding-type (ethernet packet pdh sonet-sdh); gpid (ethernet hdlc ipv4 pos-scrambling-crc-16 pos-no-scrambling-crc-16 pos-scrambling-crc-32 pos-no-scrambling-crc-32 ppp); signal-bandwidth <i>type</i> ; switching-type (fiber lambda psc-1 tdm); }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4. pos-scrambling-crc-16, pos-no-scrambling-crc-16, pos-scrambling-crc-32, and pos-no-scrambling-crc-32 options added in JUNOS Release 8.0.
Description	Define the parameters signaled during LSP setup. These usually determine the nature of the resource (label) allocated for the LSP. The options are explained separately.
Usage Guidelines	See “Configuring MPLS LSPs for GMPLS” on page 460.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

maximum-bandwidth

Syntax	maximum-bandwidth <i>bps</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth], [edit protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the maximum amount of bandwidth.
Options	<i>bps</i> —Bits per second.
Usage Guidelines	See “Configuring the Maximum and Minimum Bounds of the LSP’s Bandwidth” on page 90.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

metric

Syntax	<code>metric <i>metric</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Compare against another LSP or against an IGP route. To disable dynamic metric tracking, assign a fixed metric value to an LSP. If no metric is assigned, the LSP metric is dynamic and automatically tracks underlying IGP metrics.
Options	<i>metric</i> —LSP metric value. Default: No metric assigned (dynamic) Range: 1 through 16,777,215
Usage Guidelines	See “Configuring Static LSP Metrics” on page 82.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

minimum-bandwidth

Syntax	<code>minimum-bandwidth <i>bps</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth], [edit protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Set the minimum bandwidth for an LSP with automatic bandwidth allocation enabled.
Options	<i>bps</i> —Bits per second.
Usage Guidelines	See “Configuring the Maximum and Minimum Bounds of the LSP’s Bandwidth” on page 90.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

monitor-bandwidth

Syntax	monitor-bandwidth;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth], [edit protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Do not automatically adjust bandwidth allocation. However, the maximum average bandwidth utilization is monitored on the LSP, and the information is recorded in the MPLS statistics file.
Usage Guidelines	See “Configuring Passive Bandwidth Utilization Monitoring” on page 93.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

most-fill

See random

mpls

Syntax	mpls { ... }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols], [edit protocols]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Enable MPLS on the router.
Usage Guidelines	See “Minimum MPLS Configuration” on page 51.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

mtu-signaling

Syntax	mtu-signaling;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls path-mtu rsvp], [edit protocols mpls path-mtu rsvp]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Enable MTU signaling in RSVP.
Usage Guidelines	See “Enabling MTU Signaling in RSVP” on page 294.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

next-hop

Syntax	next-hop (<i>address</i> <i>interface-name</i> <i>address/interface-name</i>);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i> label-map <i>in-label</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i> label-map <i>in-label</i> default-route], [edit logical-systems <i>logical-system-name</i> protocols mpls static-path inet <i>address</i>], [edit protocols mpls interface <i>interface-name</i> label-map <i>in-label</i>], [edit protocols mpls interface <i>interface-name</i> label-map <i>in-label</i> default-route], [edit protocols mpls static-path inet <i>address</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	IP address of the next hop to the destination, specified as the IP address of the next hop, the interface name (for point-to-point interfaces only), or the <i>address/interface-name</i> to specify an IP address on an operational interface.
Options	<i>address</i> —IP address of the next-hop router. <i>interface-name</i> —IP address of the outgoing interface. It must be a point-to-point interface. The name can be a simple or fully qualified domain name.
Usage Guidelines	See “Configuring the Ingress Router for Static LSPs” on page 133 and “Configuring the Intermediate and Egress Routers for Static LSPs” on page 136.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

no-cspf

Syntax	no-cspf;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>], [edit protocols mpls], [edit protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>Disable constrained-path LSP computation.</p> <p>An explicit-path LSP is completely configured through operator action. Once configured, it is initiated only along the explicitly specified path.</p> <p>A constrained-path LSP relies on an ingress router to compute the complete path. The ingress router takes into account the following information during the computation:</p> <ul style="list-style-type: none"> ■ Interior gateway protocol (IGP) topology database ■ Link utilization information from extensions in the IGP link-state database ■ Administrative group information from extensions in the IGP link-state database ■ LSP requirements, including bandwidth, hop count, and administrative group <p>Constrained-path LSPs can generally avoid link failures and congested links. They also permit recomputation (therefore, a new path) during topology changes or unsuccessful setup.</p>
Default	Constrained-path LSP computation enabled.
Usage Guidelines	See “Disabling Constrained-Path LSP Computation” on page 95 and “Configuring Explicit-Path LSPs” on page 140.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

no-decrement-ttl

Syntax	no-decrement-ttl;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>], [edit protocols mpls], [edit protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Disable normal TTL decrementing, which decrements the TTL field in the IP header by 1. This statement decrements the IP TTL by 1 before encapsulating the IP packet within an MPLS packet. When the penultimate router pops off the top label, it does not use the standard write-back procedure of writing the MPLS TTL into the IP TTL field. Therefore, the IP packet is decremented by 1. The ultimate router then decrements the packet by one more for a total cloud appearance of 2, thus hiding the network topology.
Default	Normal TTL decrementing enabled; the TTL field value is decremented by 1 as the packet passes through each label-switched router in the LSP.
Usage Guidelines	See “Disabling Normal TTL Decrementing” on page 86.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Topics	no-propagate-ttl

no-exclude

See exclude

no-include-all

See include-all

no-include-any

See include-any

no-install-to-address

Syntax	no-install-to-address;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced in JUNOS Release 7.4.
Description	Prevent the egress router address configured using the to statement from being installed into the inet.3 and inet.0 routing tables.
Default	The egress router address for an LSP is installed into the inet.3 and inet.0 routing tables.
Usage Guidelines	See “Preventing the Addition of Egress Router Addresses to Routing Tables” on page 74.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Topics	to

no-propagate-ttl

Syntax	no-propagate-ttl;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Disable normal TTL decrementing. You configure this statement once per router, and it affects all RSVP-signaled or LDP-signaled LSPs. When this router acts as an ingress router for an LSP, it pushes an MPLS header with a TTL value of 255, regardless of the IP packet TTL. When the router acts as the penultimate router, it pops the MPLS header without writing the MPLS TTL into the IP packet.
Default	Normal TTL decrementing enabled; the TTL field value is decremented by 1 as the packet passes through each label-switched router in the LSP.
Usage Guidelines	See “Disabling Normal TTL Decrementing” on page 86.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Topics	no-decrement-ttl

no-record

See record**no-trap**

Syntax no-trap {
 mpls-lsp-traps;
 rfc-3812-traps;
}**Hierarchy Level** [edit logical-systems *logical-system-name* protocols mpls log-updown],
[edit protocols mpls log-updown]**Release Information** Statement introduced before JUNOS Release 7.4.
The mpls-lsp-traps and rfc-3812-traps options added in JUNOS Release 9.0.**Description** Prevents the transmission of SNMP traps.**Options** mpls-lsp-traps—Blocks the MPLS LSP traps defined in the jnx-mpls.mib, but allows the
rfc3812.mib traps.rfc-3812-traps—Blocks the traps defined in the rfc3812.mib, but allows the MPLS LSP
traps defined in the jnx-mpls.mib.**Usage Guidelines** See “Configuring System Log Messages and SNMP Traps for LSPs” on page 157 and
the *JUNOS Network Management Configuration Guide*.**Required Privilege Level** routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.**Related Topics** traceoptions

oam

Syntax oam {
 bfd-liveness-detection{
 failure-action teardown;
 minimum-interval *milliseconds*;
 minimum-receive-interval *milliseconds*;
 minimum-transmit-interval *milliseconds*;
 multiplier *detection-time-multiplier*;
 }
 lsp-ping-interval *seconds*;
 }

Hierarchy Level [edit protocols mpls],
 [edit protocols mpls label-switched-path *lsp-name*]
 [edit protocols mpls label-switched-path *lsp-name* primary *path-name*]

Release Information Statement introduced in JUNOS Release 7.6. The *lsp-ping-interval* option was introduced in JUNOS Release 9.4.

Description Enable Operation, Administration, and Maintenance (OAM) for RSVP-signaled LSPs.

Options *lsp-ping-interval seconds*—Specify the duration of the LSP ping interval in seconds.
 To issue a ping on an RSVP-signaled LSP, use the *ping mpls rsvp* command.

The remaining statements are explained separately.

Usage Guidelines See “Configuring BFD for MPLS IPv4 LSPs” on page 167.

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

optimize-aggressive

Syntax	optimize-aggressive;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	If enabled, the LSP reoptimization is based solely on the IGP metric. The reoptimization process ignores the available bandwidth ratio calculations, the least-fill 10 percent congestion improvement rule, and the hop-counts rule. This statement makes reoptimization more aggressive than the default.
Default	Aggressive optimization is disabled.
Usage Guidelines	See “Optimizing Signaled LSPs” on page 103.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

optimize-timer

Syntax	optimize-timer <i>seconds</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>], [edit protocols mpls], [edit protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>Enable periodic reoptimization of an LSP that is already set up. If topology changes occur, an existing path might become suboptimal, and a subsequent recomputation might be able to determine a better path. This option is useful only on LSPs for which constrained-path computation is enabled; that is, for which the no-cspf statement is not configured.</p> <p>To avoid extensive resource consumption that might result because of frequent path recomputations, or to avoid destabilizing the network as a result of constantly changing LSPs, we recommend that you either leave the timer value sufficiently large or disable the timer value.</p>
Default	The optimize timer is disabled.
Options	<i>seconds</i> —Length of the optimize timer, in seconds. Range: 0 through 65,535 seconds Default: 0 seconds (the optimize timer is disabled)
Usage Guidelines	See “Optimizing Signaled LSPs” on page 103.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

p2mp

Syntax	<code>p2mp p2mp-lsp-name;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify an LSP as either a point-to-multipoint LSP or as a branch LSP of a point-to-multipoint LSP by specifying the point-to-multipoint LSP path name.
Options	<i>p2mp-lsp-name</i> —Name of the point-to-multipoint LSP path that identifies the sequence of nodes that form the point-to-multipoint LSP. The name can contain up to 32 characters and can include letters, digits, periods, and hyphens. To include other characters or use a longer name, enclose the name in quotation marks. The name must be unique within the ingress router.
Usage Guidelines	See “Configuring the Primary Point-to-Multipoint LSP” on page 141.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

p2mp-lsp-next-hop

Syntax	<code>p2mp-lsp-next-hop point-to-multipoint-lsp-next-hop;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options static route <i>route-name</i>], [edit routing-options static route <i>route-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the name of the point-to-multipoint LSP to be used as a next hop for the static route.
Options	<i>point-to-multipoint-lsp-next-hop</i> —Name of the point-to-multipoint LSP.
Usage Guidelines	See “Configuring Static Unicast Routes for Point-to-Multipoint LSPs” on page 138.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

path

Syntax	<pre>path path-name { (address hostname) <strict loose>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>Create a named path and optionally specify the sequence of explicit routers that form the path.</p> <p>You must include this statement when configuring explicit LSPs.</p>
Options	<p>address—(Optional) IP address of each transit router in the LSP. You must specify the address or hostname of each transit router, although you do not need to list each transit router if its type is loose. As an option, you can include the ingress and egress routers in the path. Specify the addresses in order, starting with the ingress router (optional) or the first transit router, and continuing sequentially along the path until reaching the egress router (optional) or the router immediately before the egress router.</p> <p>Default: If you do not specify any routers explicitly, no routing limitations are imposed on the LSP.</p> <p>hostname—(Optional) See address.</p> <p>Default: If you do not specify any routers explicitly, no routing limitations are imposed on the LSP.</p> <p>loose—(Optional) Indicate that the next address in the path statement is a loose link. This means that the LSP can traverse through other routers before reaching this router.</p> <p>Default: strict</p> <p>path-name—Name that identifies the sequence of nodes that form an LSP. The name can contain up to 32 characters and can include letters, digits, periods, and hyphens. To include other characters or use a longer name, enclose the name in quotation marks. The name must be unique within the ingress router.</p> <p>strict—(Optional) Indicate that the LSP must go to the next address specified in the path statement without traversing other nodes. This is the default.</p>
Usage Guidelines	See “Creating Named Paths” on page 52.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Topics	static-path

path-mtu

Syntax	<pre>path-mtu { allow-fragmentation; rsvp { mtu-signaling; } }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>Configure MTU options for MPLS paths, including packet fragmentation and MTU signaling.</p> <p>The remaining statements are explained separately.</p>
Usage Guidelines	See “Configuring MTU Signaling in RSVP” on page 293.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

policing

Syntax	<pre>policing { filter <i>filter-name</i>; no-automatic-policing; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the policing filter for the LSP.
Options	<p>filter <i>filter-name</i>—Specify the name of the policing filter.</p> <p>no-automatic-policing—Disable automatic policing on this LSP.</p>
Usage Guidelines	See “Configuring Policers for LSPs” on page 160.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Topics	auto-policing

pop

Syntax	pop;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i> label-map (<i>in-label</i> default-route)], [edit protocols mpls interface <i>interface-name</i> label-map (<i>in-label</i> default-route)]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Remove the label from the top of the label stack. If there is another label in the stack, that label becomes the label at the top of the label stack. Otherwise, the packet is forwarded as a native protocol packet (typically, as an IP packet).
Usage Guidelines	See “Configuring the Intermediate and Egress Routers for Static LSPs” on page 136.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Topics	swap

preference

Syntax	<code>preference <i>preference</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i> label-map (<i>in-label</i> default-route)], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls static-path inet <i>address</i>], [edit protocols mpls], [edit protocols mpls interface <i>interface-name</i> label-map (<i>in-label</i> default-route)], [edit protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>], [edit protocols mpls static-path inet <i>address</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>Preference for the route.</p> <p>You can optionally configure multiple LSPs between the same pair of ingress and egress routers. This is useful for balancing the load among the LSPs because all LSPs, by default, have the same preference level. To prefer one LSP over another, set different preference levels for individual LSPs. The LSP with the lowest preference value is used. The default preference for LSPs is lower (more preferred) than all learned routes except direct interface routes.</p>
Options	<p><i>preference</i>—Preference to assign to the route. A route with a lower preference value is preferred.</p> <p>Range: 1 through 255</p> <p>Default: 5 for static MPLS LSPs, 7 for RSVP MPLS LSPs, 9 for LDP MPLS LSPs</p>
Usage Guidelines	See “Configuring Preference Values for LSPs” on page 98, “Configuring the Ingress Router for Static LSPs” on page 133, and “Configuring the Intermediate and Egress Routers for Static LSPs” on page 136.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

primary

Syntax `primary path-name {
 adaptive;
 admin-group {
 exclude [group-names];
 include-all [group-names];
 include-any [group-names];
 }
 bandwidth bps;
 class-of-service cos-value;
 hop-limit number;
 no-cspf;
 no-decrement-ttl;
 optimize-timer seconds;
 preference preference;
 priority setup-priority reservation-priority;
 (record | no-record);
 select (manual | unconditional);
 standby;
}`

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*],
 [edit protocols mpls label-switched-path *lsp-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify the primary path to use for an LSP. You can configure only one primary path.

You can optionally specify preference, CoS, and bandwidth values for the primary path, which override any equivalent values that you configure for the LSP (at the [edit mpls label-switched-path *lsp-name*] hierarchy level).

Options *path-name*—Name of a path that you created with the **path** statement.

The remaining statements are explained separately.

Usage Guidelines See “Configuring Primary and Secondary LSPs” on page 75.

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

priority

Syntax	<code>priority setup-priority reservation-priority;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols mpls],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i></code> <code> (primary secondary) <i>path-name</i>],</code> <code>[edit protocols mpls],</code> <code>[edit protocols mpls label-switched-path <i>lsp-name</i>],</code> <code>[edit protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>]</code>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure the setup priority and reservation priority for an LSP. If insufficient link bandwidth is available during session establishment, the setup priority is compared with other setup priorities for established sessions on the link to determine whether some of them should be preempted to accommodate the new session. Sessions with lower hold priorities are preempted.
Options	<p><i>reservation-priority</i>—Reservation priority, used to keep a reservation after it has been set up. A smaller number has a higher priority. The priority must be greater than or equal to the setup priority to prevent preemption loops.</p> <p>Range: 0 through 7, where 0 is the highest and 7 is the lowest priority.</p> <p>Default: 0 (Once the session is set up, no other session can preempt it.)</p> <p><i>setup-priority</i>—Setup priority.</p> <p>Range: 0 through 7, where 0 is the highest and 7 is the lowest priority.</p> <p>Default: 7 (The session cannot preempt any existing sessions.)</p>
Usage Guidelines	See “Configuring Priority and Preemption for LSPs” on page 102.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

push

Syntax	<code>push out-label;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls static-path inet <i>address</i>], [edit protocols mpls static-path inet <i>address</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Add a new label to the top of the label stack.
Options	<p><i>out-label</i>—Label value.</p> <p>Range: 0 through 1,048,575. Labels 0 through 999,999 are for internal use. Labels 1,000,000 through 1,048,575 are unassigned by the JUNOS Software and are available for static LSPs. When you configure static LSPs, you can use only this range of labels.</p>
Usage Guidelines	See “Configuring the Ingress Router for Static LSPs” on page 133.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

random

Syntax	(random least-fill most-fill);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>Configure the preferred path when several equal-cost candidate paths to a destination exist, and prefer the path with the highest available bandwidth (with the largest minimum available bandwidth ratio). The available bandwidth ratio of a link is the available bandwidth on a link divided by the maximum reservable bandwidth on the link.</p> <ul style="list-style-type: none"> ■ least-fill—Prefer the path with the most available bandwidth (with the largest minimum available bandwidth ratio). ■ most-fill—Prefer the path with the least available bandwidth (with the minimum available bandwidth ratio). The minimum available bandwidth ratio of a path is the smallest available bandwidth ratio belonging to any of the links in the path. ■ random—Choose the path at random.
Default	random
Usage Guidelines	See “Configuring CSPF Tie Breaking” on page 83.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

record

Syntax	(record no-record);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path (primary secondary) <i>path-name</i>], [edit protocols mpls], [edit protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify whether an LSP should actively record the routes in the path. Recording routes requires that all transit routers support the RSVP Record Route object. Recording routes can be useful for diagnostics and loop detection.
Default	Record routes.
Usage Guidelines	See “Disabling Path Route Recording” on page 98.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

reject

Syntax	reject;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i> label-map (<i>in-label</i> default-route)], [edit protocols mpls interface <i>interface-name</i> label-map (<i>in-label</i> default-route)]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Do not forward a packet with the matching incoming label. Instead, drop the packet and, for IP packets, send an ICMP unreachable message to the packet’s originator.
Usage Guidelines	See “Configuring the Intermediate and Egress Routers for Static LSPs” on page 136.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

retry-limit

Syntax	<code>retry-limit <i>number</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>],
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Maximum number of times the ingress router tries to establish the primary path. This counter is reset each time a primary path is created successfully. When the limit is exceeded, no more connection attempts are made. Intervention is then required to restart the connection.
Options	<i>number</i> —Maximum number of tries to establish the primary path. Range: 0 through 10,000 Default: 0 (The ingress node never stops trying to establish the primary path.)
Usage Guidelines	See “Configuring the Connection Between Ingress and Egress Routers” on page 81.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

retry-timer

Syntax	<code>retry-timer <i>seconds</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Amount of time the ingress router waits between attempts to establish the primary path.
Options	<i>seconds</i> —Amount of time between attempts to connect to the primary path. Range: 1 through 600 seconds Default: 30 seconds
Usage Guidelines	See “Configuring the Connection Between Ingress and Egress Routers” on page 81.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

revert-timer

Syntax	<code>revert-timer seconds;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4. BFD behavior modified in JUNOS Release 9.0.
Description	<p>Specify the amount of time (in seconds) that an LSP must wait before traffic reverts to a primary path. If during this time the primary path experiences any connectivity problem or stability problem, the timer is restarted.</p> <p>If you have configured BFD on the LSP, the JUNOS Software waits until the BFD session is restored before starting the revert timer counter.</p> <p>If you have configured a value of 0 seconds for the revert-timer statement and traffic is switched to the secondary path, the traffic remains on that path indefinitely. It is never switched back to the primary path unless you intervene.</p>
Options	<p><i>seconds</i>—Time in seconds.</p> <p>Range: 0 through 65,535 seconds</p> <p>Default: 60 seconds</p>
Usage Guidelines	See “Configuring the Revert Timer for LSPs” on page 76.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

rpf-check-policy

Syntax	<code>rpf-check-policy policy;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options multicast], [edit routing-options multicast]
Release Information	Statement introduced in JUNOS Release 8.0.
Description	Allows you to control whether an reverse path forwarding (RPF) check is performed for a source and group entry before installing a route in the multicast forwarding cache. This makes it possible to use point-to-multipoint LSPs to distribute multicast traffic to PIM islands situated downstream from the egress routers of the point-to-multipoint LSPs.
Options	<i>policy</i> —Name of the RPF check routing policy.
Usage Guidelines	See “Configuring a Multicast RPF Check Policy for Point-to-Multipoint LSPs” on page 145.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

rsvp-error-hold-time

Syntax	rsvp-error-hold-time <i>seconds</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>Amount of time MPLS retains RSVP PathErr messages and considers them for CSPF computations. The more time you configure, the more time a source node (ingress of an RSVP LSP) can have to learn about the failures of its LSP by monitoring PathErr messages transmitted from downstream nodes.</p> <p>Information from the PathErr messages is incorporated into subsequent LSP computations, which can improve the accuracy and speed of LSP setup. Some PathErr messages are also used to update traffic engineering database bandwidth information, reducing inconsistencies between the database and the network.</p>
Options	<p><i>seconds</i>—Amount of time MPLS retains RSVP PathErr messages and considers them for CSPF computations.</p> <p>Range: 0 through 240 seconds</p> <p>Default: 25 seconds</p>
Usage Guidelines	See “Improving Traffic Engineering Database Accuracy with RSVP PathErr Messages” on page 60.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

secondary

Syntax `secondary path-name {
 adaptive;
 admin-group {
 exclude [group-names];
 include-all [group-names];
 include-any [group-names];
 }
 bandwidth bps;
 class-of-service cos-value;
 hop-limit number;
 no-cspf;
 no-decrement-ttl;
 optimize-timer seconds;
 preference preference;
 priority setup-priority reservation-priority;
 (record | no-record);
 retry-limit number;
 retry-timer seconds;
 select (manual | unconditional);
 standby;
}`

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*],
 [edit protocols mpls label-switched-path *lsp-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify one or more secondary paths to use for the LSP. You can configure more than one secondary path. All secondary paths are equal, and the first one that is available is chosen.

You can specify secondary paths even if you have not specified any primary paths.

Optionally, you can specify preference, CoS, and bandwidth values for the secondary path, which override any equivalent values that you configure for the LSP (at the [edit mpls label-switched-path] hierarchy level).

Options *path-name*—Name of a path that you created with the `path` statement.

The remaining statements are explained separately.

Usage Guidelines See “Configuring Primary and Secondary LSPs” on page 75.

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

select

Syntax	select (manual unconditional);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the conditions under which the path is selected to carry traffic. The manual and unconditional options are mutually exclusive.
Options	<p>manual—The path is selected for carrying traffic if it is up and stable for at least the revert timer window (potentially before the revert timer has elapsed). Traffic is sent to other working paths if the current path is down or degraded (receiving errors).</p> <p>unconditional—The path is always selected for carrying traffic, even if it is currently down or degraded (receiving errors).</p>
Usage Guidelines	See “Specifying the Conditions for Path Selection” on page 77.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

signal-bandwidth

Syntax	signal-bandwidth <i>type</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> <i>lsp-attributes</i>], [edit protocols mpls label-switched-path <i>lsp-name</i> <i>lsp-attributes</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the bandwidth encoding of the signal used for path computation and admission control.
Options	type —Configure the type of bandwidth encoding used on the LSP. It can be any of the following values: 10gigether , ds1 , ds3 , e1 , e3 , ethernet , fastether , gigether , stm-1 , stm-4 , stm-16 , stm-64 , stm-256 , sts-1 , vt1-5 , or vt2 .
Usage Guidelines	See “Configuring the Signal Bandwidth Type” on page 461.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

smart-optimize-timer

Syntax	smart-optimize-timer <i>seconds</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	When a link fails and traffic is moved to an alternate path, the smart optimize timer waits for the specified period of time and then switches the traffic back to the original path (if that path is back up). If the original path fails again, the traffic is shifted to an alternate path and the smart optimization timer is disabled for one hour. This statement is disabled by default.
Options	<i>seconds</i> —Specify the number of seconds for the smart optimize timer. Range: 0 through 65,535 seconds Default: 180 seconds
Usage Guidelines	See “Configuring the Smart Optimize Timer” on page 105.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

soft-preemption

Syntax	soft-preemption;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Soft preemption attempts to establish a new path for a preempted LSP before tearing it down.
Usage Guidelines	See “Configuring MPLS Soft Preemption” on page 87.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

standby

Syntax	standby;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>], [edit protocols mpls], [edit protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Have the path remain up at all times to provide instant switchover if connectivity problems occur.
Usage Guidelines	See “Configuring Hot Standby of Secondary Paths” on page 106.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

static-path

Syntax static-path inet {
 prefix {
 class-of-service *cos-value*;
 double-push *bottom-label top-label*;
 next-hop (*address* | *interface-name* | *address/interface-name*);
 preference *preference*;
 push *out-label*;
 triple-push *bottom-label middle-label top-label*;
 }
 }

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls],
 [edit protocols mpls]

Release Information Statement introduced before JUNOS Release 7.4.

Description Statically configure an LSP. You configure the LSP on the ingress router only.

You can specify one or more **static-path** statements.

Options *prefix*—IP address that matches the packet’s destination field. You can specify one or more addresses. You can specify the prefix in one of the following ways:

- IP address; for example, 10.0.0.2
- Range of IP addresses; for example, 10.0.0.0/8

inet—Configure the path for packets with IPv4 destinations.

The remaining statements are explained separately.

Usage Guidelines See “Configuring the Ingress Router for Static LSPs” on page 133.

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

statistics

Syntax	<pre>statistics { auto-bandwidth; file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; interval <i>seconds</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Enable MPLS statistics collection and reporting.
Options	<p>auto-bandwidth—Collect statistics related to automatic bandwidth.</p> <p>file <i>filename</i>—Name of the file to receive the output. We recommend that you place MPLS tracing output in the file <code>mpls-stat</code> in the <code>/var/log</code> directory.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named <i>file</i> reaches its maximum size, it is renamed <i>file.0</i>, then <i>file.1</i>, and so on, until the maximum number of files is reached. Then, the oldest file is overwritten. Range: 2 or more Default: 2 files</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>interval <i>seconds</i>—Interval at which to periodically collect statistics. Range: 1 through 65,535 Default: 300 seconds</p> <p>no-world-readable—(Optional) Prevents users from reading the log file.</p> <p>size <i>size</i>—(Optional) Maximum size of each file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a file named <i>file</i> reaches this size, it is renamed <i>file.0</i>. When the <i>file</i> again reaches its maximum size, <i>file.0</i> is renamed <i>file.1</i> and <i>file</i> is renamed <i>file.0</i>. This renaming scheme continues until the maximum number of files is reached. Then the oldest trace file is overwritten. Syntax: Syntax: <i>xk</i> to specify KB, <i>xm</i> to specify MB, or <i>xg</i> to specify GB Range: 10 KB through the maximum file size supported on your system Default: 1 MB</p> <p>If you specify a maximum file size, you also must specify a maximum number of files with the files option.</p> <p>world-readable—(Optional) Enables users to read the log file.</p>
Usage Guidelines	See “Configuring MPLS to Gather Statistics” on page 156, and “Configuring MPLS Statistics for Automatic Bandwidth Allocation” on page 89.

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

swap

Syntax	swap <i>out-label</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i> label-map (<i>in-label</i> default-route)], [edit protocols mpls interface <i>interface-name</i> label-map (<i>in-label</i> default-route)]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Remove the label at the top of the label stack and replace it with the specified label.
Options	<i>out-label</i> —Label value. Range: 0 through 1,048,575. Labels 0 through 999,999 are for internal use. Labels 1,000,000 through 1,048,575 are unassigned by the JUNOS Software and are available for static LSPs. When you configure static LSPs, you can use only this range of labels. Default: If you do not define the <i>out-label</i> option, the original label value remains unchanged.
Usage Guidelines	See “Configuring the Intermediate and Egress Routers for Static LSPs” on page 136.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Topics	pop

swap-push

Syntax	<code>swap-push swap-label push-label;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i> label-map <i>in-label</i> default-route]], [edit protocols mpls interface <i>interface-name</i> label-map <i>in-label</i> default-route]]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Replace the stack's top label and then push one more label on top of the stack.
Options	<i>push-label</i> —Label value. <i>swap-label</i> —Label value. The following range and default values apply to both <i>push-label</i> and <i>swap-label</i> : Range: 0 through 1,048,575. Labels 0 through 999,999 are for internal use. Labels 1,000,000 through 1,048,575 are unassigned by the JUNOS Software and are available for static LSPs. When you configure static LSPs, you can use only this range of labels. Default: If you do not define the <i>swap-label</i> option, the original label value remains unchanged.
Usage Guidelines	See “Configuring the Intermediate and Egress Routers for Static LSPs” on page 136.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Topics	pop and swap

switching-type

Syntax	switching-type (fiber lambda psc-1 tdm);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> lsp-attributes], [edit protocols mpls label-switched-path <i>lsp-name</i> lsp-attributes]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the switching method for the LSP. The switching method can be one of the following values: <ul style="list-style-type: none"> ■ fiber—Fiber switching ■ lambda—Lambda switching ■ psc-1—Packet switching ■ tdm—Time-division multiplexing (TDM) switching
Default	psc-1
Usage Guidelines	See “Configuring MPLS LSPs for GMPLS” on page 460.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

te-class-matrix

Syntax

```
te-class-matrix {
  tnumber {
    priority priority;
    traffic-class {
      ctnumber priority priority;
    }
  }
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls diffserv-te],
[edit protocols mpls diffserv-te]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify the traffic engineering class matrix for a multiclass LSP or a DiffServ-aware traffic engineering LSP.

Default The default traffic engineering class matrix is:

```
te-class-matrix {
  te0 traffic-class ct0 priority 7;
  te1 traffic-class ct1 priority 7;
  te2 traffic-class ct2 priority 7;
  te3 traffic-class ct3 priority 7;
  te4 traffic-class ct0 priority 0;
  te5 traffic-class ct1 priority 0;
  te6 traffic-class ct2 priority 0;
  te7 traffic-class ct3 priority 0;
}
```

If you define any of the traffic engineering classes, all the default values are dropped.

Options *ctnumber*—Specify the number of the class type. It can be one of four values: *ct0*, *ct1*, *ct2*, or *ct3*.

priority priority—Specify the priority of the class type. It can be one of eight values from 0 through 7.

tnumber—Specify the number of the traffic engineering class. It can be one of eight values: *te0*, *te1*, *te2*, *te3*, *te4*, *te5*, *te6*, or *te7*. You must configure the traffic engineering classes in order, starting with *te0*.

traffic-class—Specify the traffic class for the traffic engineering class.

Usage Guidelines See “Configuring Traffic Engineering Classes” on page 117.

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

to

Syntax	<code>to address;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the egress router of a dynamic LSP.
Options	<i>address</i> —Address of the egress router.
Usage Guidelines	See “Configuring the Egress Router Address for LSPs” on page 74.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

traceoptions

Syntax traceoptions {
 file *filename* <files *number*> <size *size*> <world-readable | no-world-readable>;
 flag *flag*;
 }

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls],
 [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*],
 [edit protocols mpls],
 [edit protocols mpls label-switched-path *lsp-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure MPLS tracing options at the protocol level or for a label-switched path.

To specify more than one tracing operation, include multiple **flag** statements.

Default The default MPLS protocol-level tracing options are inherited from the routing protocols **traceoptions** statement included at the [edit **routing-options**] hierarchy level.

Options *filename*—Name of the file to receive the output of the tracing operation. All files are placed in the directory `/var/log`. We recommend that you place MPLS tracing output in the file `mpls-log`.

files number—(Optional) Maximum number of trace files. When a trace file named *trace-file* reaches its maximum size, it is renamed *trace-file.0*, then *trace-file.1*, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Range: 2 through 1000

Default: 2 files

If you specify a maximum number of files, you must also include the **size** statement to specify the maximum file size.

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

MPLS Tracing Flags

- **all**—Trace all operations
- **connection**—All circuit cross-connect (CCC) activity
- **connection-detail**—Detailed CCC activity
- **cspf**—CSPF computations
- **cspf-link**—Links visited during CSPF computations
- **cspf-node**—Nodes visited during CSPF computations
- **error**—MPLS error packets
- **graceful-restart**—Trace MPLS graceful restart events

- **lsping**—Trace lsping packets and return codes
- **state**—All LSP state transitions
- **timer**—Timer usage

no-world-readable—(Optional) Allow only certain users to read the log file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *trace-file* reaches this size, it is renamed *trace-file.0*. When the *trace-file* again reaches this size, *trace-file.0* is renamed *trace-file.1* and *trace-file* is renamed *trace-file.0*. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Syntax: *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 1 MB

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of files.

world-readable—(Optional) Allow any user to read the log file.

Usage Guidelines See “Tracing MPLS and LSP Packets and Operations” on page 171.

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

traffic-engineering

Syntax	traffic-engineering (bgp bgp-igp bgp-igp-both-ribs mpls-forwarding);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Select whether MPLS performs traffic engineering on BGP destinations only or on both BGP and IGP destinations. Affects only LSPs originating from this router, not transit or egress LSPs.
Default	bgp
Options	<p>bgp—On BGP destinations only. Ingress routes are installed in the inet.3 routing table.</p> <p>bgp-igp—On both BGP and IGP destinations. Ingress routes are installed in the inet.0 routing table. If IGP shortcuts are enabled, the shortcut routes are automatically installed in the inet.0 routing table.</p> <p>bgp-igp-both-ribs—On both BGP and IGP destinations. Ingress routes are installed in the inet.0 and inet.3 routing tables. This option is used to support VPNs.</p> <p>mpls-forwarding—On both BGP and IGP destinations. Use ingress routes for forwarding only, not for routing.</p>
Usage Guidelines	See “Configuring Traffic Engineering for LSPs” on page 150.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

triple-push

Syntax	<code>triple-push <bottom-label> <middle-label> <top-label>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls static-path inet <i>prefix</i>], [edit protocols mpls static-path inet <i>prefix</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Push three more labels on top of the stack.
Options	<p><i>bottom-label</i>—(Optional) Label value.</p> <p><i>middle-label</i>—(Optional) Label value.</p> <p><i>top-label</i>—(Optional) Label value.</p> <p>The labels all have the same range and default actions:</p> <p>Range: 0 through 1,048,575. Labels 0 through 999,999 are for internal use. Labels 1,000,000 through 1,048,575 are unassigned by the JUNOS Software and are available for static LSPs. When you configure static LSPs, you can use only this range of labels.</p> <p>Default: If you do not define the <i>push-label</i> option, the original label value remains unchanged.</p>
Usage Guidelines	See “Configuring the Ingress Router for Static LSPs” on page 133.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

Part 3

RSVP

- RSVP Overview on page 255
- RSVP Configuration Guidelines on page 273
- Summary of RSVP Configuration Statements on page 299

Chapter 11

RSVP Overview

This chapter discusses the following topics:

- RSVP Introduction on page 256
- RSVP Standards on page 256
- JUNOS Software RSVP Protocol Implementation on page 257
- RSVP Operation on page 258
- RSVP Operation Overview on page 258
- RSVP Authentication on page 259
- RSVP and IGP Hello Packets and Timers on page 259
- RSVP Message Types on page 259
- Path Messages on page 260
- Resv Messages on page 260
- PathTear Messages on page 260
- ResvTear Messages on page 260
- PathErr Messages on page 261
- ResvErr Messages on page 261
- ResvConfirm Messages on page 261
- RSVP Reservation Styles on page 261
- RSVP Refresh Reduction on page 262
- MTU Signaling in RSVP on page 263
- How the Correct MTU Is Signaled in RSVP on page 264
- Determining an Outgoing MTU Value on page 265
- MTU Signaling in RSVP Limitations on page 265
- Link Protection on page 265
- Fast Reroute, Node Protection, and Link Protection on page 266
- Multiple Bypass LSPs on page 267
- Node Protection on page 267
- RSVP Graceful Restart on page 268
- RSVP Graceful Restart Standard on page 269
- RSVP Graceful Restart Terminology on page 269

- RSVP Graceful Restart Operation on page 270
- Processing the Restart Cap Object on page 271

RSVP Introduction

RSVP is a resource reservation setup protocol that is used by both network hosts and routers. Hosts use RSVP to request a specific class of service (CoS) from the network for particular application flows. Routers use RSVP to deliver CoS requests to all routers along the data path. RSVP also can maintain and refresh states for a requested CoS application flow.

RSVP treats an application flow as a simplex connection. That is, the CoS request travels only in one direction—from the sender to the receiver. RSVP is a transport layer protocol that uses IP as its network layer. However, RSVP does not transport application flows. Rather, it is more of an Internet control protocol, similar to the Internet Control Message Protocol (ICMP) and Internet Group Management Protocol (IGMP). RSVP runs as a separate software process in the JUNOS Software and is not in the packet forwarding path.

RSVP is not a routing protocol, but rather is designed to operate with current and future unicast and multicast routing protocols. The routing protocols are responsible for choosing the routes to use to forward packets, and RSVP consults local routing tables to obtain routes. RSVP only ensures the CoS of packets traveling along a data path.

The receiver in an application flow requests the preferred CoS from the sender. To do this, the receiver issues an RSVP CoS request on behalf of the local application. The request propagates to all routers in reverse direction of the data paths toward the sender. In this process, RSVP requests might be merged, resulting in a protocol that scales well when there are a large number of receivers.

Because the number of receivers in an application flow is likely to change and the flow of delivery paths might change during the life of an application flow, RSVP takes a soft-state approach in its design, creating and removing the protocol states in routers and hosts incrementally over time. RSVP sends periodic refresh messages to maintain its state and to recover from occasional lost messages. In the absence of refresh messages, RSVP states automatically time out and are deleted.

RSVP Standards

RSVP is described in several RFCs and drafts.

The following RFCs provide an overview of RSVP and RSVP features:

- RFC 2205, *Resource Reservation Protocol (RSVP), Version 1, Functional Specification*
- RFC 2209, *Resource Reservation Protocol (RSVP), Version 1, Message Processing Rules*
- RFC 2210, *The Use of RSVP with IETF Integrated Services*
- RFC 2211, *Specification of the Controlled-Load Network Element Service*
- RFC 2212, *Specification of Guaranteed Quality of Service*

- RFC 2215, *General Characterization Parameters for Integrated Service Network Elements*
- RFC 2216, *Network Element Service Specification Template*
- RFC 2745, *RSVP Diagnostic Messages*
- RFC 2747, *RSVP Cryptographic Authentication* (see also RFC 3097)
- RFC 2961, *RSVP Refresh Overhead Reduction Extensions*
- RFC 3097, *RSVP Cryptographic Authentication -- Updated Message Type Value* (see also RFC 2747)
- RFC 3209, *RSVP-TE: Extensions to RSVP for LSP Tunnels* (the JUNOS Software does not support the Null Service Object for maximum transmission unit [MTU] signaling in RSVP)
- RFC 3473, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Extensions* (only Section 9, *Fault Handling*)
- RFC 3477, *Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)*
- RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels* (except node protection in facility backup)
- RFC 4125, *Maximum Allocation Bandwidth Constraints Model for Diff-Serv-aware MPLS Traffic Engineering*
- RFC 4127, *Russian Dolls Bandwidth Constraints Model for Diff-Serv-aware MPLS Traffic Engineering*
- RFC 4203, *OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)* (OSPF extensions can carry traffic engineering information over unnumbered links)
- RFC 4205, *Intermediate System to Intermediate System (IS-IS) Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)* (IS-IS extensions can carry traffic engineering information over unnumbered links)
- RFC 4561, *Definition of a Record Route Object (RRO) Node-Id Sub-Object* (The Record Route Object [RRO] node ID subobject is for use in inter-AS link and node protection configurations.)

The following Internet draft also provides information about RSVP:

- Internet draft draft-ietf-mpls-rsvp-te-p2mp-01.txt, *Extensions to RSVP-TE for Point to Multipoint TE LSPs* (expires June 2005)

To access RFCs and drafts, go to the IETF website at <http://www.ietf.org>.

JUNOS Software RSVP Protocol Implementation

The JUNOS implementation of RSVP supports RSVP version 1. The software includes support for all mandatory objects and RSVP message types, and supports message integrity and node authentications through the Integrity object.

The primary purpose of the JUNOS RSVP software is to support dynamic signaling within MPLS label-switched paths (LSPs). Supporting resource reservations over the Internet is only a secondary purpose of the JUNOS Software implementation. Since supporting resource reservations is secondary, the JUNOS RSVP software does not support the following features:

- IP multicasting sessions.
- Traffic control. The software cannot make resource reservations for real-time video or audio sessions.

With regard to the protocol mechanism, packet processing, and RSVP objects supported, the JUNOS Software implementation of the software is interoperable with other RSVP implementations.

RSVP Operation

The following sections describe RSVP operation:

- RSVP Operation Overview on page 258
- RSVP Authentication on page 259
- RSVP and IGP Hello Packets and Timers on page 259

RSVP Operation Overview

RSVP creates independent sessions to handle each data flow. A session is identified by a combination of the destination address, an optional destination port, and a protocol. Within a session, there can be one or more senders. Each sender is identified by a combination of its source address and source port. An out-of-band mechanism, such as a session announcement protocol or human communication, is used to communicate the session identifier to all senders and receivers.

A typical RSVP session involves the following sequence of events:

1. A potential sender starts sending RSVP path messages to the session address.
2. A receiver, wanting to join the session, registers itself if necessary. For example, a receiver in a multicast application would register itself with IGMP.
3. The receiver receives the path messages.
4. The receiver sends appropriate Resv messages toward the sender. These messages carry a flow descriptor, which is used by routers along the path to make reservations in their link-layer media.
5. The sender receives the Resv message and then starts sending application data.

This sequence of events is not necessarily strictly synchronized. For example, receivers can register themselves before receiving path messages from the sender, and application data can flow before the sender receives Resv messages. Application data that is delivered before the actual reservation contained in the Resv message typically is treated as best-effort, non-real-time traffic with no CoS guarantee.

RSVP Authentication

JUNOS Software supports both the RSVP authentication style described in RFC 2747 (allowing for multivendor compatibility) and the RSVP authentication style described in Internet draft draft-ietf-rsvp-md5-03.txt. The JUNOS Software uses the authentication style described in Internet draft draft-ietf-rsvp-md5-08.txt by default. If the router receives an RFC 2747-style RSVP authentication from a neighbor, it switches to this style of authentication for that neighbor. The RSVP authentication style for each neighboring router is determined separately.

RSVP and IGP Hello Packets and Timers

RSVP monitors the status of the interior gateway protocol (IGP) (IS-IS or OSPF) neighbors and relies on the IGP protocols to detect when a node fails. If an IGP protocol declares a neighbor down (because hello packets are no longer being received), RSVP also brings down that neighbor. However, the IGP protocols and RSVP still act independently when bringing a neighbor up.

In JUNOS Software, RSVP typically relies on IGP hello packet detection to check for node failures. RSVP sessions are kept up even if RSVP hello packets are no longer being received, so long as the router continues to receive IGP hello packets. RSVP sessions are maintained until either the router stops receiving IGP hello packets or the RSVP Path and Resv messages time out. Configuring a short time for the IS-IS or OSPF hello timers allows these protocols to detect node failures quickly.

RSVP hellos can be relied on when the IGP does not recognize a particular neighbor (for example, if IGP is not enabled on the interface) or if the IGP is RIP (not IS-IS or OSPF). Also, the equipment of other vendors might be configured to monitor RSVP sessions based on RSVP hello packets. This equipment might also take an RSVP session down due to a loss of RSVP hello packets.

Juniper Networks does not recommend configuring a short RSVP hello timer. If quick discovery of a failed neighbor is needed, configure short IGP (OSPF or IS-IS) hello timers.

OSPF and IS-IS have infrastructure to manage rapid hello message sending and receiving reliably, even if the routing protocols or some other process are straining the processing capability of the router. Under the same circumstances, RSVP hellos might timeout prematurely even though the neighbor is functioning normally.

RSVP Message Types

RSVP uses the following types of messages to establish and remove paths for data flows, establish and remove reservation information, confirm the establishment of reservations, and report errors:

- Path Messages on page 260
- Resv Messages on page 260
- PathTear Messages on page 260

- ResvTear Messages on page 260
- PathErr Messages on page 261
- ResvErr Messages on page 261
- ResvConfirm Messages on page 261

Path Messages

Each sender host transmits path messages downstream along the routes provided by the unicast and multicast routing protocols. Path messages follow the exact paths of application data, creating path states in the routers along the way, thus enabling routers to learn the previous-hop and next-hop node for the session. Path messages are sent periodically to refresh path states.

The refresh interval is controlled by a variable called the *refresh-time*, which is the periodical refresh timer expressed in seconds. A path state times out if a router does not receive a specified number of consecutive path messages. This number is specified by a variable called *keep-multiplier*. Path states are kept for $(\text{keep-multiplier} + 0.5) \times 1.5 \times \text{refresh-time}$ seconds.

Resv Messages

Each receiver host sends reservation request (Resv) messages upstream toward senders and sender applications. Resv messages must follow exactly the reverse path of path messages. Resv messages create and maintain a reservation state in each router along the way.

Resv messages are sent periodically to refresh reservation states. The refresh interval is controlled by the same refresh time variable, and reservation states are kept for $(\text{keep-multiplier} + 0.5) \times 1.5 \times \text{refresh-time}$ seconds.

PathTear Messages

PathTear messages remove (tear down) path states as well as dependent reservation states in any routers along a path. PathTear messages follow the same path as path messages. A PathTear typically is initiated by a sender application or by a router when its path state times out.

PathTear messages are not required, but they enhance network performance because they release network resources quickly. If PathTear messages are lost or not generated, path states eventually time out when they are not refreshed, and the resources associated with the path are released.

ResvTear Messages

ResvTear messages remove reservation states along a path. These messages travel upstream toward senders of the session. In a sense, ResvTear messages are the reverse of Resv messages. ResvTear messages typically are initiated by a receiver application or by a router when its reservation state times out.

ResvTear messages are not required, but they enhance network performance because they release network resources quickly. If ResvTear messages are lost or not generated, reservation states eventually time out when they are not refreshed, and the resources associated with the reservation are released.

PathErr Messages

When path errors occur (usually because of parameter problems in a path message), the router sends a unicast PathErr message to the sender that issued the path message. PathErr messages are advisory; these messages do not alter any path state along the way.

ResvErr Messages

When a reservation request fails, a ResvErr error message is delivered to all the receivers involved. ResvErr messages are advisory; these messages do not alter any reservation state along the way.

ResvConfirm Messages

Receivers can request confirmation of a reservation request, and this confirmation is sent with a ResvConfirm message. Because of the complex RSVP flow-merging rules, a confirmation message does not necessarily provide end-to-end confirmation of the entire path. Therefore, ResvConfirm messages are an indication, not a guarantee, of potential success.

Juniper Networks routers never request confirmation using the ResvConfirm message; however, a Juniper Networks router can send a ResvConfirm message if it receives a request from another vendor's equipment.

RSVP Reservation Styles

A reservation request includes options for specifying the reservation style. The reservation styles define how reservations for different senders within the same session are treated and how senders are selected.

Two options specify how reservations for different senders within the same session are treated:

- Distinct reservation—Each receiver establishes its own reservation with each upstream sender.
- Shared reservation—All receivers make a single reservation that is shared among many senders.

Two options specify how senders are selected:

- Explicit sender—List all selected senders.
- Wildcard sender—Select all senders, which then participate in the session.

The following reservation styles, formed by a combination of these four options, currently are defined:

- **Fixed filter (FF)**—This reservation style consists of distinct reservations among explicit senders. Examples of applications that use fixed-filter-style reservations are video applications and unicast applications, which both require flows that have a separate reservation for each sender. The fixed filter reservation style is enabled on RSVP LSPs by default.
- **Wildcard filter (WF)**—This reservation style consists of shared reservations among wildcard senders. This type of reservation reserves bandwidth for any and all senders, and propagates upstream toward all senders, automatically extending to new senders as they appear. A sample application for wildcard filter reservations is an audio application in which each sender transmits a distinct data stream. Typically, only a few senders are transmitting at any one time. Such a flow does not require a separate reservation for each sender; a single reservation is sufficient.
- **Shared explicit (SE)**—This reservation style consists of shared reservations among explicit senders. This type of reservation reserves bandwidth for a limited group of senders. A sample application is an audio application similar to that described for wildcard filter reservations.

RSVP Refresh Reduction

RSVP relies on soft-state to maintain the path and reservation states on each router. If the corresponding refresh messages are not sent periodically, the states eventually time out and reservations are deleted. RSVP also sends its control messages as IP datagrams with no reliability guarantee. It relies on periodic refresh messages to handle the occasional loss of Path or Resv messages.

The RSVP refresh reduction extensions, based on RFC 2961, addresses the following problems that result from relying on periodic refresh messages to handle message loss:

- **Scalability**—The scaling problem arises from the periodic transmission and processing overhead of refresh messages, which increases as the number of RSVP sessions increases.
- **Reliability and latency**—The reliability and latency problem stems from the loss of nonrefresh RSVP messages or one-time RSVP messages such as PathTear or PathErr. The time to recover from such a loss is usually tied to refresh interval and the keepalive timer.

The RSVP refresh reduction capability is advertised by enabling the refresh reduction (RR) capable bit in the RSVP common header. This bit is only significant between RSVP neighbors.

RSVP refresh reduction includes the following features:

- RSVP message bundling using the bundle message
- RSVP Message ID to reduce message processing overhead

- Reliable delivery of RSVP messages using Message ID, Message Ack, and Message Nack
- Summary refresh to reduce the amount of information transmitted every refresh interval

The RSVP refresh reduction specification (RFC 2961) allows you to enable some or all of the above capabilities on a router. It also describes various procedures that a router can use to detect the refresh reduction capabilities of its neighbor.

The JUNOS Software supports all of the refresh reduction extensions, some of which can be selectively enabled or disabled. The JUNOS Software supports Message ID and therefore can perform reliable message delivery only for Path and Resv messages.

For information on how to configure RSVP refresh reduction, see “Configuring RSVP Refresh Reduction” on page 275.

MTU Signaling in RSVP

The maximum transmission unit (MTU) is the largest size packet or frame, in bytes, that can be sent in a network. An MTU that is too large might cause retransmissions. Too small an MTU might cause the router to send and handle relatively more header overhead and acknowledgments. There are default values for MTUs associated with various protocols. You can also explicitly configure an MTU on an interface.

When an LSP is created across a set of links with different MTU sizes, the ingress router does not know what the smallest MTU is on the LSP path. By default, the maximum packet size for the LSP is based on the MTU for the outgoing interface for the LSP on the ingress router.

If this MTU is larger than the MTU of one of the intermediate links, traffic might be dropped, because MPLS packets cannot be fragmented. Also, the ingress router is not aware of this type of traffic loss, because the control plane for the LSP would still function normally.

To prevent this type of packet loss in MPLS LSPs, you can configure MTU signaling in RSVP. This feature is described in RFC 3209. Juniper Networks supports the Integrated Services object for MTU signaling in RSVP. The Integrated Services object is described in RFCs 2210 and 2215. MTU signaling in RSVP is disabled by default.

To avoid packet loss due to MTU mismatches, the ingress router needs to do the following:

- Signal the MTU on the RSVP LSP—To prevent packet loss from an MTU mismatch, the ingress router needs to know what the smallest MTU value is along the path taken by the LSP. Once this MTU value is obtained, the ingress router can assign it to the LSP.
- Fragment packets—Using the assigned MTU value, packets that exceed the size of the MTU can be fragmented into smaller packets on the ingress router before they are sent over the RSVP LSP.

Once both MTU signaling and packet fragmentation have been enabled on an ingress router, any route resolving to an RSVP LSP on this router uses the signaled MTU value. For information on how to configure this feature, see “Configuring MTU Signaling in RSVP” on page 293.

The following sections describe how MTU signaling in RSVP works:

- How the Correct MTU Is Signaled in RSVP on page 264
- Determining an Outgoing MTU Value on page 265
- MTU Signaling in RSVP Limitations on page 265

How the Correct MTU Is Signaled in RSVP

How the correct MTU is signaled in RSVP varies depending on whether the network devices (for example, routers) explicitly support MTU signaling in RSVP or not.

If the network devices support MTU signaling in RSVP, the following occur when you enable MTU signaling:

- The MTU is signaled from the ingress router to the egress router by means of the Adspec object. Before forwarding this object, the ingress router enters the MTU value associated with the interface over which the path message is sent. At each hop in the path, the MTU value in the Adspec object is updated to the minimum of the received value and the value of the outgoing interface.
- The ingress router uses the traffic specification (Tspec) object to specify the parameters for the traffic it is going to send. The MTU value signaled for the Tspec object at the ingress router is the maximum MTU value (9192 bytes). This value does not change en route to the egress router.
- When the Adspec object arrives at the egress router, the MTU value is correct for the path (meaning it is the smallest MTU value discovered). The egress router compares the MTU value in the Adspec object to the MTU value in the Tspec object. It signals the smaller MTU using the Flowspec object in the Resv message.
- When the Resv object arrives at the ingress router, the MTU value in this object is used as the MTU for the next hops that use the LSP.

In a network where there are devices that do not support MTU signaling in RSVP, you might have the following behaviors:

- If the egress router does not support MTU signaling in RSVP, the MTU is set to the value of the outgoing interface on the ingress router. Setting the MTU to the value of the outgoing interface is the same as the default behavior when MTU signaling is not configured.
- A Juniper Networks transit router that does not support MTU signaling in RSVP always propagates an MTU value of 1500 in the Adspec object.

Determining an Outgoing MTU Value

The outgoing MTU value is the smaller of the values received in the Adspec object compared to the MTU value of the outgoing interface. The MTU value of the outgoing interface is determined as follows:

- If you configure an MTU value under the [family mpls] hierarchy level, this value is signaled.
- If you do not configure an MTU, the `inet` MTU is signaled.

MTU Signaling in RSVP Limitations

The following are limitations to MTU signaling in RSVP:

- Changes in the MTU value might cause a temporary loss of traffic in the following situations:
 - For link protection and node protection, the MTU of the bypass is only signaled at the time the bypass becomes active. During the time it takes for the new path MTU to be propagated, packet loss might occur because of an MTU mismatch.
 - For fast reroute, the MTU of the path is updated only after the detour becomes active, causing a delay in an update to the MTU at the ingress router. Until the MTU is updated, packet loss might occur if there is an MTU mismatch.

In both cases, only packets that are larger than the detour or bypass MTU are lost.

- When an MTU is updated, it triggers a change in the next hop. Any change in the next hop causes the route statistics to be lost.
- The minimum MTU supported for MTU signaling in RSVP is 1488 bytes. This value prevents a false or incorrectly configured value from being used.
- For single-hop LSPs, the MTU value displayed by the `show` commands is the RSVP-signaled value. However, this MPLS value is ignored and the correct IP value is used.


Link Protection

Link protection helps to ensure that traffic going over a specific interface to a neighboring router can continue to reach this router if that interface fails. When link protection is configured for an interface and an LSP that traverses this interface, a bypass LSP is created that will handle this traffic if the interface fails. The bypass LSP uses a different interface and path to reach the same destination. The path used can be configured explicitly, or you can rely on CSPF. The RSVP metric for the bypass LSP is set in the range of 20,000 through 29,999 (this value is not user configurable).

If a link-protected interface fails, traffic is quickly switched to the bypass LSP. Note that a bypass LSP cannot share the same egress interface with the LSPs it monitors.

In Figure 21 on page 266, link protection is enabled on Interface B between Router 1 and Router 2. It is also enabled on LSP A, an LSP that traverses the link between Router 1 and Router 2. If the link between Router 1 and Router 2 fails, traffic from LSP A is quickly switched to the bypass LSP generated by link protection.

Figure 21: Link Protection Creating a Bypass LSP for the Protected Interface



Although LSPs traversing an interface can be configured to take advantage of link protection, it is important to note that it is specifically the interface that benefits from link protection. If link protection is enabled on an interface but not on a particular LSP traversing that interface, then if the interface fails, that LSP will also fail.



NOTE: Link protection does not work on unnumbered interfaces.

To protect traffic over the entire route taken by an LSP, you should configure fast reroute. For more information, see “Configuring Fast Reroute” on page 78.

The following sections provide more information on link protection:

- Fast Reroute, Node Protection, and Link Protection on page 266
- Multiple Bypass LSPs on page 267

Fast Reroute, Node Protection, and Link Protection

RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*, describes two different types of traffic protection for RSVP-signaled LSPs:

- One-to-one backup—In the JUNOS Software this type of traffic protection is provided by fast reroute. Each LSP requires a protecting LSP to be signaled at each hop except the egress router. This protecting LSP cannot be shared.
- Facility backup—This is sometimes called many-to-one backup. In the JUNOS Software this type of traffic protection is provided by node and link protection. Each LSP requires a protecting LSP to be signaled at each hop except the egress router. Unlike fast reroute, this protecting LSP can be shared by other LSPs.

The information above is summarized in Table 5 on page 267.

Table 5: One-to-One Backup Compared with Facility Backup

Comparison	One-to-One Backup	Facility Backup
Name of the protecting LSP	Detour LSP	Bypass LSP
Sharing of the protecting LSP	Cannot be shared	Can be shared by multiple LSPs
JUNOS configuration statements	<code>fast-reroute</code>	<code>node-link-protection</code> and <code>link-protection</code>

Multiple Bypass LSPs

By default, link protection relies on a single bypass LSP to provide path protection for an interface. However, you can also specify multiple bypass LSPs to provide link protection for an interface. You can individually configure each of these bypass LSPs or create a single configuration for all of the bypass LSPs. If you do not configure the bypass LSPs individually, they all share the same path and bandwidth constraints.

The following algorithm describes how and when an additional bypass LSP is activated for an LSP:

1. If any currently active bypass can satisfy the requirements of the LSP (bandwidth, link protection, or node-link protection), the traffic is directed to that bypass.
2. If no active bypass LSP is available, scan through the manual bypass LSPs in first-in, first-out (FIFO) order, skipping those that are already active (each manual bypass can only be activated once). The first inactive manual bypass that can satisfy the requirements is activated and traffic is directed to that bypass.
3. If no manual bypass LSPs are available and if the `max-bypasses` statement activates multiple bypass LSPs for link protection, determine whether an automatically configured bypass LSP can satisfy the requirements. If an automatically configured bypass LSP is available and if the total number of active automatically configured bypass LSPs does not exceed the maximum bypass LSP limit (configured with the `max-bypasses` statement), activate another bypass LSP.

For information on how to configure multiple bypass LSPs for link protection, see “Configuring Bypass LSPs” on page 283.

Node Protection


Node protection extends the capabilities of link protection. Link protection helps to ensure that traffic going over a specific interface to a neighboring router can continue to reach this router if that interface fails. Node protection ensures that traffic from an LSP traversing a neighboring router can continue to reach its destination even if the neighboring router fails.

When you enable node protection for an LSP, you must also enable link protection. Once enabled, node protection and link protection establish the following types of bypass LSPs:

- **Next-hop bypass LSP**—Provides an alternate route for an LSP to reach a neighboring router. This type of bypass LSP is established when you enable either node protection or link protection.
- **Next-next-hop bypass LSP**—Provides an alternate route for an LSP to get around a neighboring router en route to the destination router. This type of bypass LSP is established exclusively when node protection is configured.

In Figure 22 on page 268, both node protection and link protection are enabled on Interface B on Router 1. Both node protection and link protection are also enabled on LSP A, an LSP that traverses the link transiting Router 1, Router 2, and Router 3. If the link between Router 1 and Router 2 fails, traffic from LSP A is quickly switched to the next-hop bypass LSP generated by link protection. If Router 2 suffers a hardware or software failure, traffic from LSP A is switched to the next-next-hop bypass LSP generated by node protection.

Figure 22: Node Protection Creating a Next-Next-Hop Bypass LSP



The time needed by node protection to switch traffic to a next-next-hop bypass LSP can be significantly longer than the time needed by link protection to switch traffic to a next-hop bypass LSP. Link protection relies on a hardware mechanism to detect a link failure, allowing it to quickly switch traffic to a next-hop bypass LSP.

Node failures are often due to software problems on the node router. Node protection relies on the receipt of hello messages from a neighboring router to determine whether it is still functioning. The time it takes node protection to divert traffic partly depends on how often the node router sends hello messages and how long it takes the node-protected router to react to having not received a hello message. However, once the failure is detected, traffic can be quickly diverted to the next-next-hop bypass LSP.

RSVP Graceful Restart

RSVP graceful restart allows a router undergoing a restart to inform its adjacent neighbors of its condition. The restarting router requests a grace period from the neighbor or peer, which can then cooperate with the restarting router. The restarting router can still forward MPLS traffic during the restart period; convergence in the network is not disrupted. The restart is not visible to the rest of the network, and the restarting router is not removed from the network topology. RSVP graceful restart

can be enabled on both transit routers and ingress routers. It is available for both point-to-point LSPs and point-to-multipoint LSPs.

RSVP graceful restart is described in the following sections:

- RSVP Graceful Restart Standard on page 269
- RSVP Graceful Restart Terminology on page 269
- RSVP Graceful Restart Operation on page 270
- Processing the Restart Cap Object on page 271

RSVP Graceful Restart Standard

RSVP graceful restart is described in RFC 3473, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions* (only Section 9, “Fault Handling”).

RSVP Graceful Restart Terminology

R

Recovery time (in milliseconds)	<p>Applies only when the control channel is up (the hello exchange is complete) before the restart time. Applies only to nodal faults.</p> <p>When a graceful restart is in progress, the time left to complete a recovery is advertised. At other times, this value is zero. The maximum advertised recovery time is 2 minutes (120,000 milliseconds).</p> <p>During the recovery time, a restarting node attempts to recover its lost states with assistance from its neighbors. The neighbor of the restarting node must send the path messages with the recovery labels to the restarting node within a period of one-half the recovery time. The restarting node considers its graceful restart complete after its advertised recovery time.</p>
Restart time (in milliseconds)	<p>The default value is 60,000 milliseconds (1 minute). The restart time is advertised in the hello message. The time indicates how long a neighbor should wait to receive a hello message from a restarting router before declaring that router dead and purging states.</p> <p>The JUNOS Software can override a neighbor’s advertised restart time if the time is greater than one-third the local restart time. For example, given the default restart time of 60 seconds, a router would wait 20 seconds or less to receive a hello message from a restarting neighbor. If the restart time is zero, the restarting neighbor can immediately be declared dead.</p>

RSVP Graceful Restart Operation

For RSVP graceful restart to function, the feature must be enabled on the global routing instance. RSVP graceful restart can be disabled at the protocol level (for RSVP alone) or at the global level for all protocols.

RSVP graceful restart requires the following of a restarting router and the router's neighbors:

- For the restarting router, RSVP graceful restart attempts to maintain the routes installed by RSVP and the allocated labels, so that traffic continues to be forwarded without disruption. RSVP graceful restart is done quickly enough to reduce or eliminate the impact on neighboring nodes.
- The neighboring routers must have RSVP graceful restart helper mode enabled, thus allowing them to assist a router attempting to restart RSVP.

An object called Restart Cap that is sent in RSVP hello messages advertises a node's restart capability. The neighboring node sends a Recover Label object to the restarting node to recover its forwarding state. This object is essentially the old label that the restarting node advertised before the node went down.

The following lists the RSVP graceful restart behaviors, which vary depending on the configuration and on which features are enabled:

- If you disable helper mode, the JUNOS Software does not attempt to help a neighbor restart RSVP. Any information that arrives with a Restart Cap object from a neighbor is ignored.
- When you enable graceful restart under the routing instance configuration, the router can restart gracefully with the help of its neighbors. RSVP advertises a Restart Cap object (RSVP RESTART) in hello messages in which restart and recovery times are specified (neither value is 0).
- If you explicitly disable RSVP graceful restart under the `[protocols rsvp]` hierarchy level, the Restart Cap object is advertised with restart and recovery times specified as 0. The restart of neighboring routers is supported (unless helper mode is disabled), but the router itself does not preserve the RSVP forwarding state and cannot recover its control state.
- If after a restart RSVP realizes that no forwarding state has been preserved, the Restart Cap object is advertised with restart and recovery times specified as 0.
- If graceful restart and helper mode are disabled, RSVP graceful restart is completely disabled. The router neither recognizes nor advertises the RSVP graceful restart objects.

You cannot explicitly configure values for the restart and recovery times.

Unlike other protocols, there is no way for RSVP to determine that it has completed a restart procedure, other than a fixed timeout. All RSVP graceful restart procedures are timer-based. A `show rsvp version` command might indicate that the restart is still in progress even if all RSVP sessions are back up and the routes are restored.

Processing the Restart Cap Object

The following assumptions are made about a neighbor based on the Restart Cap object (assuming that a control channel failure can be distinguished unambiguously from a node restart):

- A neighbor that does not advertise the Restart Cap object in its hello messages cannot assist a router with state or label recovery, nor can it perform an RSVP graceful restart.
- After a restart, a neighbor advertising a Restart Cap object with a restart time equal to any value and a recovery time equal to 0 has not preserved its forwarding state. When a recovery time equals 0, the neighbor is considered dead and any states related to this neighbor are purged, regardless of the value of the restart time.
- After a restart, a neighbor advertising its recovery time with a value other than 0 can keep or has kept the forwarding state. If the local router is helping its neighbor with restart or recovery procedures, it sends a Recover Label object to this neighbor.

Chapter 12

RSVP Configuration Guidelines

This chapter describes how to configure RSVP and discusses the following configuration tasks:

- Minimum RSVP Configuration on page 273
- Configuring RSVP and MPLS on page 274
- Configuring RSVP Interfaces on page 275
- Configuring Node Protection or Link Protection for LSPs on page 281
- Configuring Inter-AS Node and Link Protection on page 281
- Configuring Link Protection on Interfaces Used by LSPs on page 282
- Configuring RSVP Graceful Restart on page 289
- Configuring Load Balancing Across RSVP LSPs on page 291
- Configuring Timers for RSVP Refresh Messages on page 292
- Preempting RSVP Sessions on page 293
- Configuring MTU Signaling in RSVP on page 293
- Configuring RSVP to Pop the Label on the Ultimate-Hop Router on page 295
- Disabling Adjacency Down and Neighbor Down Notification in IS-IS and OSPF on page 295
- Enabling Ultimate-Hop Popping on Point-to-Multipoint LSPs on page 296
- Tracing RSVP Protocol Traffic on page 297

Minimum RSVP Configuration

To enable RSVP on a single interface, include the **rsvp** statement and specify the interface using the **interface** statement. This is the minimum RSVP configuration. All other RSVP configuration statements are optional.

```
rsvp {  
    interface interface-name;  
}
```

You can include these statements at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

To enable RSVP on all interfaces, substitute **all** for the *interface-name* variable.

If you have configured interface properties on a group of interfaces and want to disable RSVP on one of the interfaces, include the **disable** statement:

```
interface interface-name {
    disable;
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp interface *interface-name*]
- [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name*]

Configuring RSVP and MPLS

The primary purpose of the JUNOS RSVP software is to support dynamic signaling within label-switched paths (LSPs). When you enable both MPLS and RSVP on a router, MPLS becomes a client of RSVP. No additional configuration is required to bind MPLS and RSVP.

You can configure MPLS to set up signaled paths by using the **label-switched-path** statement at the [edit protocols mpls] hierarchy level. Each LSP translates into a request for RSVP to initiate an RSVP session. This request is passed through the internal interface between label switching and RSVP. After examining the request information, checking RSVP states, and checking the local routing tables, RSVP initiates one session for each LSP. The session is sourced from the local router and is destined for the target of the LSP.

When an RSVP session is successfully created, the LSP is set up along the paths created by the RSVP session. If the RSVP session is unsuccessful, RSVP notifies MPLS of its status. It is up to MPLS to initiate backup paths or continue retrying the initial path.

To pass label-switching signaling information, RSVP supports four additional objects: Label Request object, Label object, Explicit Route object, and Record Route object. For an LSP to be set up successfully, all routers along the path must support MPLS, RSVP, and the four objects. Of the four objects, the Record Route object is not mandatory.

To configure MPLS and make it a client of RSVP, do the following:

- Enable MPLS on all routers that will participate in the label switching (this is, on all routers that might be part of a label-switching path).
- Enable RSVP on all routers and on all router interfaces that form the LSP.
- Configure the routers at the beginning of the LSP.

Example: Configuring RSVP and MPLS

The following shows a sample configuration for a router at the beginning of an LSP:


```
[edit]
protocols {
  mpls {
    label-switched-path sf-to-london {
      to 192.168.1.4;
    }
  }
  rsvp {
    interface so-0/0/0;
  }
}
```

The following shows a sample configuration for all the other routers that form the LSP:

```
[edit]
protocols {
  mpls {
    interface so-0/0/0;
  }
  rsvp {
    interface so-0/0/0;
  }
}
```

Configuring RSVP Interfaces

The following sections describe how to configure RSVP interfaces:

- Configuring RSVP Refresh Reduction on page 275
- Configuring the RSVP Hello Interval on page 278
- Configuring RSVP Authentication on page 278
- Configuring the Bandwidth Subscription for Class Types on page 279
- Configuring the RSVP Update Threshold on an Interface on page 279
- Configuring RSVP for Unnumbered Interfaces on page 280

Configuring RSVP Refresh Reduction

You can configure RSVP refresh reduction on each interface by including the following statements in the interface configuration:

- **aggregate**—Enable all RSVP refresh reduction features: RSVP message bundling, RSVP message ID, reliable message delivery, and summary refresh.
- **no-aggregate**—Disable RSVP message bundling and summary refresh.
- **reliable**—Enable RSVP message ID and reliable message delivery.
- **no-reliable**—Disable RSVP message ID, reliable message delivery, and summary refresh.

For more information on RSVP refresh reduction, see “RSVP Refresh Reduction” on page 262.

Table 6 on page 276 lists various combinations of the RSVP refresh reduction configuration statements and how they alter the behavior of the JUNOS Software. The table describes only the expected behavior based on the configuration on the router. The actual behavior is dictated not only by the local configuration on this router, but also on the refresh reduction capabilities of its RSVP neighbors. Note that by configuring the **aggregate** statement, you enable all RSVP refresh reduction features, including reliable message delivery.

Table 6: RSVP Refresh Reduction Behavior

Configuration Statement	Send Capability	Receive Capability
aggregate or aggregate and reliable	RR bit = 1 BundleMessage ID (Path/Resv messages) Ack/Nack (all messages) Summary Refresh	BundleAck/Nack (all messages) Summary Refresh
aggregate and no-reliable	RR bit = 1 BundleAck/Nack (all messages)	BundleMessage ID (all messages)
reliable or reliable and no-aggregate	RR bit = 0 Message ID (Path/Resv messages) Ack/Nack (all messages)	BundleMessage ID (all messages) Ack/Nack

The send capability shown in Table 6 on page 276 lists the RSVP messages and objects related to RSVP refresh reduction that the router is capable of sending. This does not mean that all these messages are exchanged between this router and a neighbor. For example, if the router is configured with the **aggregate** statement, but RSVP refresh reduction is not enabled on its neighbor, then no Summary Refresh message is sent to this neighbor even though the router is capable of sending it.

The receive capability shown in Table 6 on page 276 lists the messages and objects related to RSVP refresh reduction that the router is capable of receiving and processing without generating any errors or resulting in error conditions.

If the **no-reliable** statement is configured on the router (reliable message delivery is disabled), the router accepts RSVP messages that include the Message ID object but ignore the Message ID object and continue performing standard message processing. No error is generated in this case, and RSVP operates normally.

However, not all combinations between two neighbors with different refresh reduction capabilities function correctly. For example, a router is configured with either the **aggregate** statement and **no-reliable** statement or with the **reliable** and **no-aggregate** statements. If an RSVP neighbor sends a Summary Refresh object to this router, no error is generated, but the Summary Refresh object cannot be processed. Consequently, RSVP states can time out on this router if the neighbor is relying only on Summary Refresh to refresh those RSVP states.

We recommend, unless there are specific requirements, that you configure RSVP refresh reduction in a similar manner on each RSVP neighbor.

To enable all RSVP refresh reduction features on an interface, include the **aggregate** statement:

```
aggregate;
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp interface *interface-name*]
- [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name*]

To disable RSVP message bundling and summary refresh, include the **no-aggregate** statement:

```
no-aggregate;
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp interface *interface-name*]
- [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name*]

To enable RSVP message ID and reliable message delivery on an interface, include the **reliable** statement:

```
reliable;
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp interface *interface-name*]
- [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name*]

To disable RSVP message ID, reliable message delivery, and summary refresh, include the **no-reliable** statement:

```
no-reliable;
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp interface *interface-name*]
- [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name*]

Determining the Refresh Reduction Capability of RSVP Neighbors

To determine the RSVP refresh reduction capability of an RSVP neighbor, you need the following information:

- The RR bit advertised by the neighbor
- The local configuration of RSVP refresh reduction
- The actual RSVP messages received from the neighbor

To obtain this information, you can issue a **show rsvp neighbor detail** command. The following is a sample of output from this command:

```
user@host> show rsvp neighbor detail
RSVP neighbor: 6 learned
  Address: 192.168.224.178 via: fxp1.0 status: Up
```

```

Last changed time: 10:06, Idle: 5 sec, Up cnt: 1, Down cnt: 0
Message received: 36
Hello: sent 69, received: 69, interval: 9 sec
Remote instance: 0x60b8feba, Local instance: 0x74bc7a8d
Refresh reduction: not operational

```

```

Address: 192.168.224.186 via: fxp2.0 status: Down
Last changed time: 10:17, Idle: 40 sec, Up cnt: 0, Down cnt: 0
Message received: 6
Hello: sent 20, received: 0, interval: 9 sec
Remote instance: 0x0, Local instance: 0x2ae1b339
Refresh reduction: incomplete
Remote end: disabled, Ack-extension: enabled

```

```

Address: 192.168.224.188 via: fxp2.0 status: Up
Last changed time: 4:15, Idle: 0 sec, Up cnt: 1, Down cnt: 0
Message received: 55
Hello: sent 47, received: 31, interval: 9 sec
Remote instance: 0x6436a35b, Local instance: 0x663849f0
Refresh reduction: operational
Remote end: enabled, Ack-extension: enabled

```

For more information on the `show rsvp neighbor detail` command, see the *JUNOS Routing Protocols and Policies Command Reference*.

Configuring the RSVP Hello Interval

RSVP monitors the status of the interior gateway protocol (IGP) (IS-IS or OSPF) neighbors and relies on the IGP protocols to detect when a node fails. If an IGP protocol declares a neighbor down (because hello packets are no longer being received), RSVP also brings down that neighbor. However, the IGP protocols and RSVP still act independently when bringing a neighbor up.

For Juniper Networks routers, configuring a shorter or longer RSVP hello interval has no impact on whether or not an RSVP session is brought down. RSVP sessions are kept up even if RSVP hello packets are no longer being received. RSVP sessions are maintained until either the router stops receiving IGP hello packets or the RSVP Path and Resv messages time out.

However, the RSVP hello interval might impact when another vendor's equipment brings down an RSVP session. For example, a neighboring non-Juniper Networks router might be configured to monitor RSVP hello packets.

To modify how often RSVP sends hello packets, include the `hello-interval` statement:

```
hello-interval seconds;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section.

Configuring RSVP Authentication

All RSVP protocol exchanges can be authenticated to guarantee that only trusted neighbors participate in setting up reservations. By default, RSVP authentication is disabled.

RSVP authentication uses a Hashed Message Authentication Code (HMAC)-MD5 message-based digest. This scheme produces a message digest based on a secret authentication key and the message contents. (The message contents also include a sequence number.) The computed digest is transmitted with RSVP messages. Once you have configured authentication, all received and transmitted RSVP messages with all neighbors are authenticated on this interface.

MD5 authentication provides protection against forgery and message modification. It also can prevent replay attacks. However, it does not provide confidentiality, because all messages are sent in clear text.

By default, authentication is disabled. To enable authentication, configure a key on each interface by including the `authentication-key` statement:

```
authentication-key key;
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp interface *interface-name*]
- [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name*]

Configuring the Bandwidth Subscription for Class Types

By default, RSVP allows 100 percent the bandwidth for a class type to be used for RSVP reservations. When you oversubscribe a class type for a multiclass LSP, the aggregate demand of all RSVP sessions is allowed to exceed the actual capacity of the class type.

For detailed instructions on how to configure the bandwidth subscription for class types, see “Configuring the Bandwidth Subscription Percentage for LSPs” on page 124.

Configuring the RSVP Update Threshold on an Interface

The interior gateway protocols (IGPs) maintain the traffic engineering database, but the current available bandwidth on the traffic engineering database links originates from RSVP. When a link’s bandwidth changes, RSVP informs the IGPs, which can then update the traffic engineering database and forward the new bandwidth information to all network nodes. The network nodes then know how much bandwidth is available on the traffic engineering database link (local or remote), and CSPF can correctly compute the paths.

However, IGP updates can consume excessive system resources. Depending on the number of nodes in a network, it might not be desirable to perform an IGP update for small changes in bandwidth. By configuring the `update-threshold` statement at the [edit protocols rsvp] hierarchy level, you can adjust the threshold at which a change in the reserved bandwidth triggers an IGP update.

You can configure a value of from 1 percent through 20 percent (the default is 10 percent) for when to trigger an IGP update. If the change in the reserved bandwidth is greater than or equal to the configured threshold percentage of the static bandwidth on that interface, then an IGP update occurs. For example, if you have configured

the `update-threshold` statement to be 15 percent and the router discovers that the reserved bandwidth on a link has changed by 10 percent of the link bandwidth, RSVP does not trigger an IGP update. However, if the reserved bandwidth on a link changes by 20 percent of the link bandwidth, RSVP triggers an IGP update.

To adjust the threshold at which a change in the reserved bandwidth triggers an IGP update, include the `update-threshold` statement:

```
update-threshold percentage;
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp interface *interface-name*]
- [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name*]

Because of the update threshold, it is possible for Constrained Shortest Path First (CSPF) to compute a path using outdated traffic engineering database bandwidth information on a link. If RSVP attempts to establish an LSP over that path, it might find that there is insufficient bandwidth on that link. When this happens, RSVP triggers an IGP traffic engineering database update, flooding the updated bandwidth information on the network. CSPF can then recompute the path by using the updated bandwidth information, and attempt to find a different path, avoiding the congested link. Note that this functionality is the default and does not need any additional configuration.

You can configure the `rsvp-error-hold-time` statement at the [edit protocols mpls] hierarchy level or the [edit logical-systems *logical-system-name* protocols mpls] hierarchy level to improve the accuracy of the traffic engineering database (including the accuracy of bandwidth estimates for LSPs) using information provided by PathErr messages. See “Improving Traffic Engineering Database Accuracy with RSVP PathErr Messages” on page 60.

Configuring RSVP for Unnumbered Interfaces

The JUNOS Software supports RSVP traffic engineering over unnumbered interfaces. Traffic engineering information about unnumbered links is carried in the IGP traffic engineering extensions for OSPF and IS-IS as described in RFC 4203, *OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)*, and RFC 4205, *Intermediate System to Intermediate System (IS-IS) Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)*. Unnumbered links can also be specified in the MPLS traffic engineering signaling as described in RFC 3477, *Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)*. This feature allows you avoid having to configure IP addresses for each interface participating in the RSVP-signaled network.

To configure RSVP for unnumbered interfaces, you must configure the router with a router ID using the `router-id` statement specified at the [edit routing-options] hierarchy level. The router ID must be available for routing (you can typically use the loopback address). The RSVP control messages for the unnumbered links are sent using the router ID address (rather than a randomly selected address).

To configure link protection and fast reroute on a router with unnumbered interfaces enabled, you must configure at least two addresses. We recommend that you

configure a secondary interface on the loopback in addition to configuring the router ID.

Configuring Node Protection or Link Protection for LSPs

When you configure node protection or link protection on a router, bypass LSPs are created to the next-hop or next-next-hop routers for the LSPs traversing the router. You must configure node protection or link protection for each LSP that you want protected. To extend protection along the entire path used by an LSP, you must configure protection on each router that the LSP traverses.

To configure node protection on a router for a specified LSP, include the `node-link-protection` statement:

```
node-link-protection;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls label-switched-path *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*]

To configure link protection on a router for a specified LSP, include the `link-protection` (MPLS) statement:

```
link-protection;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls label-switched-path *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*]



NOTE: To complete the configuration of node or link protection, you must also configure link protection on all unidirectional RSVP interfaces that the LSPs traverse, as described in “Configuring Link Protection on Interfaces Used by LSPs” on page 282.

Configuring Inter-AS Node and Link Protection

To interoperate with other vendors’ equipment, the JUNOS Software supports the record route object (RRO) node ID subobject for use in inter-AS link and node protection configurations. The RRO node ID subobject is defined in RFC 4561, *Definition of a Record Route Object (RRO) Node-Id Sub-Object*. This functionality is enabled by default in JUNOS Release 9.4 and later.

If you have Juniper Networks routers running JUNOS Release 9.4 and later releases in the same MPLS-TE network as routers running JUNOS Release 8.4 and earlier releases, you might need to disable the RRO node ID subobject by configuring the `no-node-id-subobject` statement:

```
no-node-id-subobject;
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp]
- [edit logical-systems *logical-system-name* protocols rsvp]

Configuring Link Protection on Interfaces Used by LSPs

When you configure node protection or link protection on a router for LSPs as described in “Configuring Node Protection or Link Protection for LSPs” on page 281, you also must configure the `link-protection` statement on the RSVP interfaces used by the LSPs.

To configure link protection on the interfaces used by the LSPs, include the `link-protection (MPLS)` statement:

```
link-protection {
  disable;
  admin-group
    exclude group-names;
    include-all group-names;
    include-any group-names;
  }
  bandwidth bps;
  bypass bypass-name {
    bandwidth bps;
    hop-limit number;
    no-cspf;
    path address <strict | loose>;
    priority setup-priority reservation-priority;
    to address;
  }
  class-of-service cos-value;
  hop-limit number;
  max-bypasses number;
  no-cspf;
  no-node-protection;
  optimize-timer seconds;
  path address <strict | loose>;
  priority setup-priority reservation-priority;
  subscription percent {
    ct0 percent;
    ct1 percent;
    ct2 percent;
    ct3 percent;
  }
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp interface *interface-name*]
- [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name*]

All the statements under `link-protection` are optional.

The following sections describe how to configure link protection:

- Configuring Bypass LSPs on page 283
- Configuring Administrative Groups for Bypass LSPs on page 284
- Configuring the Bandwidth for Bypass LSPs on page 284
- Configuring Class of Service for Bypass LSPs on page 285
- Configuring the Hop Limit for Bypass LSPs on page 285
- Configuring the Maximum Number of Bypass LSPs on page 286
- Disabling CSPF for Bypass LSPs on page 287
- Disabling Node Protection for Bypass LSPs on page 287
- Configuring the Optimization Interval for Bypass LSPs on page 287
- Configuring an Explicit Path for Bypass LSPs on page 288
- Configuring the Amount of Bandwidth Subscribed for Bypass LSPs on page 288
- Configuring Priority and Preemption for Bypass LSPs on page 289

Configuring Bypass LSPs

You can configure specific bandwidth and path constraints for a bypass LSP. You can also individually configure each bypass LSP generated when you enable multiple bypass LSPs. If you do not configure the bypass LSPs individually, they all share the same path and bandwidth constraints (if any).

If you specify the **bandwidth**, **hop-limit**, and **path** statements for the bypass LSP, these values take precedence over the values configured at the `[edit protocols rsvp interface interface-name link-protection]` hierarchy level. The other attributes (**subscription**, **no-node-protection**, and **optimize-timer**) are inherited from the general constraints.

To configure a bypass LSP, include the **bypass** statement:

```
bypass bypass-name {
  bandwidth bps;
  class-of-service cos-value;
  hop-limit number;
  no-cspf;
  path address <strict | loose>;
  priority setup-priority reservation-priority;
  to address;
}
```

You can include this statement at the following hierarchy levels:

- `[edit protocols rsvp interface interface-name link-protection]`
- `[edit logical-systems logical-system-name protocols rsvp interface interface-name link-protection]`

Configuring the Next-Hop or Next-Next-Hop Node Address for Bypass LSPs

If you configure a bypass LSP, you must also configure the `to` statement. The `to` statement specifies the address for the interface of the immediate next-hop node (for link protection) or the next-next-hop node (for node-link protection). The address specified determines whether this is a link protection bypass or a node-link protection bypass. On multiaccess networks (for example, a LAN), this address is also used to specify which next-hop node is being protected.

Configuring Administrative Groups for Bypass LSPs

Administrative groups, also known as link coloring or resource class, are manually assigned attributes that describe the “color” of links, such that links with the same color conceptually belong to the same class. You can use administrative groups to implement a variety of policy-based LSP setups. You can configure administrative groups for bypass LSPs. For more information about configuring administrative groups, see “Configuring Administrative Groups” on page 96.

To configure administrative groups for bypass LSPs, include the `admin-group` statement:

```
admin-group {
  exclude group-names;
  include-all group-names;
  include-any group-names;
}
```

To configure an administrative group for all of the bypass LSPs, include the `admin-group` statement at the following hierarchy levels:

- [edit protocols rsvp interface *interface-name* link-protection]
- [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection]

To configure an administrative groups for a specific bypass LSP, include the `admin-group` statement at the following hierarchy levels:

- [edit protocols rsvp interface *interface-name* link-protection bypass *bypass-name*]
- [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection bypass *bypass-name*]

Configuring the Bandwidth for Bypass LSPs

You can specify the amount of bandwidth allocated for automatically generated bypass LSPs or you can individually specify the amount of bandwidth allocated for each LSP.

To specify the bandwidth allocation, include the `bandwidth` statement:

```
bandwidth bps;
```

For automatically generated bypass LSPs, include the **bandwidth** statement at the following hierarchy levels:

- [edit protocols rsvp interface *interface-name* link-protection]
- [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection]

For individually configured bypass LSPs, include the **bandwidth** statement at the following hierarchy levels:

- [edit protocols rsvp interface *interface-name* link-protection bypass *bypass-name*]
- [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection bypass *bypass-name*]

If you have enabled multiple bypass LSPs, this statement is required. See also “Configuring the Maximum Number of Bypass LSPs” on page 286.

Configuring Class of Service for Bypass LSPs

You can specify the class-of-service value for bypass LSPs by including the **class-of-service** statement:

```
class-of-service cos-value;
```

To apply a class-of-service value to all the automatically generated bypass LSPs, include the **class-of-service** statement at the following hierarchy levels:

- [edit protocols rsvp interface *interface-name* link-protection]
- [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection]

To configure a class-of-service value for a specific bypass LSPs, include the **class-of-service** statement at the following hierarchy levels:

- [edit protocols rsvp interface *interface-name* link-protection bypass *bypass-name*]
- [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection bypass *bypass-name*]

Configuring the Hop Limit for Bypass LSPs

You can specify the maximum number of hops a bypass can traverse. By default, each bypass can traverse a maximum of 255 hops (the ingress and egress routers count as one hop each, so the minimum hop limit is two).

To configure the hop limit for bypass LSPs, include the **hop-limit** statement:

```
hop-limit number;
```

For automatically generated bypass LSPs, include the **hop-limit** statement at the following hierarchy levels:

- [edit protocols rsvp interface *interface-name* link-protection]
- [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection]

For individually configured bypass LSPs, include the **hop-limit** statement at the following hierarchy levels:

- [edit protocols rsvp interface *interface-name* link-protection bypass *bypass-name*]
- [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection bypass *bypass-name*]

Configuring the Maximum Number of Bypass LSPs

You can specify the maximum number of dynamic bypass LSPs permitted for protecting an interface using the **max-bypasses** statement at the [edit protocols rsvp interface *interface-name* link-protection] hierarchy level. When this statement is configured, multiple bypasses for link protection are enabled. Call admission control (CAC) is also enabled.

By default, this option is disabled and only one bypass is enabled for each interface. You can configure a value of between 0 through 99 for the **max-bypasses** statement. Configuring a value of 0 prevents the creation of any dynamic bypass LSPs for the interface. If you configure a value of 0 for the **max-bypasses** statement, you need to configure one or more static bypass LSPs to enable link protection on the interface.

If you configure the **max-bypasses** statement, you must also configure the **bandwidth** statement (discussed in “Configuring the Bandwidth for Bypass LSPs” on page 284).

To configure the maximum number of bypass LSPs for a protected interface, include the **max-bypasses** statement:

```
max-bypasses number;
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp interface *interface-name* link-protection]
- [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection]

Disabling CSPF for Bypass LSPs

Under certain circumstances, you might need to disable CSPF computation for bypass LSPs and use the configured Explicit Route Object (ERO) if available. For example, a bypass LSP might need to traverse multiple OSPF areas or IS-IS levels, preventing the CSPF computation from working. To ensure that link and node protection function properly in this case, you have to disable CSPF computation for the bypass LSP.

You can disable CSPF computation for all bypass LSPs or for specific bypass LSPs.

To disable CSPF computation for bypass LSPs, include the **no-cspf** statement:

```
no-cspf;
```

For a list of hierarchy levels where you can include this statement, see the statement summary for this statement.

Disabling Node Protection for Bypass LSPs

You can disable node protection on the RSVP interface. Link protection remains active. When this option is configured, the router can only initiate a next-hop bypass, not a next-next-hop bypass.

To disable node protection for bypass LSPs, include the **no-node-protection** statement:

```
no-node-protection;
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp interface *interface-name* link-protection]
- [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection]

Configuring the Optimization Interval for Bypass LSPs

You can configure an optimization interval for bypass LSPs. At the end of this interval, an optimization process is initiated that attempts to either minimize the number of bypasses currently in use, minimize the total amount of bandwidth reserved for all of the bypasses, or both. You can configure an optimization interval from 1 through 65,535 seconds. A default value of 0 disables bypass LSP optimization.

To configure the optimization interval for bypass LSPs, include the **optimize-timer** statement:

```
optimize-timer seconds;
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp interface *interface-name* link-protection]
- [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection]

Configuring an Explicit Path for Bypass LSPs

By default, when you establish a bypass LSP to an adjacent neighbor, CSPF is used to discover the least-cost path. The **path** statement allows you to configure an explicit path (a sequence of strict or loose routes), giving you control over where and how the bypass LSP is established. To configure an explicit path, include the **path** statement:

```
path address <strict | loose>;
```

For automatically generated bypass LSPs, include the **path** statement at the following hierarchy levels:

- [edit protocols rsvp interface *interface-name* link-protection]
- [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection]

For individually configured bypass LSPs, include the **path** statement at the following hierarchy levels:

- [edit protocols rsvp interface *interface-name* link-protection bypass *bypass-name*]
- [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection bypass *bypass-name*]

Configuring the Amount of Bandwidth Subscribed for Bypass LSPs

You can configure the amount of bandwidth subscribed to bypass LSPs. You can configure the bandwidth subscription for the whole bypass LSP or for each class type that might traverse the bypass LSP. You can configure any value between 1 percent and 65,535 percent. By configuring a value less than 100 percent, you are undersubscribing the bypass LSPs. By configuring a value greater than 100 percent, you are oversubscribing the bypass LSPs.

The ability to oversubscribe the bandwidth for the bypass LSPs makes it possible to more efficiently use network resources. You can configure the bandwidth for the bypass LSPs based on the average network load as opposed to the peak load.

To configure the amount of bandwidth subscribed for bypass LSPs, include the **subscription** statement:

```
subscription percentage {
  ct0 percentage;
  ct1 percentage;
  ct2 percentage;
  ct3 percentage;
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp interface *interface-name* link-protection]

- [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection]

Configuring Priority and Preemption for Bypass LSPs

When there is insufficient bandwidth to establish a more important LSP, you might want to tear down a less important existing LSP to release the bandwidth. You do this by preempting the existing LSP.

For more detailed information on configuring setup priority and reservation priority for LSPs, see “Configuring Priority and Preemption for LSPs” on page 102.

To configure the bypass LSP’s priority and preemption properties, include the **priority** statement:

```
priority setup-priority reservation-priority;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Configuring RSVP Graceful Restart

The following RSVP graceful restart configurations are possible:

- Graceful restart and helper mode are both enabled (the default).
- Graceful restart is enabled but helper mode is disabled. A router configured in this way can restart gracefully, but cannot help a neighbor with its restart and recovery procedures.
- Graceful restart is disabled but helper mode is enabled. A router configured in this way cannot restart gracefully, but can help a restarting neighbor.
- Graceful restart and helper mode both are disabled. This configuration completely disables RSVP graceful restart (including restart and recovery procedures and helper mode). The router behaves like a router that does not support RSVP graceful restart.



NOTE: In order to turn on RSVP graceful restart, you must set the global graceful restart timer to at least 180 seconds.

The following sections describe how to configure RSVP graceful restart:

- Enabling Graceful Restart for All Routing Protocols on page 290
- Disabling Graceful Restart for RSVP on page 290
- Disabling RSVP Helper Mode on page 290
- Configuring the Maximum Helper Recovery Time on page 290
- Configuring the Maximum Helper Restart Time on page 290

Enabling Graceful Restart for All Routing Protocols

To enable graceful restart for RSVP, you need to enable graceful restart for all the protocols that support graceful restart on the router. For more information about graceful restart, see the *JUNOS Routing Protocols Configuration Guide*.

To enable graceful restart on the router, include the `graceful-restart` statement:

```
graceful-restart;
```

You can include this statement at the following hierarchy levels:

- `[edit routing-options]`
- `[edit logical-systems logical-system-name routing-options]`

Disabling Graceful Restart for RSVP

By default, RSVP graceful restart and RSVP helper mode are enabled when you enable graceful restart. However, you can disable one or both of these capabilities.

To disable RSVP graceful restart and recovery, include the `disable` statement at the `[edit protocols rsvp graceful-restart]` hierarchy level:

```
disable;
```

Disabling RSVP Helper Mode

To disable RSVP helper mode, include the `helper-disable` statement at the `[edit protocols rsvp graceful-restart]` hierarchy level:

```
helper-disable;
```

Configuring the Maximum Helper Recovery Time

To configure the amount of time the router retains the state of its RSVP neighbors while they undergo a graceful restart, include the `maximum-helper-recovery-time` statement at the `[edit protocols rsvp graceful-restart]` hierarchy level. This value is applied to all neighboring routers, so it should be based on the time required by the slowest RSVP neighbor to recover.

```
maximum-helper-recovery-time seconds;
```

Configuring the Maximum Helper Restart Time

To configure the delay between when the router discovers that a neighboring router has gone down and when it declares the neighbor down, include the `maximum-helper-restart-time` statement at the `[edit protocols rsvp graceful-restart]` hierarchy level. This value is applied to all neighboring routers, so it should be based on the time required by the slowest RSVP neighbor to restart.

maximum-helper-restart-time seconds;

Configuring Load Balancing Across RSVP LSPs

By default, when you have configured several RSVP LSPs to the same egress router, the LSP with the lowest metric is selected and carries all traffic. If all of the LSPs have the same metric, one of the LSPs is selected at random and all traffic is forwarded over it.

Alternatively, you can load balance traffic across all of the LSPs by enabling per-packet load balancing.

To enable per-packet load balancing on an ingress LSP, configure the **policy-statement** statement as follows:

```
[edit policy-options]
policy-statement policy-name {
  then {
    load-balance per-packet;
  }
  accept;
}
```

You then need to apply this statement as an export policy to the forwarding table. For more information on how to configure the **policy-statement** statement, see the *JUNOS Policy Framework Configuration Guide*.

Once per-packet load balancing is applied, traffic is distributed equally between the LSPs (by default).

You need to configure per-packet load balancing if you want to enable PFE fast reroute. To enable PFE fast reroute, include the policy statement for per-packet load balancing shown in this section in the configuration of each of the routers where a reroute might take place. See also “Configuring Fast Reroute” on page 78.

You can also load-balance the traffic between the LSPs in proportion to the amount of bandwidth configured for each LSP. This capability can better distribute traffic in networks with asymmetric bandwidth capabilities across external links, since the configured bandwidth of an LSP typically reflects the traffic capacity of that LSP.

To configure RSVP LSP load balancing, include the **load-balance** statement with the **bandwidth** option:

```
load-balance {
  bandwidth;
}
```

You can configure this statement at the following hierarchy levels:

- [edit protocols rsvp]
- [edit logical-systems *logical-system-name* protocols rsvp]

Keep the following information in mind when you use the **load-balance** statement:

- If you configure the **load-balance** statement, the behavior of currently running LSPs is not altered. To force currently running LSPs to use the new behavior, you can issue a **clear mpls lsp** command.
- The **load-balance** statement only applies to ingress LSPs that have per-packet load balancing enabled.
- For differentiated services aware traffic engineered LSPs, the bandwidth of an LSP is calculated by summing the bandwidth of all of the class types.

Configuring Timers for RSVP Refresh Messages

RSVP uses two related timing parameters:

- **refresh-time**—The refresh time controls the interval between the generation of successive refresh messages. The default value for the refresh time is 45 seconds. This number is derived from the **refresh-time** statement's default value of 30, multiplied by a fixed value of 1.5. This computation differs from RFC 2205, which states that the refresh time should be multiplied by a random value in the range from 0.5 through 1.5.

Refresh messages include path and Resv messages. Refresh messages are sent periodically so that reservation states in neighboring nodes do not time out. Each path and Resv message carries the refresh timer value, and the receiving node extracts this value from the messages.

- **keep-multiplier**—The keep multiplier is a small, locally configured integer from 1 through 255. The default value is 3. It indicates the number of messages that can be lost before a particular state is declared stale and must be deleted. The keep multiplier directly affects the lifetime of an RSVP state.

To determine the lifetime of a reservation state, use the following formula:

$$\text{lifetime} = (\text{keep-multiplier} + 0.5) \times (1.5 \times \text{refresh-time})$$

In the worst case, $(\text{keep-multiplier} - 1)$ successive refresh messages must be lost before a reservation state is deleted.

Juniper Networks does not recommend configuring a short RSVP hello timer. If quick discovery of a failed neighbor is needed, configure short IGP (OSPF or IS-IS) hello timers.

By default, the refresh timer value is 30 seconds. To modify this value, include the **refresh-time** statement:

```
refresh-time seconds;
```

You can include this statement at the following hierarchy levels:

- `[edit protocols rsvp]`
- `[edit logical-systems logical-system-name protocols rsvp]`

The default value of the keep multiplier is 3. To modify this value, include the `keep-multiplier` statement:

```
keep-multiplier number;
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp]
- [edit logical-systems *logical-system-name* protocols rsvp]

Preempting RSVP Sessions

Whenever bandwidth is insufficient to handle all RSVP sessions, you can control the preemption of RSVP sessions. By default, an RSVP session is preempted only by a new higher-priority session.

To always preempt a session when the bandwidth is insufficient, include the `preemption` statement with the `aggressive` option:

```
preemption aggressive;
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp]
- [edit logical-systems *logical-system-name* protocols rsvp]

To disable RSVP session preemption, include the `preemption` statement with the `disabled` option:

```
preemption disabled;
```

To return to the default (that is, preempt a session only for a new higher-priority session), include the `preemption` statement with the `normal` option:

```
preemption normal;
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp]
- [edit logical-systems *logical-system-name* protocols rsvp]

Configuring MTU Signaling in RSVP

To configure maximum transmission unit (MTU) signaling in RSVP, you need to configure MPLS to allow IP packets to be fragmented before they are encapsulated in MPLS. You also need to configure MTU signaling in RSVP. For troubleshooting purposes, you can configure MTU signaling alone without enabling packet fragmentation.

To configure MTU signaling in RSVP, include the `path-mtu` statement:

```

path-mtu {
    allow-fragmentation;
    rsvp {
        mtu-signaling;
    }
}

```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls]
- [edit logical-systems *logical-system-name* protocols mpls]

The following sections describe how to enable packet fragmentation and MTU signaling in RSVP:

- Enabling MTU Signaling in RSVP on page 294
- Enabling Packet Fragmentation on page 294

Enabling MTU Signaling in RSVP

To enable MTU signaling in RSVP, include the `rsvp mtu-signaling` statement:

```
rsvp mtu-signaling;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls path-mtu]
- [edit logical-systems *logical-system-name* protocols mpls path-mtu]

Once you have committed the configuration, changes in the MTU signaling behavior for RSVP take effect the next time the path is refreshed.

You can configure the `mtu-signaling` statement by itself at the [edit protocols mpls path-mtu rsvp] hierarchy level. This can be useful for troubleshooting. If you configure just the `mtu-signaling` statement, you can use the `show rsvp session detail` command to determine what the smallest MTU is on an LSP. The `show rsvp session detail` command displays the MTU value received and sent in the Adspec object.

Enabling Packet Fragmentation

To allow IP packets to be fragmented before they are encapsulated in MPLS, include the `allow-fragmentation` statement:

```
allow-fragmentation;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls path-mtu]
- [edit logical-systems *logical-system-name* protocols mpls path-mtu]



NOTE: Do not configure the `allow-fragmentation` statement alone. Always configure it in conjunction with the `mtu-signaling` statement.

Configuring RSVP to Pop the Label on the Ultimate-Hop Router

You can control the label value advertised on the egress router of an LSP. The default advertised label is label 3 (Implicit Null label). If label 3 is advertised, the penultimate-hop router removes the label and sends the packet to the egress router. When ultimate-hop popping is enabled, label 0 (IP version 4 [IPv4] Explicit Null label) is advertised. Ultimate-hop popping ensures that any packets traversing an MPLS network include a label.

To configure ultimate-hop popping for RSVP, include the `explicit-null` statement:

```
explicit-null;
```

You can include this statement at the following hierarchy levels:

- `[edit protocols mpls]`
- `[edit logical-systems logical-system-name protocols mpls]`



NOTE: Juniper Networks routers queue packets based on the incoming label. Routers from other vendors might queue packets differently. Keep this in mind when working with networks containing routers from multiple vendors.

For more information about labels, see “Label Description” on page 25 and “Label Allocation” on page 26.

Disabling Adjacency Down and Neighbor Down Notification in IS-IS and OSPF

Whenever IS-IS is deactivated, the IS-IS adjacencies are brought down. IS-IS signals to RSVP to bring down any RSVP neighbors associated with the IS-IS adjacencies, and this further causes the associated LSPs signaled by RSVP to go down as well.

A similar process occurs whenever OSPF is deactivated. The OSPF neighbors are brought down. OSPF signals to RSVP to bring down any of the RSVP neighbors associated with the OSPF neighbors, and this further causes the associated LSPs signaled by RSVP to go down as well.

If you need to migrate from IS-IS to OSPF or from OSPF to IS-IS, the IGP notification to RSVP for an adjacency or neighbor down event needs to be ignored. Using the `no-adjacency-down-notification` or `no-neighbor-down-notification` statements, you can disable IS-IS adjacency down notification or OSPF neighbor down notification, respectively, until the migration is complete. The network administrator is responsible for configuring the statements before the migration, and then removing them from the configuration afterward, so that IGP notification can function normally.

To disable adjacency down notification in IS-IS, include the `no-adjacency-down-notification` statement:

```
no-adjacency-down-notification;
```

You can include this statement at the following hierarchy levels:

- [edit protocols isis interface *interface-name*]
- [edit logical-systems *logical-system-name* protocols isis interface *interface-name*]

To disable neighbor down notification in OSPF, include the `no-neighbor-down-notification` statement:

```
no-neighbor-down-notification;
```

You can include this statement at the following hierarchy levels:

- [edit protocols ospf area *area-id* interface *interface-name*]
- [edit logical-systems *logical-system-name* protocols ospf area *area-id* interface *interface-name*]

Enabling Ultimate-Hop Popping on Point-to-Multipoint LSPs

By default, for both point-to-point and point-to-multipoint LSPs, penultimate-hop popping is used for MPLS traffic. MPLS labels are removed from packets on the router just before the egress router of the LSP. The plain IP packets are then forwarded to the egress router. For ultimate-hop popping, the egress router is responsible for both removing the MPLS label and processing the plain IP packet.

It can be beneficial to enable ultimate-hop popping on point-to-multipoint LSPs, particularly when transit traffic is traversing the same egress device. If you enable ultimate-hop popping, a single copy of traffic can be sent over the incoming link, saving significant bandwidth. By default, ultimate-hop popping is disabled. Ultimate-hop popping is not available for point-to-point LSPs.

You enable ultimate-hop popping for point-to-multipoint LSPs by configuring the `tunnel-services` statement. When you enable ultimate-hop popping, the JUNOS Software selects one of the available virtual loopback tunnel (VT) interfaces to loop back the packets to the PFE for IP forwarding. By default, the VT interface selection process is performed automatically. Bandwidth admission control is used to limit the number of LSPs that can be used on one VT interface. Once all the bandwidth is consumed on one interface, the JUNOS Software selects another VT interface with sufficient bandwidth for admission control.

If an LSP requires more bandwidth than is available from any of the VT interfaces, ultimate-hop popping cannot be enabled and penultimate-hop popping is enabled instead.

You can explicitly configure which VT interfaces handle the RSVP traffic by including the `devices` option for the `tunnel-services` statement. The `devices` option allows you to specify which VT interfaces are to be used by RSVP. If you do not configure this option, all of the VT interfaces available to the router can be used.

For ultimate-hop popping on point-to-multipoint LSPs to function properly, the egress router must have a PIC that provides tunnel services, such as the tunnel services PIC or the adaptive services PIC. Tunnel services are needed for popping the final MPLS label and for returning packets for IP address lookups.

If you configure the **tunnel-services** statement on an operating router, only the behavior of newly signaled LSPs changes. Existing LSPs are not affected. To force all existing LSPs to use ultimate-hop popping, issue a **clear mpls lsp** command. Note that this causes all of the MPLS LSPs on the router to be signaled again.

To enable ultimate-hop popping for the egress point-to-multipoint LSPs on a router, configure the **tunnel-services** statement:

```
tunnel-services {
    devices device-names;
}
```

You can configure this statement at the [edit protocols rsvp] hierarchy level.

To enable ultimate-hop popping for egress point-to-multipoint LSPs, you must also configure the **interface** statement with the **all** option:

```
interface all;
```

You must configure this statement at the [edit protocols rsvp] hierarchy level.

Tracing RSVP Protocol Traffic

To trace RSVP protocol traffic, include the **traceoptions** statement:

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp]
- [edit logical-systems *logical-system-name* protocols rsvp]

Use the **file** statement to specify the name of the file that receives the output of the tracing operation. All files are placed in the directory **/var/log**. We recommend that you place RSVP tracing output in the file **rsvp-log**.

You can specify the following RSVP-specific flags in the RSVP **traceoptions** statement:

- **all**—All tracing operations.
- **error**—All detected error conditions
- **event**—RSVP-related events (helps to trace events related to RSVP graceful restart)
- **lmp**—RSVP-Link Management Protocol (LMP) interactions
- **packets**—All RSVP packets

- path—All path messages
- pathtear—PathTear messages
- resv—Resv messages
- resvtear—ResvTear messages
- route—Routing information
- state—Session state transitions

For general information about tracing and global tracing options, see the *JUNOS Routing Protocols Configuration Guide*.

Examples: Tracing RSVP Protocol Traffic

Trace RSVP path messages in detail:

```
[edit]
protocols {
  rsvp {
    traceoptions {
      file rsvp size 10m files 5;
      flag path;
    }
  }
}
```

Trace all RSVP messages:

```
[edit]
protocols {
  rsvp {
    traceoptions {
      file rsvp size 10m files 5;
      flag packets;
    }
  }
}
```

Trace all RSVP error conditions:

```
[edit]
protocols {
  rsvp {
    traceoptions {
      file rsvp size 10m files 5;
      flag error;
    }
  }
}
```


Chapter 13

Summary of RSVP Configuration Statements

This chapter provides a reference for each RSVP configuration statement. The statements are organized alphabetically.

admin-group

Syntax admin-group {
 exclude [*group-names*];
 include-all [*group-names*];
 include-any [*group-names*];
 }

Hierarchy Level [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection],
 [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection bypass *bypass-name*],
 [edit protocols rsvp interface *interface-name* link-protection],
 [edit protocols rsvp interface *interface-name* link-protection bypass *bypass-name*]

Release Information Statement introduced in JUNOS Release 9.2.

Description Allows you to configure administrative groups for bypass label-switched paths (LSPs). You can configure administrative groups either globally for all bypass LSPs traversing an interface or for just a specific bypass LSP.

Options exclude *group-names*—Specifies the administrative groups to exclude for a bypass LSP.

 include-all *group-names*—Specifies the administrative groups whose links the bypass LSP must traverse.

 include-any *group-names*—Specifies the administrative groups whose links the bypass LSP can traverse.

Usage Guidelines See “Configuring Administrative Groups for Bypass LSPs” on page 284.

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

aggregate

Syntax	(aggregate no-aggregate);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols rsvp peer-interface <i>peer-interface-name</i>], [edit protocols rsvp interface <i>interface-name</i>], [edit protocols rsvp peer-interface <i>peer-interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>Control the use of RSVP aggregate messages on an interface or peer interface:</p> <ul style="list-style-type: none"> ■ aggregate—Use RSVP aggregate messages. ■ no-aggregate—Do not use RSVP aggregate messages. <p>Aggregate messages can pack multiple RSVP messages into a single transmission, thereby reducing network overhead and enhancing efficiency. The number of supportable sessions and processing overhead are significantly improved when aggregation is enabled.</p> <p>Not all routers connected to a subnet need to support aggregation simultaneously. Each RSVP router negotiates its intention to use aggregate messages on a per-neighbor basis. Only when both routers agree are aggregate messages sent.</p>
Default	Aggregation is disabled.
Usage Guidelines	See “Configuring RSVP Refresh Reduction” on page 275.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

authentication-key

Syntax	authentication-key <i>key</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols rsvp peer-interface <i>peer-interface-name</i>], [edit protocols rsvp interface <i>interface-name</i>], [edit protocols rsvp peer-interface <i>peer-interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>Authentication key (password). Neighboring routers use the password to verify the authenticity of packets sent from this interface or peer interface.</p> <p>RSVP uses HMAC-MD5 authentication, which is defined in RFC 2104, <i>HMAC: Keyed-Hashing for Message Authentication</i>.</p> <p>All routers that are connected to the same IP subnet must use the same authentication scheme and password.</p>
Options	<i>key</i> —Authentication password. It can be 1 through 16 contiguous digits or letters. Separate decimal digits with periods. Separate hexadecimal digits with periods and precede the string with 0x. If you include spaces in the password, enclose the entire password in quotation marks (" ").
Usage Guidelines	See “Configuring RSVP Authentication” on page 278.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

bandwidth

Syntax	<code>bandwidth <i>bps</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection], [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>], [edit protocols rsvp interface <i>interface-name</i>], [edit protocols rsvp interface <i>interface-name</i> link-protection], [edit protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>]</p>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>For certain logical interfaces (such as Asynchronous Transfer Mode [ATM], Permanent Virtual Circuit [PVC], or Frame Relay), you cannot determine the correct bandwidth from the hardware. This statement allows you to specify the actual available bandwidth.</p> <p>This statement also allows you to specify the bandwidth for a bypass label switched path (LSP). If you have configured multiple bypasses, this statement is mandatory and is applied to all of the bypass LSPs.</p>
Default	The hardware raw bandwidth is used.
Options	<p><i>bps</i>—Bandwidth in bits per second. You can specify this as an integer value. If you do so, count your zeros carefully, or you can use the abbreviations k (for a thousand), m (for a million), or g (for a billion [also called a thousand million]).</p> <p>Range: Any positive integer Default: 0 (no bandwidth is reserved)</p>
Usage Guidelines	See “Configuring the Bandwidth for Bypass LSPs” on page 284, “Configuring Link Protection on Interfaces Used by LSPs” on page 282, and “Configuring Bypass LSPs” on page 283.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

bypass

Syntax	<pre>bypass <i>bypass-name</i> { bandwidth <i>bps</i>; hop-limit <i>number</i>; no-cspf; path <i>address</i> <strict loose>; priority <i>setup-priority reservation-priority</i>; to <i>address</i>; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection], [edit protocols rsvp interface <i>interface-name</i> link-protection]</p>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>Allows you to configure specific bandwidth and path constraints for a bypass LSP. It is possible to individually configure multiple bypass LSPs. If you do not configure the bypass LSPs individually, they all share the same path and bandwidth constraints.</p> <p>If you specify the bandwidth, hop-limit, and path statements for the bypass LSP, these values take precedence over the values configured at the [edit protocols rsvp interface <i>interface-name</i> link-protection] hierarchy level. The other attributes (subscription, no-node-protection, and optimize-timer) are inherited from the general constraints.</p>
Options	<p>to <i>address</i>—(Required) Specify the address for the interface of the immediate next-hop node (for link protection) or the next-next-hop node (for node-link protection). The address specified determines whether this is a link protection bypass or a node-link protection bypass. On multiaccess networks (for example, a LAN), this address is also used to specify which next-hop node is being protected.</p> <p>The remaining statements are explained separately.</p>
Usage Guidelines	See “Configuring Bypass LSPs” on page 283.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

class-of-service

Syntax	<code>class-of-service cos-value;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>],</p> <p>[edit protocols rsvp interface <i>interface-name</i> link-protection],</p> <p>[edit protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>]</p>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>Class-of-service (CoS) value given to all packets in the bypass LSP. You can specify a single CoS value for all the bypass LSPs traversing an interface. You can also configure CoS values for specific bypass LSPs traversing an interface.</p> <p>The CoS value might affect the scheduling or queuing algorithm of traffic traveling along an LSP.</p>
Options	<p><i>cos-value</i>—CoS value. A higher value typically corresponds to a higher level of service.</p> <p>Range: 0 through 7</p> <p>Default: If you do not specify a CoS value, the IP precedence bits from the packet's IP header are used as the packet's CoS value.</p>
Usage Guidelines	See "Configuring Class of Service for Bypass LSPs" on page 285.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

disable

Syntax	disable;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit logical-systems <i>logical-system-name</i> protocols rsvp graceful-restart], [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection], [edit logical-systems <i>logical-system-name</i> protocols rsvp peer-interface <i>peer-interface-name</i>], [edit protocols rsvp], [edit protocols rsvp graceful-restart], [edit protocols rsvp interface <i>interface-name</i>], [edit protocols rsvp interface <i>interface-name</i> link-protection], [edit protocols rsvp peer-interface <i>peer-interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Explicitly disable RSVP or RSVP graceful restart. Explicitly disable link protection on the specified interface.
Default	RSVP is enabled on interfaces and peer interfaces configured with the RSVP interface statement. RSVP graceful restart is enabled on the router. Link protection is disabled.
Usage Guidelines	See “Minimum RSVP Configuration” on page 273, “Configuring RSVP Graceful Restart” on page 289, and “Configuring Link Protection on Interfaces Used by LSPs” on page 282.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

fast-reroute optimize-timer

Syntax	fast-reroute optimize-timer <i>seconds</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
Release Information	Statement added in JUNOS Release 7.5.
Description	Configure the optimize timer for fast reroute. The optimize timer triggers a periodic optimization process that recomputes the fast reroute detour LSPs to use network resources more efficiently.
Options	<i>seconds</i> —Specify the number of seconds between fast reroute detour LSP optimizations. Range: 0 through 65,535 seconds Default: 0 (disabled)
Usage Guidelines	See “Configuring the Optimization Interval for Fast Reroute Paths” on page 79.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

graceful-deletion-timeout

Syntax	graceful-deletion-timeout <i>seconds</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the time, in seconds, before completing graceful deletion of signaling.
Options	<i>seconds</i> —Time before completing graceful deletion of signaling. Range: 1 through 300 seconds Default: 30 seconds
Usage Guidelines	See “Configuring the Graceful Deletion Timeout Interval” on page 463.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

graceful-restart

Syntax	<pre> graceful-restart { disable; helper-disable; maximum-helper-recovery-time seconds; maximum-helper-restart-time seconds; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options], [edit protocols rsvp], [edit routing-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Enable graceful restart on the router. You must configure the graceful-restart statement at the [edit routing-options] hierarchy level to enable graceful restart on the router.
Options	<p>disable—Disable graceful restart on the router or for RSVP.</p> <p>helper-disable—Disable RSVP graceful restart helper mode (this option is only available at the [edit protocols rsvp] hierarchy level). Default: Helper mode is enabled by default.</p> <p>maximum-helper-recovery-time seconds—The maximum length of time the router stores the state of neighboring routers when they undergo a graceful restart. The value applies to all neighboring routers, so it should be based on the time that the slowest RSVP neighbor requires for restart. Default: 180 seconds Range: 1 through 3600 seconds</p> <p>maximum-helper-restart-time seconds—The maximum length of time the router waits between when it discovers that a neighboring router has gone down and when it declares the neighbor down. This value is applied to all neighboring routers, so it should be based on the time that the slowest RSVP neighbor requires for restart. Default: 20 seconds Range: 1 through 1800 seconds</p>
Usage Guidelines	See “Configuring RSVP Graceful Restart” on page 289.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

hello-interval

Syntax	hello-interval <i>seconds</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols rsvp peer-interface <i>peer-interface-name</i>], [edit protocols rsvp interface <i>interface-name</i>], [edit protocols rsvp peer-interface <i>peer-interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Enable the sending of hello packets on the interface.
Options	<i>seconds</i> —Length of time between hello packets. A value of 0 disables the sending of hello packets on the interface. Range: 1 through 60 seconds Default: 9 seconds
Usage Guidelines	See “Configuring the RSVP Hello Interval” on page 278.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

hop-limit

Syntax	hop-limit <i>number</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection], [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>], [edit protocols rsvp interface <i>interface-name</i> link-protection], [edit protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the maximum number of hops a bypass can traverse. By default, each bypass can traverse a maximum of 255 hops, including the ingress and egress routers.
Options	<i>number</i> —Maximum number of hops a bypass can traverse. Range: 2 through 255 hops Default: 255 hops
Usage Guidelines	See “Configuring the Hop Limit for Bypass LSPs” on page 285.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

interface

Syntax `interface interface-name {`
 `disable;`
 `(aggregate | no-aggregate);`
 `authentication-key key;`
 `bandwidth bps;`
 `hello-interval seconds;`
 `link-protection {`
 `disable;`
 `admin-group {`
 `exclude [group-names];`
 `include-all [group-names];`
 `include-any [group-names];`
 `}`
 `bandwidth bps;`
 `bypass bypass-name {`
 `bandwidth bps {`
 `ct0 bps;`
 `ct1 bps;`
 `ct2 bps;`
 `ct3 bps;`
 `}`
 `class-of-service cos-value;`
 `hop-limit number;`
 `no-cspf;`
 `path address <strict | loose>;`
 `priority setup-priority reservation-priority;`
 `to address;`
 `}`
 `class-of-service cos-value;`
 `hop-limit number;`
 `max-bypasses number;`
 `no-cspf;`
 `no-node-protection;`
 `optimize-timer seconds;`
 `path address <strict | loose>;`
 `priority setup-priority reservation-priority;`
 `subscription percentage;`
 `}`
 `(reliable | no-reliable);`
 `subscription percentage {`
 `ct0 percentage;`
 `ct1 percentage;`
 `ct2 percentage;`
 `ct3 percentage;`
 `}`
 `update-threshold threshold;`
 `}`

Hierarchy Level [edit logical-systems *logical-system-name* protocols rsvp],
 [edit protocols rsvp]

Release Information	Statement introduced before JUNOS Release 7.4.
Description	Enable RSVP on one or more router interfaces.
Default	RSVP is disabled on all interfaces.
Options	<p><i>interface-name</i>—Name of an interface. To configure all interfaces, specify all. For details about specifying interfaces, see the <i>JUNOS Network Interfaces Configuration Guide</i>.</p> <p>The remaining statements are explained separately.</p>
Usage Guidelines	See “Minimum RSVP Configuration” on page 273.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

keep-multiplier

Syntax	keep-multiplier <i>number</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Set the keep multiplier value.
Options	<p><i>number</i>—Multiplier value.</p> <p>Range: 1 through 255</p> <p>Default: 3</p>
Usage Guidelines	See “Configuring Timers for RSVP Refresh Messages” on page 292.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

link-protection (RSVP)

Syntax

```
link-protection {
  disable;
  admin-group {
    exclude [ group-names ];
    include-all [ group-names ];
    include-any [ group-names ];
  }
  bandwidth bps;
  bypass bypass-name {
    bandwidth bps {
      ct0 bps;
      ct1 bps;
      ct2 bps;
      ct3 bps;
    }
    class-of-service cos-value;
    hop-limit number;
    no-cspf;
    path address <strict | loose>;
    priority setup-priority reservation-priority;
    to address;
  }
  class-of-service cos-value;
  hop-limit number;
  max-bypasses number;
  no-cspf;
  no-node-protection;
  optimize-timer seconds;
  path address <strict | loose>;
  priority setup-priority reservation-priority;
  subscription percentage;
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name*],
[edit protocols rsvp interface *interface-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Enable link protection on the specified interface. Using link protection, you can configure a network to reroute traffic quickly around broken links. To fully enable link protection, you also need to configure the **link-protection** statement at the [edit protocols mpls label-switched-path *lsp-name*] hierarchy level. You can configure single or multiple bypasses for protected interface.

Default Link protection is disabled.

Options **no-node-protection**—Disables node-link protection on the RSVP interface. Link protection remains active. When this option is configured, the router can only initiate a next-hop bypass, not a next-next-hop bypass.

The remaining statements are explained separately.

Usage Guidelines	See “Configuring Link Protection on Interfaces Used by LSPs” on page 282.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Topics	link-protection (MPLS)

load-balance

Syntax	load-balance { bandwidth; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Load-balance traffic between RSVP LSPs.
Options	bandwidth—Load-balance traffic between RSVP LSPs based on the bandwidth configured for each LSP.
Usage Guidelines	See “Configuring Load Balancing Across RSVP LSPs” on page 291.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

max-bypasses

Syntax	max-bypasses <i>number</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i>], [edit protocols rsvp interface <i>interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4. Range modified in JUNOS Release 9.3.
Description	Specify the maximum number of dynamic bypass LSPs permitted for protecting this interface. When this option is configured, multiple bypasses for link protection are enabled. Call admission control (CAC) is also enabled. The limit on bypasses configured applies only to dynamically generated bypass LSPs. By default, this option is disabled and only one dynamic bypass LSP is enabled for each interface. If you configure max-bypasses , you must also configure the bandwidth statement.
Options	<i>number</i> —Configure the maximum number of bypass LSPs. If you configure a value of 0, no dynamic bypass LSPs are allowed to be established for the interface. Only static bypass LSPs can be configured. Default: 1 Range: 0 through 99
Usage Guidelines	See “Configuring the Maximum Number of Bypass LSPs” on page 286.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

no-adjacency-down-notification

Syntax	no-adjacency-down-notification;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i>], [edit protocols isis interface <i>interface-name</i>]
Release Information	Statement introduced in JUNOS Release 8.0.
Description	Disables adjacency down notification for IS-IS to allow for migration from IS-IS to OSPF without disruption of the RSVP neighbors and associated RSVP-signaled LSPs.
Usage Guidelines	See “Disabling Adjacency Down and Neighbor Down Notification in IS-IS and OSPF” on page 295.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

no-aggregate

See aggregate.

no-cspf

Syntax no-cspf;

Hierarchy Level [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection],
[edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection bypass *bypass-name*],
[edit protocols rsvp interface *interface-name* link-protection],
[edit protocols rsvp interface *interface-name* link-protection bypass *bypass-name*]

Release Information Statement introduced in JUNOS Release 7.5.

Description Disable CSPF computation on all bypass LSPs or on a specific bypass LSP. You need to disable CSPF for link protection to function properly on interarea paths.

Default CSPF is enabled.

Usage Guidelines See “Disabling CSPF for Bypass LSPs” on page 287.

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

no-neighbor-down-notification

Syntax no-neighbor-down-notification;

Hierarchy Level [edit logical-systems *logical-system-name* protocols ospf area *area-id* interface *interface-name*],
[edit protocols ospf area *area-id* interface *interface-name*]

Release Information Statement introduced in JUNOS Release 8.0.

Description Disables neighbor down notification for OSPF to allow for migration from OSPF to IS-IS without disruption of the RSVP neighbors and associated RSVP-signaled LSPs.

Usage Guidelines See “Disabling Adjacency Down and Neighbor Down Notification in IS-IS and OSPF” on page 295.

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

no-node-id-subobject

Syntax	no-node-id-subobject;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
Release Information	Statement introduced in JUNOS Release 9.4.
Description	Disables the record route object (RRO) node ID subobject for compatibility with earlier versions of the JUNOS Software. To interoperate with other vendors' equipment, the JUNOS Software supports the RRO node ID subobject for use in inter-AS link and node protection configurations.
Usage Guidelines	See “Configuring Inter-AS Node and Link Protection” on page 281.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

no-p2mp-sublsp

Syntax	no-p2mp-sublsp;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
Release Information	Statement introduced in JUNOS Release 9.2.
Description	Rejects Resv messages that include the S2L_SUB_LSP object. By default, Resv messages which include the S2L_SUB_LSP object are accepted. However, in a network which includes Juniper Networks devices running both JUNOS 9.2 and later releases and JUNOS 9.1 and earlier releases, it is necessary to configure the no-p2mp-sublsp statement on the JUNOS 9.2 and later devices to ensure that P2MP LSPs function properly.
Default	Resv messages which include the S2L_SUB_LSP object are accepted.
Usage Guidelines	See “Preserving Point-to-Multipoint LSP Functioning with Different JUNOS Software Releases” on page 147.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

no-reliable

See reliable

node-link-protection

Syntax	node-link-protection;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Enable node and link protection on the specified LSP. To fully enable node and link protection, you also need to include the <code>link-protection</code> statement at the [edit protocols rsvp interface <i>interface-name</i>] hierarchy level.
Default	Node and link protection is disabled.
Usage Guidelines	See “Configuring Node Protection or Link Protection for LSPs” on page 281.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

optimize-timer

Syntax	optimize-timer <i>seconds</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection], [edit protocols rsvp interface <i>interface-name</i> link-protection]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure an optimize timer for a bypass LSP. The optimize timer initiates a periodic optimization process that reshuffles data LSPs among bypass LSPs to achieve the most efficient use of network resources. The optimization process attempts to either minimize the number of bypasses currently in use, minimize the total amount of bandwidth reserved for all bypasses, or both.
Options	<i>seconds</i> —Specify the number of seconds between optimizations. Range: 0 through 65,535 seconds Default: 0 (disabled)
Usage Guidelines	See “Configuring the Optimization Interval for Bypass LSPs” on page 287.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

path

Syntax	<code>path address <strict loose>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection], [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>], [edit protocols rsvp interface <i>interface-name</i> link-protection], [edit protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure an explicit path (a sequence of strict or loose routes) to control where and how a bypass LSP is established. If multiple bypasses are configured, they all will use the same explicit path.
Default	No path is configured. CSPF automatically calculates the path the bypass LSP takes.
Options	<p>address—IP address of each transit router in the LSP. You must specify the address or hostname of each transit router, although you do not need to list each transit router if its type is loose. As an option, you can include the ingress and egress routers in the path. Specify the addresses in order, starting with the ingress router (optional) or the first transit router, and continuing sequentially along the path until reaching the egress router (optional) or the router immediately before the egress router.</p> <p>Default: If you do not specify any routers explicitly, no routing limitations are imposed on the bypass LSP.</p> <p>loose—The next address in the path statement is loose. The LSP can traverse other routers before reaching this router.</p> <p>Default: strict</p> <p>strict—The LSP must go to the next address specified in the path statement without traversing other nodes. This is the default.</p>
Usage Guidelines	See “Configuring an Explicit Path for Bypass LSPs” on page 288.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

peer-interface

Syntax `peer-interface peer-interface-name {
 disable;
 (aggregate | no-aggregate);
 authentication-key key;
 hello-interval seconds;
 (reliable | no-reliable);
}`

Hierarchy Level [edit logical-systems *logical-system-name* protocols rsvp],
 [edit protocols rsvp]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the name of the LMP peer device.

The remaining statements are explained separately.

Options The options are explained separately.

Usage Guidelines See “Configuring RSVP and OSPF for LMP Peer Interfaces” on page 457.

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

preemption

Syntax	preemption { (aggressive disabled normal); soft-preemption { cleanup-timer <i>seconds</i> ; } }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Control RSVP session preemption.
Default	normal
Options	<p>aggressive—Preempt RSVP sessions whenever bandwidth is insufficient to handle all sessions. A session is preempted whenever bandwidth is lowered or a new higher-priority session is established.</p> <p>disabled—Do not preempt RSVP sessions.</p> <p>normal—Preempt RSVP sessions to accommodate new higher-priority sessions when bandwidth is insufficient to handle all sessions.</p> <p>The remaining statements are explained separately.</p>
Usage Guidelines	See “Preempting RSVP Sessions” on page 293.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

priority

Syntax	<code>priority setup-priority reservation-priority;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>],</p> <p>[edit protocols rsvp interface <i>interface-name</i> link-protection],</p> <p>[edit protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>],</p>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>Configure the setup priority and reservation priority for a bypass LSP. If insufficient link bandwidth is available during session establishment, the setup priority is compared with other setup priorities for established sessions on the link to determine whether some of them should be preempted to accommodate the new session. The session with the lower-hold priority is preempted.</p>
Options	<p><i>reservation-priority</i>—Reservation priority, used to keep a reservation after it has been set up. A smaller number has a higher priority. The priority must be greater than or equal to the setup priority to prevent preemption loops.</p> <p>Range: 0 through 7, where 0 is the highest and 7 is the lowest priority.</p> <p>Default: 0 (Once the session is set up, no other session can preempt it.)</p> <p><i>setup-priority</i>—Setup priority.</p> <p>Range: 0 through 7, where 0 is the highest and 7 is the lowest priority.</p> <p>Default: 7 (The session cannot preempt any existing sessions.)</p>
Usage Guidelines	See “Configuring Priority and Preemption for Bypass LSPs” on page 289 and “Configuring Priority and Preemption for LSPs” on page 102.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

refresh-time

Syntax	<code>refresh-time seconds;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Set the refresh time.
Options	<i>seconds</i> —Refresh time. Range: 1 through 65,535 Default: 30 seconds
Usage Guidelines	See “Configuring Timers for RSVP Refresh Messages” on page 292.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

reliable

Syntax	<code>(reliable no-reliable);</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols rsvp peer-interface <i>peer-interface-name</i>], [edit protocols rsvp interface <i>interface-name</i>], [edit protocols rsvp peer-interface <i>peer-interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Enable reliable message delivery on the interface.
Usage Guidelines	See “Configuring RSVP Refresh Reduction” on page 275.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

rsvp

Syntax	<code>rsvp { ... }</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols], [edit protocols]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Enable RSVP routing on the router. You must include the rsvp statement in the configuration to enable RSVP on the router.
Default	RSVP is disabled on the router.
Usage Guidelines	See “Minimum RSVP Configuration” on page 273.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

soft-preemption

Syntax	<code>soft-preemption { cleanup-timer <i>seconds</i>; }</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp preemption], [edit protocols rsvp preemption]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Soft preemption attempts to establish a new path for a preempted LSP before tearing it down.
Options	cleanup-timer—A value of 0 disables soft preemption. Range: 0 through 180 seconds Default: 30 seconds
Usage Guidelines	See “Configuring MPLS Soft Preemption” on page 87.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

subscription

Syntax	<pre>subscription percentage { ct0 percentage; ct1 percentage; ct2 percentage; ct3 percentage; }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection], [edit protocols rsvp interface <i>interface-name</i>], [edit protocols rsvp interface <i>interface-name</i> link-protection]</pre>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configures the amount of bandwidth subscribed to a class type (when you have enabled Differentiated Services), or bypass LSP (when you have enabled link protection). The subscription is the percentage of the link bandwidth that can be used for the RSVP reservation process.
Options	<p><i>ctnumber percentage</i>—Percent of the class-type bandwidth that RSVP allows to be used for reservations. If you specify a value greater than 100, you are oversubscribing the class type. You can specify bandwidth subscriptions for class types 0 through 3. This option is not available for bypass LSPs. Range: 0 through 65,000 Default: 100 percent</p> <p><i>percentage</i>—Percent of the class-type or bypass LSP bandwidth that RSVP allows to be used for reservations. If you specify a value greater than 100, you are oversubscribing the class type or bypass LSP. Range: 0 through 65,000 Default: 100 percent</p>
Usage Guidelines	See “Configuring the Bandwidth Subscription Percentage for LSPs” on page 124 and “Configuring the Amount of Bandwidth Subscribed for Bypass LSPs” on page 288.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

traceoptions

Syntax traceoptions {
 file *filename* <files *number*> <size *size*> <world-readable | no-world-readable>;
 flag *flag* <flag-modifier> <disable>;
 }

Hierarchy Level [edit logical-systems *logical-system-name* protocols rsvp],
 [edit protocols rsvp]

Release Information Statement introduced before JUNOS Release 7.4.

Description Enable RSVP-level trace options.

Default The default RSVP-level trace options are those inherited from the routing protocols **traceoptions** statement included at the [edit routing-options] hierarchy level.

Options **disable**—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

filename—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`. We recommend that you place RSVP tracing output in the file **rsvp-log**.

files *number*—(Optional) Maximum number of trace files. When a trace file named *trace-file* reaches its maximum size, it is renamed *trace-file.0*, then *trace-file.1*, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Range: 2 through 1000

Default: 2 files

If you specify a maximum number of files, you must also include the **size** statement to specify the maximum file size.

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

- **all**—All tracing operations
- **error**—All detected error conditions
- **event**—RSVP-related events
- **lmp**—RSVP-LMP interactions
- **packets**—All RSVP packets
- **path**—All path messages
- **pathtear**—PathTear messages
- **resv**—Resv messages
- **resvtear**—ResvTear messages

- **route**—Routing information
- **state**—Session state transitions

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

no-world-readable—(Optional) Enable only certain users to read the log file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *trace-file* reaches this size, it is renamed *trace-file.0*. When the *trace-file* again reaches this size, *trace-file.0* is renamed *trace-file.1* and *trace-file* is renamed *trace-file.0*. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Syntax: *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 1 MB

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of files.

world-readable—(Optional) Enable any user to read the log file.

Usage Guidelines See “Tracing RSVP Protocol Traffic” on page 297.

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

tunnel-services

Syntax	tunnel-services { devices <i>device-names</i> ; }
Hierarchy Level	[edit protocols rsvp]
Release Information	Statement introduced in JUNOS Release 8.1.
Description	Enables ultimate-hop popping on point-to-multipoint LSPs. The JUNOS Software selects one of the available virtual tunnel (VT) interfaces to de-encapsulate the egress traffic. By default, the selection process is performed automatically.
Default	Ultimate-hop popping is disabled.
Options	devices <i>device-names</i> —Specifies which VT interfaces are used to handle the RSVP traffic. Range: 0 to 8 devices
Usage Guidelines	See “Enabling Ultimate-Hop Popping on Point-to-Multipoint LSPs” on page 296.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

update-threshold

Syntax	update-threshold <i>threshold</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i>], [edit protocols rsvp interface <i>interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Adjust the threshold at which a change in bandwidth triggers an interior gateway protocol (IGP) update.
Options	<i>threshold</i> —Specifies the percentage change in bandwidth to trigger an IGP update. Range: 1 through 20 percent Default: 10 percent
Usage Guidelines	See “Configuring the RSVP Update Threshold on an Interface” on page 279.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

Part 4

LDP

- LDP Overview on page 329
- LDP Configuration Guidelines on page 337
- Summary of LDP Configuration Statements on page 367

Chapter 14

LDP Overview

This chapter discusses the following topics:

- LDP Introduction on page 329
- LDP Standards on page 330
- JUNOS Software LDP Protocol Implementation on page 330
- LDP Operation on page 331
- Tunneling LDP LSPs in RSVP LSPs on page 331
- Tunneling LDP LSPs in RSVP LSPs Overview on page 331
- Label Operations on page 332
- LDP Message Types on page 333
- Discovery Messages on page 333
- Session Messages on page 334
- Advertisement Messages on page 334
- Notification Messages on page 334
- LDP Graceful Restart on page 334

LDP Introduction

The Label Distribution Protocol (LDP) is a protocol for distributing labels in non-traffic-engineered applications. LDP allows routers to establish label-switched paths (LSPs) through a network by mapping network-layer routing information directly to data link layer-switched paths.

These LSPs might have an endpoint at a directly attached neighbor (comparable to IP hop-by-hop forwarding), or at a network egress node, enabling switching through all intermediary nodes. LSPs established by LDP can also traverse traffic-engineered LSPs created by RSVP.

LDP associates a forwarding equivalence class (FEC) with each LSP it creates. The FEC associated with an LSP specifies which packets are mapped to that LSP. LSPs are extended through a network as each router chooses the label advertised by the next hop for the FEC and splices it to the label it advertises to all other routers. This process forms a tree of LSPs that converge on the egress router.

LDP Standards

LDP is described in the following RFC and Internet draft:

- RFC 3036, *LDP Specification*

The JUNOS Software supports the required elements of RFC 3036, except the following:

- loop detection
- CR-LDP

The RFC establishes three modes that the JUNOS Software only partially supports:

- Label distribution control mode
- Label retention mode
- Label advertisement mode

The following values for these modes are supported:

- Label distribution control mode: ordered
- Label retention mode: liberal
- Label advertisement mode: downstream unsolicited

The following values for these modes are not supported:

- Label distribution control mode: independent
- Label retention mode: conservative
- Label advertisement mode: downstream on demand
- RFC 3212, *Constraint-Based LSP Setup using LDP*
- RFC 3215, *LDP State Machine*
- RFC 3478, *Graceful Restart Mechanism for Label Distribution Protocol*

To access Internet RFCs and drafts, go to the IETF website at <http://www.ietf.org/>.

JUNOS Software LDP Protocol Implementation

The JUNOS Software implementation of LDP supports LDP version 1. The JUNOS Software supports a simple mechanism for tunneling between routers in an interior gateway protocol (IGP), to eliminate the required distribution of external routes within the core. The JUNOS Software allows an MPLS tunnel next hop to all egress routers in the network, with only an IGP running in the core to distribute routes to egress routers. Edge routers run BGP but do not distribute external routes to the core. Instead, the recursive route lookup at the edge resolves to an LSP switched to the egress router. No external routes are necessary on the transit LDP routers.

LDP Operation

You must configure LDP for each interface on which you want LDP to run. LDP creates LSP trees rooted at each egress router for the router ID address that is the subsequent BGP next hop. The ingress point is at every router running LDP. This process provides an `inet.3` route to every egress router. If BGP is running, it will attempt to resolve next hops by using the `inet.3` table first, which binds most, if not all, of the BGP routes to MPLS tunnel next hops.

Two adjacent routers running LDP become neighbors. If the two routers are connected by more than one interface, they become neighbors on each interface. When LDP routers become neighbors, they establish an LDP session to exchange label information. If per-router labels are in use on both routers, only one LDP session is established between them, even if they are neighbors on multiple interfaces. For this reason, an LDP session is not related to a particular interface.

LDP operates in conjunction with a unicast routing protocol. LDP installs LSPs only when both LDP and the routing protocol are enabled. For this reason, you must enable both LDP and the routing protocol on the same set of interfaces. If this is not done, LSPs might not be established between each egress router and all ingress routers, which might result in loss of BGP-routed traffic.

You can apply policy filters to labels received from and distributed to other routers through LDP. Policy filters provide you with a mechanism to control the establishment of LSPs.

For LDP to run on an interface, MPLS must be enabled on a logical interface on that interface. For more information, see the *JUNOS Network Interfaces Configuration Guide*.

Tunneling LDP LSPs in RSVP LSPs

You can tunnel LDP LSPs over RSVP LSPs. The following sections describe how tunneling of LDP LSPs in RSVP LSPs works:

- Tunneling LDP LSPs in RSVP LSPs Overview on page 331
- Label Operations on page 332

Tunneling LDP LSPs in RSVP LSPs Overview

If you are using RSVP for traffic engineering, you can run LDP simultaneously to eliminate the distribution of external routes in the core. The LSPs established by LDP are tunneled through the LSPs established by RSVP. LDP effectively treats the traffic-engineered LSPs as single hops.

When you configure the router to run LDP across RSVP-established LSPs, LDP automatically establishes sessions with the router at the other end of the LSP. LDP control packets are routed hop-by-hop, rather than carried through the LSP. This routing allows you to use simplex (one-way) traffic-engineered LSPs. Traffic in the

opposite direction flows through LDP-established LSPs that follow unicast routing rather than through traffic-engineered tunnels.

If you configure LDP over RSVP LSPs, you can still configure multiple OSPF areas and IS-IS levels in the traffic engineered core and in the surrounding LDP cloud.

Label Operations

Figure 23 on page 332 depicts an LDP LSP being tunneled through an RSVP LSP. (For definitions of label operations, see “Label Description” on page 25.) The shaded inner oval represents the RSVP domain, whereas the outer oval depicts the LDP domain. RSVP establishes an LSP through routers B, C, D, and E, with the sequence of labels L3, L4. LDP establishes an LSP through Routers A, B, E, F, and G, with the sequence of labels L1, L2, L5. LDP views the RSVP LSP between Routers B and E as a single hop.

When the packet arrives at Router A, it enters the LSP established by LDP, and a label (L1) is pushed onto the packet. When the packet arrives at Router B, the label (L1) is swapped with another label (L2). Because the packet is entering the traffic-engineered LSP established by RSVP, a second label (L3) is pushed onto the packet.

This outer label (L3) is swapped with a new label (L4) at the intermediate router (C) within the RSVP LSP tunnel, and when the penultimate router (D) is reached, the top label is popped. Router E swaps the label (L2) with a new label (L5), and the penultimate router for the LDP-established LSP (F) pops the last label.

Figure 23: Swap and Push When LDP LSPs Are Tunneled Through RSVP LSPs





Figure 24 on page 333 depicts double push label operation (L1L2), which is used when the ingress router (A) of the LDP and the RSVP LSPs are the same router. Note that Router D is the penultimate hop for the LDP-established LSP, so L2 is popped from the packet by Router D.

Figure 24: Double Push When LDP LSPs Are Tunneled Through RSVP LSPs

LDP Message Types

LDP uses the message types described in the following sections to establish and remove mappings and to report errors. All LDP messages have a common structure that uses a type, length, and value (TLV) encoding scheme.

- Discovery Messages on page 333
- Session Messages on page 334
- Advertisement Messages on page 334
- Notification Messages on page 334

Discovery Messages

Discovery messages announce and maintain the presence of a router in a network. Routers indicate their presence in a network by sending hello messages periodically. Hello messages are transmitted as UDP packets to the LDP port at the group multicast address for all routers on the subnet.

LDP uses the following discovery procedures:

- Basic discovery—A router periodically sends LDP link hello messages through an interface. LDP link hello messages are sent as UDP packets addressed to the LDP discovery port. Receipt of an LDP link hello message on an interface identifies an adjacency with the LDP peer router.
- Extended discovery—LDP sessions between routers not directly connected are supported by LDP extended discovery. A router periodically sends LDP targeted hello messages to a specific address. Targeted hello messages are sent as UDP packets addressed to the LDP discovery port at the specific address. The targeted router decides whether to respond to or ignore the targeted hello message. A targeted router that chooses to respond does so by periodically sending targeted hello messages to the initiating router.

Session Messages

Session messages establish, maintain, and terminate sessions between LDP peers. When a router establishes a session with another router learned through the hello message, it uses the LDP initialization procedure over TCP transport. When the initialization procedure is completed successfully, the two routers are LDP peers and can exchange advertisement messages.

Advertisement Messages

Advertisement messages create, change, and delete label mappings for FECs. Requesting a label or advertising a label mapping to a peer is a decision made by the local router. In general, the router requests a label mapping from a neighboring router when it needs one and advertises a label mapping to a neighboring router when it wants the neighbor to use a label.

Notification Messages

Notification messages provide advisory information and signal error information. LDP sends notification messages to report errors and other events of interest. There are two kinds of LDP notification messages:

- Error notifications, which signal fatal errors. If a router receives an error notification from a peer for an LDP session, it terminates the LDP session by closing the TCP transport connection for the session and discarding all label mappings learned through the session.
- Advisory notifications, which pass information to a router about the LDP session or the status of some previous message received from the peer.

LDP Graceful Restart

LDP graceful restart enables a router whose LDP control plane is undergoing a restart to continue to forward traffic while recovering its state from neighboring routers. It also enables a router on which helper mode is enabled to assist a neighboring router that is attempting to restart LDP.

During session initialization, a router advertises its ability to perform LDP graceful restart or to take advantage of a neighbor performing LDP graceful restart by sending the graceful restart TLV. This TLV contains two fields relevant to LDP graceful restart: the reconnect time and the recovery time. The values of the reconnect and recovery times indicate the graceful restart capabilities supported by the router.

The reconnect time is configured in the JUNOS Software as 60 seconds and is not user-configurable. When a router discovers that a neighboring router is restarting, it waits until the end of the recovery time before attempting to reconnect. The recovery time is the length of time a router waits for LDP to restart gracefully. The recovery time period begins when an initialization message is sent or received. This time period is also typically the length of time that a neighboring router maintains its information about the restarting router, allowing it to continue to forward traffic.

You can configure LDP graceful restart in both the master instance for the LDP protocol and for a specific routing instance. You can disable graceful restart at the global level for all protocols, at the protocol level for LDP only, and on a specific routing instance. LDP graceful restart is disabled by default, because at the global level, graceful restart is disabled by default. However, helper mode (the ability to assist a neighboring router attempting a graceful restart) is enabled by default.

The following are some of the behaviors associated with LDP graceful restart:

- Outgoing labels are not maintained in restarts. New outgoing labels are allocated.
- When a router is restarting, no label-map messages are sent to neighbors that support graceful restart until the restarting router has stabilized (label-map messages are immediately sent to neighbors that do not support graceful restart). However, all other messages (keepalive, address-message, notification, and release) are sent as usual. Distributing these other messages prevents the router from distributing incomplete information.
- Helper mode and graceful restart are independent. You can disable graceful restart in the configuration, but still allow the router to cooperate with a neighbor attempting to restart gracefully.

Chapter 15

LDP Configuration Guidelines

This chapter describes the minimum required LDP configuration and discusses the following configuration tasks:

- Minimum LDP Configuration on page 338
- Enabling and Disabling LDP on page 338
- Configuring the LDP Timer for Hello Messages on page 338
- Configuring the Delay Before LDP Neighbors Are Considered Down on page 339
- Enabling Strict Targeted Hello Messages for LDP on page 340
- Configuring the Interval for LDP Keepalive Messages on page 341
- Configuring the LDP Keepalive Timeout on page 341
- Configuring LDP Route Preferences on page 341
- Configuring LDP Graceful Restart on page 342
- Filtering Inbound LDP Label Bindings on page 343
- Filtering Outbound LDP Label Bindings on page 345
- Specifying the Transport Address Used by LDP on page 347
- Configuring the Prefixes Advertised into LDP from the Routing Table on page 348
- Configuring FEC Deaggregation on page 349
- Configuring Policers for LDP FECs on page 349
- Configuring LDP IPv4 FEC Filtering on page 350
- Configuring BFD for LDP LSPs on page 351
- Configuring ECMP-Aware BFD for RSVP LSPs on page 353
- Configuring a Failure Action for the BFD Session on an LDP LSP on page 353
- Configuring the Holddown Interval for the BFD Session on page 354
- Configuring OAM Ingress Policies for LDP on page 354
- Configuring LDP LSP Traceroute on page 354
- Collecting LDP Statistics on page 355
- Tracing LDP Protocol Traffic on page 358
- Configuring Miscellaneous LDP Properties on page 360

Minimum LDP Configuration

To enable LDP on a single interface, include the `ldp` statement and specify the interface using the `interface` statement. This is the minimum LDP configuration. All other LDP configuration statements are optional.

```
ldp {
  interface interface-name;
}
```

To enable LDP on all interfaces, specify `all` for *interface-name*.

For a list of hierarchy levels at which you can include these statements, see the statement summary sections.

Enabling and Disabling LDP

LDP is routing-instance-aware. To enable LDP on a specific interface, include the following statements:

```
ldp {
  interface interface-name;
}
```

For a list of hierarchy levels at which you can include these statements, see the statement summary sections.

To enable LDP on all interfaces, specify `all` for *interface-name*.

If you have configured interface properties on a group of interfaces and want to disable LDP on one of the interfaces, include the `interface` statement with the `disable` option:

```
interface interface-name {
  disable;
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section.

Configuring the LDP Timer for Hello Messages

LDP hello messages enable LDP nodes to discover one another and to detect the failure of a neighbor or the link to the neighbor. Hello messages are sent periodically on all interfaces where LDP is enabled.

There are two types of LDP hello messages:

- Link hello messages—Sent through the LDP interface as UDP packets addressed to the LDP discovery port. Receipt of an LDP link hello message on an interface identifies an adjacency with the LDP peer router.

- Targeted hello messages—Sent as UDP packets addressed to the LDP discovery port at a specific address. Targeted hello messages are used to support LDP sessions between routers that are not directly connected. A targeted router determines whether to respond or ignore a targeted hello message. A targeted router that chooses to respond does so by periodically sending targeted hello messages back to the initiating router.

By default, LDP sends hello messages every 5 seconds for link hello messages and every 15 seconds for targeted hello messages. You can configure the LDP timer to alter how often both types of hello messages are sent. However, you cannot configure a time for the LDP timer that is greater than the LDP hold time. For more information, see “Configuring the Delay Before LDP Neighbors are Considered Down” on page 339.

Configuring the LDP Timer for Link Hello Messages

To modify how often LDP sends link hello messages, specify a new link hello message interval for the LDP timer using the `hello-interval` statement:

```
hello-interval seconds;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Configuring the LDP Timer for Targeted Hello Messages

To modify how often LDP sends targeted hello messages, specify a new targeted hello message interval for the LDP timer by configuring the `hello-interval` statement as an option for the `targeted-hello` statement:

```
targeted-hello {
    hello-interval seconds;
}
```

For a list of hierarchy levels at which you can include these statements, see the statement summary sections for these statements.

Configuring the Delay Before LDP Neighbors Are Considered Down

The hold time determines how long an LDP node should wait for a hello message before declaring a neighbor to be down. This value is sent as part of a hello message so that each LDP node tells its neighbors how long to wait. The values sent by each neighbor do not have to match.

The hold time should normally be at least three times the hello interval. The default is 15 seconds for link hello messages and 45 seconds for targeted hello messages. However, it is possible to configure an LDP hold time that is close to the value for the hello interval.



NOTE: By configuring an LDP hold time close to the hello interval (less than three times the hello interval), LDP neighbor failures might be detected more quickly. However, this also increases the possibility that the router might declare an LDP neighbor down that is still functioning normally. For more information, see “Configuring the LDP Timer for Hello Messages” on page 338.

Configuring the LDP Hold Time for Link Hello Messages

To modify how long an LDP node should wait for a link hello message before declaring the neighbor down, specify a new time in seconds using the **hold-time** statement:

```
hold-time seconds;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Configuring the LDP Hold Time for Targeted Hello Messages

To modify how long an LDP node should wait for a targeted hello message before declaring the neighbor down, specify a new time in seconds using the **hold-time** statement as an option for the **targeted-hello** statement:

```
targeted-hello {
  hold-time seconds;
}
```

For a list of hierarchy levels at which you can include these statements, see the statement summary sections for these statements.

Enabling Strict Targeted Hello Messages for LDP

Use strict targeted hello messages to prevent LDP sessions from being established with remote neighbors that have not been specifically configured. If you configure the **strict-targeted-hellos** statement, an LDP peer does not respond to targeted hello messages coming from a source that is not one of its configured remote neighbors. Configured remote neighbors can include:

- Endpoints of RSVP tunnels for which LDP tunneling is configured
- Layer 2 circuit neighbors

If an unconfigured neighbor sends a hello message, the LDP peer ignores the message and logs an error (with the **error** trace flag) indicating the source. For example, if the LDP peer received a targeted hello from the Internet address 10.0.0.1 and no neighbor with this address is specifically configured, the following message is printed to the LDP log file:

```
LDP: Ignoring targeted hello from 10.0.0.1
```

To enable strict targeted hello messages, include the **strict-targeted-hellos** statement:

```
strict-targeted-hellos;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Configuring the Interval for LDP Keepalive Messages

The keepalive interval determines how often a message is sent over the session to ensure that the keepalive timeout is not exceeded. If no other LDP traffic is sent over the session in this much time, a keepalive message is sent. The default is 10 seconds.

To modify the keepalive interval, include the **keepalive-interval** statement:

```
keepalive-interval seconds;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Configuring the LDP Keepalive Timeout

After an LDP session is established, messages must be exchanged periodically to ensure that the session is still working. The keepalive timeout defines the amount of time that the neighbor LDP node waits before deciding that the session has failed. This value is usually set to at least three times the keepalive interval. The default is 30 seconds.

To modify the keepalive interval, include the **keepalive-timeout** statement:

```
keepalive-timeout seconds;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

The value configured for the **keepalive-timeout** statement is displayed as the hold time when you issue the **show ldp session detail** command.

Configuring LDP Route Preferences

When several protocols calculate routes to the same destination, route preferences are used to select which route is installed in the forwarding table. The route with the lowest preference value is selected. The preference value can be a number in the range 0 through 255. By default, LDP routes have a preference value of 9.

To modify the route preferences, include the **preference** statement:

```
preference preference;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Configuring LDP Graceful Restart

By default, graceful restart helper mode is enabled, but graceful restart is disabled. Thus, the default behavior of a router is to assist neighboring routers attempting a graceful restart, but not to attempt a graceful restart itself.

To configure LDP graceful restart, see the following sections:

- Enabling Graceful Restart on page 342
- Disabling LDP Graceful Restart or Helper Mode on page 342
- Configuring Recovery Time and Maximum Recovery Time on page 343

Enabling Graceful Restart

To enable LDP graceful restart, you also need to enable graceful restart on the router. To enable graceful restart, include the `graceful-restart` statement:

```
graceful-restart;
```

You can include this statement at the following hierarchy levels:

- [edit routing-options]
- [edit logical-systems *logical-system-name* routing-options]

The `graceful-restart` statement enables graceful restart for all protocols supporting this feature on the router. For more information about graceful restart, see the *JUNOS Routing Protocols Configuration Guide*.

By default, LDP graceful restart is enabled when you enable graceful restart at both the LDP protocol level and on all the routing instances. However, you can disable both LDP graceful restart and LDP graceful restart helper mode.

Disabling LDP Graceful Restart or Helper Mode

To disable LDP graceful restart and recovery, include the `disable` statement:

```
ldp {
  graceful-restart {
    disable;
  }
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

You can disable helper mode at the LDP protocols level only. You cannot disable helper mode for a specific routing instance. To disable LDP helper mode, include the `helper-disable` statement:

```
ldp {
```

```

    graceful-restart {
        helper-disable;
    }
}

```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

The following LDP graceful restart configurations are possible:

- LDP graceful restart and helper mode are both enabled.
- LDP graceful restart is disabled but helper mode is enabled. A router configured in this way cannot restart gracefully but can help a restarting neighbor.
- LDP graceful restart and helper mode are both disabled. The router does not use LDP graceful restart or the graceful restart type, length, and value (TLV) sent in the initialization message. The router behaves as a router that cannot support LDP graceful restart.

A configuration error is issued if you attempt to enable graceful restart and disable helper mode.

Configuring Recovery Time and Maximum Recovery Time

The recovery time is the amount of time a router waits for LDP to restart gracefully. The recovery time period begins when an initialization message is sent or received. This period is also typically the amount of time that a neighboring router maintains its information about the restarting router, allowing it to continue to forward traffic.

To prevent a neighboring router from being adversely affected if it receives a false value for the recovery time from the restarting router, you can configure the maximum recovery time on the neighboring router. A neighboring router maintains its state for the shorter of the two times. For example, router A is performing an LDP graceful restart. It has sent a recovery time of 900 seconds to neighboring router B. However, router B has its maximum recovery time configured at 400 seconds. Router B will only wait for 400 seconds before it purges its LDP information from router A.

To configure recovery time, include the **recovery-time** statement and the **maximum-neighbor-recovery-time** statement:

```

    graceful-restart {
        maximum-neighbor-recovery-time seconds;
        recovery-time seconds;
    }
}

```

For a list of hierarchy levels at which you can configure these statements, see the statement summary sections for these statements.

Filtering Inbound LDP Label Bindings

You can filter received LDP label bindings, applying policies to accept or deny bindings advertised by neighboring routers. To configure received-label filtering, include the **import** statement:

```
import [ policy-names ];
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

The named policy (configured at the [edit **policy-options**] hierarchy level) is applied to all label bindings received from all LDP neighbors. All filtering is done with **from** statements. Table 7 on page 344 lists the only **from** operators that apply to LDP received-label filtering.

Table 7: from Operators That Apply to LDP Received-Label Filtering

from Operator	Description
interface	Matches on bindings received from a neighbor that is adjacent over the specified interface
neighbor	Matches on bindings received from the specified LDP router ID
next-hop	Matches on bindings received from a neighbor advertising the specified interface address
route-filter	Matches on bindings with the specified prefix

If a binding is filtered, it still appears in the LDP database, but is not considered for installation as part of a label-switched path (LSP).

Generally, applying policies in LDP can be used only to block the establishment of LSPs, not to control their routing. This is because the path that an LSP follows is determined by unicast routing, and not by LDP. However, when there are multiple equal-cost paths to the destination through different neighbors, you can use LDP filtering to exclude some of the possible next hops from consideration. (Otherwise, LDP chooses one of the possible next hops at random.)

LDP sessions are not bound to interfaces or interface addresses. LDP advertises only per-router (not per-interface) labels; so if multiple parallel links exist between two routers, only one LDP session is established, and it is not bound to a single interface. When a router has multiple adjacencies to the same neighbor, take care to ensure that the filter does what is expected. (Generally, using **next-hop** and **interface** is not appropriate in this case.)

If a label has been filtered (meaning that it has been rejected by the policy and is not used to construct an LSP), it is marked as filtered in the database:

```
user@host> show ldp database
Input label database, 10.10.255.1:0-10.10.255.6:0
Label Prefix
3 10.10.255.6/32 (Filtered)
Output label database, 10.10.255.1:0-10.10.255.6:0
Label Prefix
3 10.10.255.1/32 (Filtered)
```

For more information about how to configure policies for LDP, see the *JUNOS Policy Framework Configuration Guide*.

Examples: Filtering Inbound LDP Label Bindings

Accept only /32 prefixes from all neighbors:

```
[edit]
protocols {
  ldp {
    import only-32;
    ...
  }
}
policy-options {
  policy-statement only-32 {
    term first {
      from {
        route-filter 0.0.0.0/0 upto /31;
      }
      then reject;
    }
    then accept;
  }
}
```

Accept 131.108/16 or longer from router ID 10.10.255.2 and accept all prefixes from all other neighbors:

```
[edit]
protocols {
  ldp {
    import nosy-neighbor;
    ...
  }
}
policy-options {
  policy-statement nosy-neighbor {
    term first {
      from {
        neighbor 10.10.255.2;
        route-filter 131.108.0.0/16 orlonger accept;
        route-filter 0.0.0.0/0 orlonger reject;
      }
    }
    then accept;
  }
}
```

Filtering Outbound LDP Label Bindings

You can configure export policies to filter LDP outbound labels. You can filter outbound label bindings by applying routing policies to block bindings from being

advertised to neighboring routers. To configure outbound label filtering, include the **export** statement:

```
export [policy-name];
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

The named export policy (configured at the [edit policy-options] hierarchy level) is applied to all label bindings transmitted to all LDP neighbors. The only **from** operator that applies to LDP outbound label filtering is **route-filter**, which matches bindings with the specified prefix. The only **to** operators that apply to outbound label filtering are the operators in Table 8 on page 346.

Table 8: to Operators for LDP Outbound-Label Filtering

to Operator	Description
interface	Matches on bindings sent to a neighbor that is adjacent over the specified interface
neighbor	Matches on bindings sent to the specified LDP router ID
next-hop	Matches on bindings sent to a neighbor advertising the specified interface address

If a binding is filtered, the binding is not advertised to the neighboring router, but it can be installed as part of an LSP on the local router. You can apply policies in LDP to block the establishment of LSPs, but not to control their routing. The path an LSP follows is determined by unicast routing, not by LDP.

LDP sessions are not bound to interfaces or interface addresses. LDP advertises only per-router (not per-interface) labels. If multiple parallel links exist between two routers, only one LDP session is established, and it is not bound to a single interface.

Do not use the **next-hop** and **interface** operators when a router has multiple adjacencies to the same neighbor.

Filtered labels are marked in the database:

```
user@host> show ldp database
Input label database, 10.10.255.1:0-10.10.255.3:0
Label Prefix
100007 10.10.255.2/32
3 10.10.255.3/32
Output label database, 10.10.255.1:0-10.10.255.3:0
Label Prefix
3 10.10.255.1/32
100001 10.10.255.6/32 (Filtered)
```

For more information about how to configure policies for LDP, see the *JUNOS Policy Framework Configuration Guide*.

Examples: Filtering Outbound LDP Label Bindings

Block transmission of the route for 10.10.255.6/32 to any neighbors:

```
[edit protocols]
ldp {
    export block-one;
}
policy-options {
    policy-statement block-one {
        term first {
            from {
                route-filter 10.10.255.6/32 exact;
            }
            then reject;
        }
        then accept;
    }
}
```

Send only 131.108/16 or longer to router ID 10.10.255.2, and send all prefixes to all other routers:

```
[edit protocols]
ldp {
    export limit-lsps;
}
policy-options {
    policy-statement limit-lsps {
        term allow-one {
            from {
                route-filter 131.108.0.0/16 orlonger;
            }
            to {
                neighbor 10.10.255.2;
            }
            then accept;
        }
        term block-the-rest {
            to {
                neighbor 10.10.255.2;
            }
            then reject;
        }
        then accept;
    }
}
```

Specifying the Transport Address Used by LDP

You can control the transport address used by LDP. The transport address is the address used for the TCP session over which LDP is running. To configure transport address control, include the `transport-address` statement:

```
transport-address (router-id | interface);
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

If you specify the **router-id** option, the address of the router identifier is used as the transport address (unless otherwise configured, the router identifier is typically the same as the loopback address). If you specify the **interface** option, the interface address is used as the transport address for any LDP sessions to neighbors that can be reached over that interface. Note that the router identifier is used as the transport address by default.

You cannot specify the **interface** option when there are multiple parallel links to the same LDP neighbor, because the LDP specification requires that the same transport address be advertised on all interfaces to the same neighbor. If LDP detects multiple parallel links to the same neighbor, it disables interfaces to that neighbor one by one until the condition is cleared, either by disconnecting the neighbor on an interface or by specifying the **router-id** option.

Configuring the Prefixes Advertised into LDP from the Routing Table

You can control the set of prefixes that are advertised into LDP and cause the router to be the egress router for those prefixes. By default, only the loopback address is advertised into LDP. To configure the set of prefixes from the routing table to be advertised into LDP, include the **egress-policy** statement:

```
egress-policy policy-name;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.



NOTE: If you configure an egress policy for LDP that does not include the loopback address, it is no longer advertised in LDP. To continue to advertise the loopback address, you need to explicitly configure it as a part of the LDP egress policy.

The named policy (configured at the [edit policy-options] or [edit logical-systems *logical-system-name* policy-options] hierarchy level) is applied to all routes in the routing table. Those routes that match the policy are advertised into LDP. You can control the set of neighbors to which those prefixes are advertised by using the **export** statement. Only **from** operators are considered; you can use any valid **from** operator. For more information, see the *JUNOS Routing Protocols Configuration Guide*.

Example: Configuring the Prefixes Advertised into LDP

Advertise all connected routes into LDP:

```
[edit protocols]
ldp {
  egress-policy connected-only;
}
policy-options {
```

```

policy-statement connected-only {
  from {
    protocol direct;
  }
  then accept;
}

```

Configuring FEC Deaggregation

When an LDP egress router advertises multiple prefixes, the prefixes are bound to a single label and aggregated into a single forwarding equivalence class (FEC). By default, LDP maintains this aggregation as the advertisement traverses the network.

By default, because an LSP cannot be split across multiple next hops and all the prefixes are bound into a single LSP, you cannot load-balance across equal-cost paths.

To change the default to load-balance across equal-cost paths, you need to deaggregate the FECs. Deaggregating the FECs causes each prefix to be bound to a separate label and become a separate LSP.

To configure deaggregated FECs, include the **deaggregate** statement:

```
deaggregate;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

For all LDP sessions, you can configure deaggregated FECs only globally.

Deaggregating a FEC allows the resulting multiple LSPs to be distributed across multiple equal-cost paths and distributes LSPs across the multiple next hops on the egress segments but installs only one next hop per LSP.

To aggregate FECs, include the **no-deaggregate** statement:

```
no-deaggregate;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

For all LDP sessions, you can configure aggregated FECs only globally.

Configuring Policers for LDP FECs

You can configure the JUNOS Software to track and police traffic for LDP FECs. LDP FEC policers can be used to do any of the following:

- Track or police the ingress traffic for an LDP FEC.
- Track or police the transit traffic for an LDP FEC.
- Track or police LDP FEC traffic originating from a specific forwarding class.

- Track or police LDP FEC traffic originating from a specific virtual routing and forwarding (VRF) site.
- Discard false traffic bound for a specific LDP FEC.

To police traffic for an LDP FEC, you must first configure a filter. Specifically, you need to configure either the `interface` statement or the `interface-set` statement at the `[edit firewall family protocol-family filter filter-name term term-name from]` hierarchy level. The `interface` statement allows you to match the filter to a single interface. The `interface-set` statement allows you to match the filter to multiple interfaces.

For more information on how to configure the `interface` statement, the `interface-set` statement, and policers for LDP FECs, see the *JUNOS Policy Framework Configuration Guide*.

Once you have configured the filters, you need to include them in the `policing` statement configuration for LDP. To configure policers for LDP FECs, include the `policing` statement:

```
policing {
  fec fec-address {
    ingress-traffic filter-name;
    transit-traffic filter-name;
  }
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

The `policing` statement includes the following options:

- `fec`—Specify the FEC address for the LDP FEC you want to police.
- `ingress-filter`—Specify the name of the ingress traffic filter.
- `transit-traffic`—Specify the name of the transit traffic filter.

Configuring LDP IPv4 FEC Filtering

By default, when a targeted LDP session is established, the JUNOS Software always exchanges both the IPv4 forwarding equivalence classes (FECs) and the Layer 2 circuit FECs over the targeted LDP session. For an LDP session to an indirectly connected neighbor, you might only want to export Layer 2 circuit FECs to the neighbor if the session was specifically configured to support Layer 2 circuits or VPLS.

In a mixed vendor network where all non-BGP prefixes are advertised into LDP, the LDP database can become large. For this type of environment, it can be useful to prevent the advertisement of IPv4 FECs over LDP sessions formed due to Layer 2 circuit or LDP VPLS configuration. Similarly, it can be useful to filter any IPv4 FECs received in this sort of environment.

If all the LDP neighbors associated with an LDP session are Layer 2 only, you can configure the JUNOS Software to only advertise Layer 2 circuit FECs by configuring the `l2-smart-policy` statement. This feature also automatically filters out the IPv4 FECs

received on this session. If you have configured an explicit export or import policy, this feature is disabled.

If one of the LDP session's neighbors is formed due to a discovered adjacency or if the adjacency is formed due to an LDP tunneling configuration on one or more RSVP LSPs, the IPv4 FECs are advertised and received using the default behavior.

To prevent LDP from exporting IPv4 FECs over LDP sessions with Layer 2 neighbors only and to filter out IPv4 FECs received over such sessions, include the `I2-smart-policy` statement:

```
I2-smart-policy;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary for this statement.

Configuring BFD for LDP LSPs

You can configure Bidirectional Forwarding Detection (BFD) for LDP LSPs. BFD can be turned on for all or a subset of LDP IPv4 FECs.

An error is logged whenever a BFD session for a path fails. The following shows how BFD for LDP LSP log messages might appear:

```
RPD_LDP_BFD_UP: LDP BFD session for FEC 10.255.16.14/32 is up
RPD_LDP_BFD_DOWN: LDP BFD session for FEC 10.255.16.14/32 is down
```

You can also configure BFD for RSVP LSPs, as described in “Configuring BFD for RSVP-Signaled LSPs” on page 167.

To enable BFD for LDP LSPs, include the `oam` and `bfd-liveness-detection` statements. You can enable BFD for all LDP LSPs on the router or just for the LDP LSPs associated with a specific FEC.

```
oam {
  bfd-liveness-detection {
    detection-time threshold milliseconds;
    ecmp;
    failure-action {
      remove-nexthop;
      remove-route;
    }
    holddown-interval seconds;
    ingress-policy ingress-policy-name;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    minimum-transmit-interval milliseconds;
    multiplier detection-time-multiplier;
    no-adaptation;
    transmit-interval {
      minimum-interval milliseconds;
      threshold milliseconds;
    }
  }
}
```

```

fec fec-address;
lsp-ping-interval seconds;
periodic-traceroute {
    disable;
    exp exp-value;
    fanout fanout-value;
    frequency minutes;
    paths number-of-paths;
    retries retry-attempts;
    source address;
    ttl ttl-value;
    wait seconds;
}
}

```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

The `oam` statement includes the following options:

- `fec`—Specify the FEC address.
- `lsp-ping-interval`—Specify the duration of the LSP ping interval in seconds. To issue a ping on an LDP-signaled LSP, use the `ping mpls ldp` command. For more information, see the *JUNOS System Basics and Services Command Reference*.

The `bfd-liveness-detection` statement includes the following options:

- `ecmp`—Causes LDP to establish BFD sessions for all ECMP paths configured for the specified FEC. If you configure the `ecmp` option, you must also configure the `periodic-traceroute` statement for the specified FEC. If you do not do so, the commit fails. You can configure the `periodic-traceroute` statement at the global hierarchy level (`[edit protocols ldp oam]`) while only configuring the `ecmp` option for a specific FEC (`[edit protocols ldp oam fec address bfd-liveness-detection]`).
- `holddown-interval`—Specifies the duration the BFD session should remain up before adding the route or next hop. Specifying a time of 0 seconds causes the route or next hop to be added as soon as the BFD session comes back up.
- `minimum-interval`—Specifies the minimum transmit and receive interval. If you configure the `minimum-interval` option, you do not need to configure the `minimum-receive-interval` option or the `minimum-transmit-interval` option.
- `minimum-receive-interval`—Specifies the minimum receive interval. The range is from 1 through 255,000 milliseconds.
- `minimum-transmit-interval`—Specifies the minimum transmit interval. The range is from 1 through 255,000 milliseconds.
- `multiplier`—Specifies the detection time multiplier. The range is from 1 through 255.

Configuring ECMP-Aware BFD for RSVP LSPs

When you configure BFD for a FEC, a BFD session is established for only one active local next-hop for the router. However, you can configure multiple BFD sessions, one for each FEC associated with a specific equal-cost multipath (ECMP) path. For this to function properly, you also need to configure LDP LSP periodic traceroute. (See “Configuring LDP LSP Traceroute” on page 354.) RSVP LSP traceroute is used to discover ECMP paths. A BFD session is initiated for each ECMP path discovered. Whenever a BFD session for one of the ECMP paths fails, an error is logged.

RSVP LSP traceroute is run periodically to check the integrity of the ECMP paths. The following might occur when a problem is discovered:

- If the latest RSVP LSP traceroute for a FEC differs from the previous traceroute, the BFD sessions associated with that FEC (the BFD sessions for address ranges that have changed from previous run) are brought down and new BFD sessions are initiated for the destination addresses in the altered ranges.
- If the RSVP LSP traceroute returns an error (for example, a timeout), all the BFD sessions associated with that FEC are torn down.

To configure RSVP to establish BFD sessions for all ECMP paths configured for the specified FEC, include the `ecmp` statement.

```
ecmp;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Along with the `ecmp` statement, you must also include the `periodic-traceroute` statement, either in the global LDP OAM configuration (at the `[edit protocols ldp oam]` or `[edit logical-systems logical-system-name protocols ldp oam]` hierarchy level) or in the configuration for the specified FEC (at the `[edit protocols ldp oam fec address]` or `[edit logical-systems logical-system-name protocols ldp oam fec address]` hierarchy level). Otherwise, the commit operation fails.

Configuring a Failure Action for the BFD Session on an LDP LSP

You can configure route and next-hop properties in the event of a BFD session failure event on an LDP LSP. The failure event could be an existing BFD session that has gone down or could be a BFD session that never came up. LDP adds back the route or next hop when the relevant BFD session comes back up.

You can configure one of the following failure action options for the `failure-action` statement in the event of a BFD session failure on the LDP LSP:

- **remove-nexthop**—Removes the route corresponding to the next hop of the LSP's route at the ingress node when a BFD session failure event is detected.
- **remove-route**—Removes the route corresponding to the LSP from the appropriate routing tables when a BFD session failure event is detected. If the LSP is

configured with ECMP and a BFD session corresponding to any path goes down, the route is removed.

To configure a failure action in the event of a BFD session failure on an LDP LSP, include either the **remove-nexthop** option or the **remove-route** option for the **failure-action** statement:

```
failure-action {
  remove-nexthop;
  remove-route;
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Configuring the Holddown Interval for the BFD Session

You can specify the duration the BFD session should be up before adding a route or next hop by configuring the **holddown-interval** statement at either the **[edit protocols ldp oam bfd-liveness-detection]** hierarchy level or at the **[edit protocols ldp oam fec address bfd-liveness-detection]** hierarchy level. Specifying a time of 0 seconds causes the route or next hop to be added as soon as the BFD session comes back up.

```
holddown-interval seconds;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Configuring OAM Ingress Policies for LDP

Using the **ingress-policy** statement, you can configure an Operation, Administration, and Management (OAM) policy to choose which forwarding equivalence classes (FECs) need to have OAM enabled. If the FEC passes through the policy or if the FEC is explicitly configured, OAM is enabled for a FEC. For FECs chosen using a policy, the BFD parameters configured under **[edit protocols ldp oam bfd-liveness-detection]** are applied.

You configure the OAM ingress policy at the **[edit policy-options]** hierarchy level. To configure an OAM ingress policy, include the **ingress-policy** statement:

```
ingress-policy ingress-policy-name;
```

You can configure this statement at the following hierarchy levels:

- **[edit protocols ldp oam]**
- **[edit logical-systems *logical-system-name* protocols ldp oam]**

Configuring LDP LSP Traceroute

You can trace the route followed by an LDP-signaled LSP. LDP LSP traceroute is based on RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*. This

feature allows you to periodically trace all paths in a FEC. The FEC topology information is stored in a database accessible from the CLI.

A topology change does not automatically trigger a trace of an LDP LSP. However, you can manually initiate a traceroute. If the traceroute request is for an FEC that is currently in the database, the contents of the database are updated with the results.

The periodic traceroute feature applies to all FECs specified by the `oam` statement configured at the `[edit protocols ldp]` hierarchy level. To configure periodic LDP LSP traceroute, include the `periodic-traceroute` statement:

```
periodic-traceroute {
  disable;
  exp exp-value;
  fanout fanout-value;
  frequency minutes;
  paths number-of-paths;
  retries retry-attempts;
  source address;
  ttl ttl-value;
  wait seconds;
}
```

You can configure this statement at the following hierarchy levels:

- `[edit protocols ldp oam]`
- `[edit protocols ldp oam fec address]`

The `periodic-traceroute` statement includes the following options:

- `exp`—Specify the class of service to use when sending probes.
- `fanout`—Specify the maximum number of next hops to search per node.
- `frequency`—Specify the interval between traceroute attempts.
- `paths`—Specify the maximum number of paths to search.
- `retries`—Specify the number of attempts to send a probe to a specific node before giving up.
- `source`—Specify the IPv4 source address to use when sending probes.
- `ttl`—Specify the maximum time-to-live value. Nodes that are beyond this value are not traced.
- `wait`—Specify the wait interval before resending a probe packet.

Collecting LDP Statistics

LDP traffic statistics show the volume of traffic that has passed through a particular FEC on a router.

When you configure the `traffic-statistics` statement at the `[edit protocols ldp]` hierarchy level, the LDP traffic statistics are gathered periodically and written to a file. You can configure how often statistics are collected (in seconds) by using the `interval` option.

The default collection interval is 5 minutes. You must configure an LDP statistics file; otherwise LDP traffic statistics are not gathered. If the LSP goes down, the LDP statistics are reset.

To collect LDP traffic statistics, include the **traffic-statistics** statement:

```
traffic-statistics {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  interval interval;
  no-penultimate-hop;
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

This section includes the following topics:

- LDP Statistics Output on page 356
- Disabling LDP Statistics on the Penultimate-Hop Router on page 357
- LDP Statistics Limitations on page 357

LDP Statistics Output

The following is sample output from an LDP statistics file:

FEC	Type	Packets	Bytes	Shared
10.255.350.448/32	Transit	0	0	No
	Ingress	0	0	No
10.255.350.450/32	Transit	0	0	Yes
	Ingress	0	0	No
10.255.350.451/32	Transit	0	0	No
	Ingress	0	0	No
220.220.220.1/32	Transit	0	0	Yes
	Ingress	0	0	No
220.220.220.2/32	Transit	0	0	Yes
	Ingress	0	0	No
220.220.220.3/32	Transit	0	0	Yes
	Ingress	0	0	No

May 28 15:02:05, read 12 statistics in 00:00:00 seconds

The following describes each column of data collected in the LDP statistics file:

- **Bytes**—Number of bytes of data passed by the FEC since its LSP came up.
- **FEC**—FEC for which LDP traffic statistics are collected.
- **Packets**—Number of packets passed by the FEC since its LSP came up.
- **read**—This number (which appears next to the date and time) might differ from the actual number of the statistics displayed. Some of the statistics are summarized before being displayed.
- **Shared**—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.

- **Type**—Type of traffic originating from a router, either **Ingress** (originating from this router) or **Transit** (forwarded through this router).

Disabling LDP Statistics on the Penultimate-Hop Router

Gathering LDP traffic statistics at the penultimate-hop router can consume excessive system resources, on next-hop routes in particular. This problem is exacerbated if you have configured the **deaggregate** statement in addition to the **traffic-statistics** statement. For routers reaching their limit of next-hop route usage, we recommend configuring the **no-penultimate-hop** option for the **traffic-statistics** statement:

```
traffic-statistics {  
  no-penultimate-hop;  
}
```

For a list of hierarchy levels at which you can configure the **traffic-statistics** statement, see the statement summary section for this statement.



NOTE: When you configure the **no-penultimate-hop** option, no statistics are available for the FECs that are the penultimate hop for this router.

Whenever you include or remove this option from the configuration, the LDP sessions are taken down and then restarted.

The following is sample output from an LDP statistics file showing routers on which the **no-penultimate-hop** option is configured:

FEC	Type	Packets	Bytes	Shared
10.255.245.218/32	Transit	0	0	No
	Ingress	4	246	No
10.255.245.221/32	Transit	statistics disabled		
	Ingress	statistics disabled		
13.1.1.0/24	Transit	statistics disabled		
	Ingress	statistics disabled		
13.1.3.0/24	Transit	statistics disabled		
	Ingress	statistics disabled		

LDP Statistics Limitations

The following are issues related to collecting LDP statistics by configuring the **traffic-statistics** statement:

- You cannot clear the LDP statistics.
- If you shorten the specified interval, a new LDP statistics request is issued only if the statistics timer expires later than the new interval.

- A new LDP statistics collection operation cannot start until the previous one has finished. If the interval is short or if the number of LDP statistics is large, the time gap between the two statistics collections might be longer than the interval.

When an LSP goes down, the LDP statistics are reset.

Tracing LDP Protocol Traffic

The following sections describe how to configure the trace options to examine LDP protocol traffic:

- Tracing LDP Protocol Traffic at the Protocol and Routing Instance Levels on page 358
- Tracing LDP Protocol Traffic Within FECs on page 359
- Examples: Tracing LDP Protocol Traffic on page 359

Tracing LDP Protocol Traffic at the Protocol and Routing Instance Levels

To trace LDP protocol traffic, you can specify options in the global `traceoptions` statement at the `[edit routing-options]` hierarchy level, and you can specify LDP-specific options by including the `traceoptions` statement:

```
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  flag flag <flag-modifier> <disable>;
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Use the `file` statement to specify the name of the file that receives the output of the tracing operation. All files are placed in the directory `/var/log`. We recommend that you place LDP-tracing output in the file `ldp-log`.

The following trace flags display the operations associated with the sending and receiving of various LDP messages. Each can carry one or more of the following modifiers:

- **address**—Trace the operation of address and address withdrawal messages.
- **binding**—Trace label-binding operations.
- **error**—Trace error conditions.
- **event**—Trace protocol events.
- **initialization**—Trace the operation of initialization messages.
- **label**—Trace the operation of label request, label map, label withdrawal, and label release messages.
- **notification**—Trace the operation of notification messages.
- **packet**—Trace the operation of address, address withdrawal, initialization, label request, label map, label withdrawal, label release, notification, and periodic

messages. This modifier is equivalent to setting the **address**, **initialization**, **label**, **notification**, and **periodic** modifiers.

- **path**—Trace label-switched path operations.
- **periodic**—Trace the operation of hello and keepalive messages.
- **route**—Trace the operation of route messages.
- **state**—Trace protocol state transitions.

Tracing LDP Protocol Traffic Within FECs

You can trace LDP protocol traffic within a specific FEC. You can filter LDP trace statements based on a FEC. The following trace flags are available for this purpose: **route**, **path**, and **binding**.

The following example illustrates how you might configure the LDP **traceoptions** statement to filter LDP trace statements based on a FEC:

```
[edit protocols ldp traceoptions]
flag route filter match-on fec policy "filter-policy-for-ldp-fec";
```

This feature has the following limitations:

- The filtering capability is only available for FECs composed of IP version 4 (IPv4) prefixes.
- Layer 2 circuit FECs cannot be filtered.
- When you configure both route tracing and filtering, MPLS routes are not displayed (they are blocked by the filter).
- Filtering is determined by the policy and the configured value for the **match-on** option. When configuring the policy, be sure that the default behavior is always **reject**.
- The only **match-on** option is **fec**. Consequently, the only type of policy you should include is a route-filter policy.

Examples: Tracing LDP Protocol Traffic

Trace LDP path messages in detail:

```
[edit]
protocols {
  ldp {
    traceoptions {
      file ldp size 10m files 5;
      flag path;
    }
  }
}
```

Trace all LDP outgoing messages:

```
[edit]
protocols {
  ldp {
    traceoptions {
      file ldp size 10m files 5;
      flag packets;
    }
  }
}
```

Trace all LDP error conditions:

```
[edit]
protocols {
  ldp {
    traceoptions {
      file ldp size 10m files 5;
      flag error;
    }
  }
}
```

Trace all LDP incoming messages and all label-binding operations:

```
[edit]
protocols {
  ldp {
    traceoptions {
      file ldp size 10m files 5 world-readable;
      flag packets receive;
      flag binding;
    }
    interface all {
    }
  }
}
```

Configuring Miscellaneous LDP Properties

The following sections describe how to configure a number of miscellaneous LDP properties:

- [Configuring LDP to Use the IGP Route Metric on page 361](#)
- [Preventing Addition of Ingress Routes to the inet.0 Routing Table on page 361](#)
- [Multiple-Instance LDP and Carrier-of-Carriers VPNs on page 361](#)
- [Configuring MPLS and LDP to Pop the Label on the Ultimate-Hop Router on page 362](#)
- [Enabling LDP over RSVP-Established LSPs on page 362](#)
- [Enabling LDP over RSVP-Established LSPs in Heterogeneous Networks on page 363](#)
- [Configuring the TCP MD5 Signature for LDP Sessions on page 363](#)
- [Disabling SNMP Traps for LDP on page 364](#)

- Configuring LDP Synchronization with the IGP on LDP Links on page 364
- Configuring LDP Synchronization with the IGP on the Router on page 365
- Configuring the Label Withdrawal Timer on page 365
- Ignoring the LDP Subnet Check on page 366

Configuring LDP to Use the IGP Route Metric

Use the `track-igp-metric` statement if you want the interior gateway protocol (IGP) route metric to be used for the LDP routes instead of the default LDP route metric (the default LDP route metric is 1).

To use the IGP route metric, include the `track-igp-metric` statement:

```
track-igp-metric;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Preventing Addition of Ingress Routes to the inet.0 Routing Table

By configuring the `no-forwarding` statement, you can prevent ingress routes from being added to the `inet.0` routing table instead of the `inet.3` routing table even if you enabled the `traffic-engineering bgp-igp` statement at the `[edit protocols mpls]` or the `[edit logical-systems logical-system-name protocols mpls]` hierarchy level. By default, the `no-forwarding` statement is disabled.

To omit ingress routes from the `inet.0` routing table, include the `no-forwarding` statement:

```
no-forwarding;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Multiple-Instance LDP and Carrier-of-Carriers VPNs

By configuring multiple LDP routing instances, you can use LDP to advertise labels in a carrier-of-carriers VPN from a service provider provider edge (PE) router to a customer carrier customer edge (CE) router. This is especially useful when the carrier customer is a basic Internet service provider (ISP) and wants to restrict full Internet routes to its PE routers. By using LDP instead of BGP, the carrier customer shields its other internal routers from the Internet. Multiple-instance LDP is also useful when a carrier customer wants to provide Layer 2 or Layer 3 VPN services to its customers.

For an example of how to configure multiple LDP routing instances for carrier-of-carriers VPNs, see the *JUNOS Feature Guide*.

Configuring MPLS and LDP to Pop the Label on the Ultimate-Hop Router

The default advertised label is label 3 (Implicit Null label). If label 3 is advertised, the penultimate-hop router removes the label and sends the packet to the egress router. If ultimate-hop popping is enabled, label 0 (IPv4 Explicit Null label) is advertised. Ultimate-hop popping ensures that any packets traversing an MPLS network include a label.

To configure ultimate-hop popping, include the `explicit-null` statement:

```
explicit-null;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.



NOTE: Juniper Networks routers queue packets based on the incoming label. Routers from other vendors might queue packets differently. Keep this in mind when working with networks containing routers from multiple vendors.

For more information about labels, see “Label Description” on page 25 and “Label Allocation” on page 26.

Enabling LDP over RSVP-Established LSPs

You can run LDP over LSPs established by RSVP, effectively tunneling the LDP-established LSP through the one established by RSVP. To do so, enable LDP on the `lo0.0` interface (see “Enabling and Disabling LDP” on page 338). You must also configure the LSPs over which you want LDP to operate by including the `ldp-tunneling` statement at the `[edit protocols mpls label-switched-path lsp-name]` hierarchy level:

```
[edit]
protocols {
  mpls {
    label-switched-path lsp-name {
      from source;
      to destination;
      ldp-tunneling;
    }
  }
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

For more information about tunneling LDP LSPs, see “Tunneling LDP LSPs in RSVP LSPs” on page 331.

Enabling LDP over RSVP-Established LSPs in Heterogeneous Networks

Some other vendors use an OSPF metric of 1 for the loopback address. Juniper Networks routers use an OSPF metric of 0 for the loopback address. This might require that you manually configure the RSVP metric when deploying LDP tunneling over RSVP LSPs in heterogeneous networks.

When a Juniper Networks router is linked to another vendor's router through an RSVP tunnel, and LDP tunneling is also enabled, by default the Juniper Networks router might not use the RSVP tunnel to route traffic to the LDP destinations downstream of the other vendor's egress router if the RSVP path has a metric of 1 larger than the physical OSPF path.

To ensure that LDP tunneling functions properly in heterogeneous networks, you can configure OSPF to ignore the RSVP LSP metric by including the `ignore-lsp-metrics` statement:

```
ignore-lsp-metrics;
```

You can configure this statement at the following hierarchy levels:

- [edit protocols ospf traffic-engineering shortcuts]
- [edit logical-systems *logical-system-name* protocols ospf traffic-engineering shortcuts]

To enable LDP over RSVP LSPs, you also still need to complete the procedure in “Enabling LDP over RSVP-Established LSPs” on page 362.

Configuring the TCP MD5 Signature for LDP Sessions

You can configure an MD5 signature for an LDP TCP connection to protect against the introduction of spoofed TCP segments into LDP session connection streams.

A router using the MD5 signature option is configured with a password for each peer for which authentication is required. The password is stored encrypted.

LDP hello adjacencies can still be created even when peering interfaces are configured with different security signatures. However, the TCP session cannot be authenticated and is never established.



NOTE: If you apply an MD5 signature to an LDP interface with an established session, it drops the TCP connection and all the associated label bindings to the FEC entries for that session. The session regenerates the database information for that session once both interfaces agree on a common security method and password.

To configure an MD5 signature for an LDP TCP connection, include the `session` and `authentication-key` statement:

```
session address {
  authentication-key md5-authentication-key;
```

```
}
```

For a list of hierarchy levels at which you can include these statements, see the statement summary section for the **session** statement.

Use the **session** statement to configure the address for the remote end of the LDP session.

The *md5-authentication-key* (password) can be up to 69 characters long. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks.

Disabling SNMP Traps for LDP

Whenever an LDP LSP makes a transition from up to down, or down to up, the router sends an SNMP trap. However, it is possible to disable the LDP SNMP traps on a router, logical system, or routing instance.

For information about the LDP SNMP traps and the proprietary LDP MIB, see the *JUNOS Network Management Configuration Guide*.

To disable SNMP traps for LDP, specify the **trap disable** option for the **log-updown** statement:

```
log-updown {
  trap disable;
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Configuring LDP Synchronization with the IGP on LDP Links

LDP is a protocol for distributing labels in non-traffic-engineered applications. Labels are distributed along the best path determined by the IGP. If synchronization between LDP and the IGP is not maintained, the LSP goes down. When LDP is not fully operational on a given link (a session is not established and labels are not exchanged), the IGP advertises the link with the maximum cost metric. The link is not preferred but remains in the network topology.

LDP synchronization is supported only on active point-to-point interfaces and LAN interfaces configured as point-to-point under the IGP. LDP synchronization is not supported during graceful restart.

To advertise the maximum cost metric until LDP is operational for synchronization, include the **ldp-synchronization** statement:

```
ldp-synchronization {
  disable;
  hold-time seconds;
}
```

To disable synchronization, include the **disable** statement. To configure the time period to advertise the maximum cost metric for a link that is not fully operational, include the **hold-time** statement.

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

Configuring LDP Synchronization with the IGP on the Router

To configure LDP synchronization with the IGP on the router, include the **igp-synchronization** statement and specify a number of seconds as the value of the **holddown-interval** option:

```
igp-synchronization holddown-interval seconds;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

Configuring the Label Withdrawal Timer

The label withdrawal timer delays sending a label withdrawal message for a FEC to a neighbor. When an IGP link to a neighbor fails, the label associated with the FEC has to be withdrawn from all the upstream routers if the neighbor is the next hop for the FEC. After the IGP converges and a label is received from a new next hop, the label is readvertised to all the upstream routers. This is the typical network behavior. By delaying label withdrawal by a small amount of time (for example, until the IGP converges and the router receives a new label for the FEC from the downstream next hop), the label withdrawal and sending a label mapping soon could be avoided. The **label-withdrawal-delay** statement allows you to configure this delay time. By default, the delay is 60 seconds.

If the router receives the new label before the timer runs out, the label withdrawal timer is canceled. However, if the timer runs out, the label for the FEC is withdrawn from all of the upstream routers.

By default, LDP waits for 60 seconds before withdrawing labels to avoid ressignaling LSPs multiple times while the IGP is reconverging. To configure the label withdrawal delay time in seconds, include the **label-withdrawal-delay** statement:

```
label-withdrawal-delay seconds;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

Ignoring the LDP Subnet Check

In JUNOS Release 8.4 and later releases, an LDP source address subnet check is performed during the neighbor establishment procedure. The source address in the LDP link hello packet is matched against the interface address. This causes an interoperability issue with some other vendors' equipment.

To disable the subnet check, include the **allow-subnet-mismatch** statement:

```
allow-subnet-mismatch;
```

This statement can be included at the following hierarchy levels:

- [edit protocols ldp interface *interface-name*]
- [edit logical-systems *logical-system-name* protocols ldp interface *interface-name*]

Chapter 16

Summary of LDP Configuration Statements

This chapter provides a reference for each LDP configuration statement. The statements are organized alphabetically.

allow-subnet-mismatch

Syntax	allow-subnet-mismatch;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp interface <i>interface-name</i>], [edit protocols ldp interface <i>interface-name</i>]
Release Information	Statement introduced in JUNOS Release 9.3.
Description	Ignore the LDP subnet check. For JUNOS Release 8.4 and later releases, an LDP source address subnet check was added for the neighbor establishment procedure. The source address in the LDP link hello packet is matched against the interface address.
Default	The source address in the LDP link hello packet is matched against the interface address.
Usage Guidelines	See “Ignoring the LDP Subnet Check” on page 366.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

authentication-key

Syntax	authentication-key <i>md5-authentication-key</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp session <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp session <i>address</i>], [edit protocols ldp session <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols ldp session <i>address</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure the MD5 authentication signature. The maximum length of the authentication signature is 69 characters.
Usage Guidelines	See “Configuring the TCP MD5 Signature for LDP Sessions” on page 363.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

bfd-liveness-detection

Syntax	<pre> bfd-liveness-detection { detection-time threshold <i>milliseconds</i>; ecmp; failure-action { remove-nexthop; remove-route; } holddown-interval <i>seconds</i>; minimum-interval <i>milliseconds</i>; minimum-receive-interval <i>milliseconds</i>; minimum-transmit-interval <i>milliseconds</i>; multiplier <i>detection-time-multiplier</i>; no-adaptation; transmit-interval { minimum-interval <i>milliseconds</i>; threshold <i>milliseconds</i>; } } </pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp oam], [edit logical-systems <i>logical-system-name</i> protocols ldp oam fec <i>address</i>], [edit protocols ldp oam], [edit protocols ldp oam fec <i>address</i>]
Release Information	Statement introduced in JUNOS Release 7.6. Support for the bfd-liveness-detection statement at the [edit protocols ldp oam fec <i>address</i>] hierarchy level and the ecmp option were added in JUNOS Release 9.0. Support for the failure-action statement with the remove-nexthop and remove-route options and the holddown-interval statement were added in JUNOS Release 9.4.
Description	Enable Bidirectional Forwarding Detection (BFD) for all MPLS LSPs or for just a specific LSP.
Options	<p>minimum-interval—Minimum transmit and receive interval. Range: 50 through 255,000 milliseconds Default: 50</p> <p>minimum-receive-interval—Minimum receive interval. Range: 50 through 255,000 milliseconds Default: 50</p> <p>minimum-transmit-interval—Minimum transmit interval. Range: 50 through 255,000 milliseconds Default: 50</p> <p>multiplier—Detection time multiplier. Range: 50 through 255 Default: 3</p>

The other options are explained separately.

Usage Guidelines See “Configuring BFD for LDP LSPs” on page 351.

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

deaggregate

Syntax deaggregate | no-deaggregate;

Hierarchy Level [edit logical-systems *logical-system-name* protocols ldp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp],
[edit protocols ldp],
[edit routing-instances *routing-instance-name* protocols ldp]

Release Information Statement introduced before JUNOS Release 7.4.

Description Control forwarding equivalence class (FEC) deaggregation on the router.

Default Deaggregation is disabled on the router.

Options deaggregate—Deaggregate FECs.
no-deaggregate—Aggregate FECs.

Usage Guidelines See “Configuring FEC Deaggregation” on page 349.

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

disable

Syntax	disable;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart], [edit logical-systems <i>logical-system-name</i> protocols ldp interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options graceful-restart], [edit protocols ldp graceful-restart], [edit protocols ldp interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> routing-options graceful-restart]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Explicitly disable LDP on an interface, or explicitly disable LDP graceful restart.
Default	LDP is enabled on interfaces configured with the LDP interface statement. LDP graceful restart is automatically enabled when graceful restart is enabled under the [edit routing-options] hierarchy level.
Usage Guidelines	See “Enabling and Disabling LDP” on page 338 and “Configuring LDP Graceful Restart” on page 342.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

ecmp

Syntax	ecmp;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp oam bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> protocols ldp oam fec <i>address</i> bfd-liveness-detection], [edit protocols ldp oam bfd-liveness-detection], [edit protocols ldp oam fec <i>address</i> bfd-liveness-detection]
Release Information	Statement introduced in JUNOS Release 8.5.
Description	Cause RSVP to establish BFD sessions for all ECMP paths configured for the specified FEC. If you configure the ecmp statement, you must also configure the periodic-traceroute statement for the specified FEC. If you do not do so, the commit fails. You can configure the periodic-traceroute statement at the global hierarchy level ([edit protocols ldp oam]) while only configuring the ecmp statement for a specific FEC ([edit protocols ldp oam fec <i>address</i> bfd-liveness-detection]).
Usage Guidelines	See “Configuring ECMP-Aware BFD for RSVP LSPs” on page 353.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

egress-policy

Syntax	egress-policy [<i>policy-names</i>];
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Control the prefixes advertised into LDP.
Default	Only the loopback address is advertised.
Options	<i>policy-names</i> —Name of one or more routing policies.
Usage Guidelines	See “Configuring the Prefixes Advertised into LDP from the Routing Table” on page 348.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

explicit-null

Syntax	explicit-null;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Advertise label 0 to the egress router of a label-switched path (LSP).
Default	If you do not include the explicit-null statement in the MPLS configuration, label 3 (implicit null) is advertised.
Usage Guidelines	See “Configuring MPLS and LDP to Pop the Label on the Ultimate-Hop Router” on page 362.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

export

Syntax	export [<i>policy-names</i>];
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Apply policy filters to outbound LDP label bindings. Filters are applied to all label bindings from all neighbors.
Options	<i>policy-names</i> —Name of one or more routing policies.
Usage Guidelines	See “Filtering Outbound LDP Label Bindings” on page 345.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

failure-action

Syntax	failure-action { remove-nexthop; remove-route; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp oam bfd-livenesss-detection], [edit logical-systems <i>logical-system-name</i> protocols ldp oam fec <i>address</i> bfd-livenesss-detection], [edit protocols ldp oam bfd-livenesss-detection], [edit protocols ldp oam fec <i>address</i> bfd-livenesss-detection]
Release Information	Statement introduced in JUNOS Release 9.4.
Description	Configure route and next-hop properties in the event of a BFD session failure event on an LDP LSP. The failure event could be an existing BFD session that has gone down or could be a BFD session that never came up. LDP adds back the route or next hop when the relevant BFD session comes back up.
Options	<p>remove-nexthop—Removes a route corresponding to a next hop of the LSP's route at the ingress node when a BFD session failure event is detected.</p> <p>remove-route—Removes the route corresponding to an LSP from the appropriate routing tables when a BFD session failure event is detected. If the LSP is configured with ECMP and a BFD session corresponding to any path goes down, the route is removed.</p>
Usage Guidelines	See “Configuring a Failure Action for the BFD Session on an LDP LSP” on page 353.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

graceful-restart

Syntax	<pre>graceful-restart { disable; helper-disable; maximum-neighbor-recovery-time <i>value</i>; recovery-time <i>value</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Enable LDP graceful restart on the LDP master protocol instance or for a specific routing instance.
Usage Guidelines	See “Configuring LDP Graceful Restart” on page 342.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

hello-interval

Syntax	hello-interval <i>seconds</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols ldp targeted-hello], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp targeted-hello], [edit protocols ldp interface <i>interface-name</i>], [edit protocols ldp targeted-hello], [edit routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i>] [edit routing-instances <i>routing-instance-name</i> protocols ldp targeted-hello]
Release Information	Statement introduced before JUNOS Release 7.4. Support for LDP targeted hellos added in JUNOS Release 9.5.
Description	Controls the LDP timer that regulates how often hello messages are sent. You can control the rate both link hello messages and targeted hello messages are sent depending on the hierarchy level at which you configure the hello-interval statement.
Options	<i>seconds</i> —Length of time between transmission of hello packets. Range: 1 through 65,535 seconds Default: 5 seconds for link hello messages, 15 seconds for targeted hello messages
Usage Guidelines	See “Configuring the LDP Timer for Hello Messages” on page 338.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

helper-disable

Syntax	helper-disable;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart], [edit protocols ldp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Disable helper mode for LDP graceful restart. When helper mode is disabled, a router cannot help a neighboring router that is attempting to restart LDP.
Default	Helper mode is enabled by default on all routing protocols (including LDP) that support graceful restart.
Usage Guidelines	See “Configuring LDP Graceful Restart” on page 342.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

holddown-interval

Syntax	holddown-interval <i>holddown-interval</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp oam bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> protocols ldp oam fec <i>address</i> bfd-liveness-detection], [edit protocols ldp oam bfd-liveness-detection], [edit protocols ldp oam fec <i>address</i> bfd-liveness-detection]
Release Information	Statement introduced in JUNOS Release 9.4.
Description	Specifies how long the BFD session should be up before adding the route or next hop. Specifying a time of 0 seconds causes the route or next hop to be added as soon as the BFD session comes back up.
Options	<i>holddown-interval</i> —Number of seconds the BFD session should remain up before adding the route or next hop. Default: 0 seconds Range: 0 through 65,535 seconds
Usage Guidelines	See “Configuring the Holddown Interval for the BFD Session” on page 354.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

hold-time

Syntax	hold-time <i>seconds</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols ldp targeted-hello], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp targeted-hello], [edit protocols ldp interface <i>interface-name</i>], [edit protocols ldp targeted-hello], [edit routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i>] [edit routing-instances <i>routing-instance-name</i> protocols ldp targeted-hello]
Release Information	Statement introduced before JUNOS Release 7.4. Support for LDP targeted hellos added in JUNOS Release 9.5.
Description	Specify how long an LDP node should wait for a hello message before declaring a neighbor to be down. This value is sent as part of a hello message so that each LDP node tells its neighbors how long to wait. You can specify times for both link hello messages and targeted hello messages depending on which hierarchy level you configure the hold-time statement.
Options	<i>seconds</i> —Hold-time value. Range: 1 through 65,535 seconds Default: 15 seconds for link hello messages, 45 seconds for targeted hello messages
Usage Guidelines	See “Configuring the Delay Before LDP Neighbors Are Considered Down” on page 339.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

ignore-lsp-metrics

Syntax	ignore-lsp-metrics;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ospf traffic-engineering shortcuts], [edit protocols ospf traffic-engineering shortcuts]
Release Information	Statement introduced in JUNOS Release 7.5.
Description	Cause OSPF to ignore the RSVP LSP metric. Some other vendors use an OSPF metric of 1 for the loopback address. Juniper Networks routers use an OSPF metric of 0 for the loopback address. This can cause interoperability problems when you configure LDP tunneling over RSVP LSPs in heterogeneous networks.
Usage Guidelines	See “Enabling LDP over RSVP-Established LSPs in Heterogeneous Networks” on page 363.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

igp-synchronization

Syntax	igp-synchronization holddown-interval <i>seconds</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced in JUNOS Release 9.5.
Description	Configures the time the JUNOS Software waits to allow LDP to synchronize with the IGP. For large networks with numerous FECs, it might be necessary to configure a longer value to allow enough time for network convergence.
Options	<i>holddown-interval seconds</i> —Specify the time to wait for LDP to synchronize with the IGP in seconds. Default: 10 seconds Range: 10 through 60 seconds
Usage Guidelines	See “Configuring LDP Synchronization with the IGP on the Router” on page 365.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

import

Syntax	import [<i>policy-names</i>];
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Apply policy filters to received LDP label bindings. Filters are applied to all label bindings from all neighbors.
Options	<i>policy-names</i> —Name of one or more routing policies.
Usage Guidelines	See “Filtering Inbound LDP Label Bindings” on page 343.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

ingress-policy

Syntax	ingress-policy [<i>ingress-policy-names</i>];
Hierarchy Level	[edit logical-system <i>logical-system-name</i> protocols ldp oam] [edit protocols ldp oam]
Release Information	Statement introduced in JUNOS Release 9.4.
Description	Enable an OAM ingress policy.
Options	<i>ingress-policy-names</i> —Specify the names of the ingress policies.
Usage Guidelines	See “Configuring OAM Ingress Policies for LDP” on page 354.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

interface

Syntax	<pre>interface <i>interface-name</i> { disable; hello-interval <i>seconds</i>; hold-time <i>seconds</i>; transport-address (interface loopback); }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Enable LDP on one or more router interfaces.
Default	LDP is disabled on all interfaces.
Options	<i>interface-name</i> —Name of an interface. To configure all interfaces, specify all . The remaining statements are explained separately.
Usage Guidelines	See “Enabling and Disabling LDP” on page 338.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

keepalive-interval

Syntax	keepalive-interval <i>seconds</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Set the keepalive interval value.
Options	<i>seconds</i> —Keepalive value. Range: 1 through 65,535 Default: 10 seconds
Usage Guidelines	See “Configuring the Interval for LDP Keepalive Messages” on page 341.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

keepalive-timeout

Syntax	keepalive-timeout <i>seconds</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Set the keepalive timeout value. The keepalive timeout defines the amount of time that the neighbor LDP node waits before determining that the session has failed.
Options	<i>seconds</i> —Keepalive timeout value. Range: 1 through 65,535 Default: 30 seconds
Usage Guidelines	See “Configuring the LDP Keepalive Timeout” on page 341.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

I2-smart-policy

Syntax	I2-smart-policy;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced in JUNOS Release 8.4.
Description	Prevent LDP from exporting IPv4 FECs over sessions with Layer 2 neighbors only. IPv4 FECs received over such sessions are filtered out.
Usage Guidelines	See “Configuring LDP IPv4 FEC Filtering” on page 350.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

label-withdrawal-delay

Syntax	label-withdrawal-delay <i>seconds</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced in JUNOS Release 9.1.
Description	Delay the withdrawal of labels to reduce router workload during IGP convergence.
Options	<i>seconds</i> —Configure the number of seconds to wait before withdrawing labels for the LDP LSPs. Default: 60 seconds Range: 0 through 300 seconds
Usage Guidelines	See “Configuring the Label Withdrawal Timer” on page 365.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

ldp

Syntax	<code>ldp { ... }</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols], [edit protocols], [edit routing-instances <i>routing-instance-name</i> protocols]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Enable LDP routing on the router. You must include the ldp statement in the configuration to enable LDP on the router.
Default	LDP is disabled on the router.
Usage Guidelines	See “Minimum LDP Configuration” on page 338 and “Enabling and Disabling LDP” on page 338.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

ldp-synchronization

Syntax	<code>ldp-synchronization { disable; hold-time seconds; }</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ospf interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf interface <i>interface-name</i>], [edit protocols ospf interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols ospf interface <i>interface-name</i>]
Release Information	Statement introduced in JUNOS Release 7.5.
Description	Enable synchronization by advertising the maximum cost metric until LDP is operational on the link.
Options	The other statements are explained separately.
Usage Guidelines	“Configuring LDP Synchronization with the IGP on LDP Links” on page 364
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

log-updown

Syntax	log-updown { trap disable; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Disable LDP traps on the router, logical system, or routing instance.
Options	trap disable—Disable LDP traps. Default: LDP traps are enabled on the router.
Usage Guidelines	See “Disabling SNMP Traps for LDP” on page 364.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

maximum-neighbor-recovery-time

Syntax	maximum-neighbor-recovery-time <i>seconds</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart], [edit protocols ldp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart]
Release Information	Statement introduced before JUNOS Release 7.4. Statement changed from maximum-recovery-time to maximum-neighbor-recovery-time in JUNOS Release 9.1.
Description	Specify the maximum amount of time to wait before giving up an attempt to gracefully restart.
Options	<i>seconds</i> —Configure the maximum recovery time, in seconds. Range: 120 through 1800 seconds Default: 140 seconds
Usage Guidelines	See “Configuring Recovery Time and Maximum Recovery Time” on page 343.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

no-deaggregate

See deaggregate.

no-forwarding

Syntax no-forwarding;

Hierarchy Level [edit logical-systems *logical-system-name* protocols ldp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp],
[edit protocols ldp],
[edit routing-instances *routing-instance-name* protocols ldp]

Release Information Statement introduced before JUNOS Release 7.4.

Description Do not add ingress routes to the inet.0 routing table even if traffic-engineering bgp-igp (configured at the [edit protocols mpls] hierarchy level) is enabled.

Default The no-forwarding statement is disabled. Ingress routes are added to the inet.0 routing table instead of the inet.3 routing table when traffic-engineering bgp-igp is enabled.

Usage Guidelines See “Preventing Addition of Ingress Routes to the inet.0 Routing Table” on page 361.

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

oam

```

Syntax  oam {
            bfd-liveness-detection {
                detection-time threshold milliseconds;
                ecmp;
                failure-action {
                    remove-nexthop;
                    remove-route;
                }
                holddown-interval milliseconds;
                ingress-policy ingress-policy-name
                minimum-interval milliseconds;
                minimum-receive-interval milliseconds;
                minimum-transmit-interval milliseconds;
                multiplier detection-time-multiplier;
                no-adaptation;
                transmit-interval {
                    minimum-interval milliseconds;
                    threshold milliseconds;
                }
            }
            fec fec-address;
            ingress-policy ingress-policy-name;
            lsp-ping-interval seconds;
            periodic-traceroute {
                disable;
                exp exp-value;
                fanout fanout-value;
                frequency minutes;
                paths number-of-paths;
                retries retry-attempts;
                source address;
                ttl tll-value;
                wait seconds;
            }
        }

```

Hierarchy Level [edit protocols ldp]

Release Information Statement introduced in JUNOS Release 7.6. The `lsp-ping-interval` option was introduced in JUNOS Release 9.4.

Description Enable OAM for all of the LDP LSPs or for a specific LDP LSP.

Options `fec fec-address`—Specify the forwarding equivalence class (FEC) address.

`lsp-ping-interval seconds`—Specify the duration of the LSP ping interval in seconds.
To issue a ping on an LDP-signaled LSP, use the `ping mpls ldp` command.

The remaining statements are explained separately.

Usage Guidelines See “Configuring BFD for LDP LSPs” on page 351.

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

periodic-traceroute

Syntax	<pre> periodic-traceroute { disable; exp <i>exp-value</i>; fanout <i>fanout-value</i>; frequency <i>minutes</i>; paths <i>number-of-paths</i>; retries <i>retry-attempts</i>; source <i>address</i>; ttl <i>ttl-value</i>; wait <i>seconds</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp oam], [edit logical-systems <i>logical-system-name</i> protocols ldp oam fec <i>fec-address</i>], [edit protocols ldp oam], [edit protocols ldp oam fec <i>fec-address</i>]
Release Information	Statement introduced in JUNOS Release 8.4. Support added at the [edit protocols ldp oam] and [edit logical-systems <i>logical-system-name</i> protocols ldp oam] hierarchy levels in JUNOS 9.0.
Description	Enable tracing of forwarding equivalence classes (FECs) for LDP LSPs.
Options	<p>disable—Disable tracing for a specific FEC. This option is available at the [edit protocols ldp oam fec <i>fec-address</i> periodic-traceroute] and [edit logical-systems <i>logical-system-name</i> protocols ldp oam fec <i>fec-address</i> periodic-traceroute] hierarchy levels only.</p> <p>exp <i>exp-value</i>—Specify the class of service to use when sending probes. Default: 7 Range: 0 through 7</p> <p>fanout <i>fanout-value</i>—Specify the maximum number of next hops to search per node. Default: 16 Range: 1 through 16</p> <p>frequency <i>minutes</i>—Specify the interval between traceroute attempts. Default: 60 minutes Range: 15 through 120 minutes</p> <p>paths <i>number-of-paths</i>—Specify the maximum number of paths to search. Default: 3 Range: 1 through 255</p> <p>retries <i>retry-attempts</i>—Specify the number of attempts to send a probe to a specific node before giving up. Default: 3 Range: 1 through 9</p>

source address—Specify the IPv4 source address to use when sending probes.

ttl value—Specify the maximum time-to-live value. Nodes that are beyond this value are not traced.

Default: 64

Range: 1 through 255

wait seconds—Specify the wait interval before resending a probe packet.

Default: 10 seconds

Range: 5 through 15 seconds

Usage Guidelines See “Configuring LDP LSP Traceroute” on page 354.

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

policing

Syntax

```
policing {
  fec fec-address {
    ingress-traffic filter-name;
    transit-traffic filter-name;
  }
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols ldp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp],
[edit protocols ldp],
[edit routing-instances *routing-instance-name* protocols ldp]

Release Information Statement introduced before JUNOS Release 7.4.

Description Enable policing of forwarding equivalence classes (FECs) for LDP.

Options **fec *fec-address***—Specify the address for the FEC.

ingress-traffic *filter-name*—Specify the name of the filter for policing ingress FEC traffic.

transit-traffic *filter-name*—Specify the name of the filter for policing transit FEC traffic.

Usage Guidelines See “Configuring Policers for LDP FECs” on page 349.

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

preference

Syntax	<code>preference preference;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Set the route preference level for LDP routes.
Options	<i>preference</i> —Preferred value. Range: 0 through 255 Default: 9
Usage Guidelines	See “Configuring LDP Route Preferences” on page 341.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

recovery-time

Syntax	<code>recovery-time seconds;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart], [edit protocols ldp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the amount of time a router waits for LDP to restart gracefully.
Options	<i>seconds</i> —Configure the recovery time, in seconds. Range: 120 through 1800 seconds Default: 140 seconds
Usage Guidelines	See “Configuring Recovery Time and Maximum Recovery Time” on page 343.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

session

Syntax	session <i>address</i> { authentication-key <i>authentication-key</i> ; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the LDP session to which you want to attach the Transmission Control Protocol (TCP) MD5 signature. Configure the address for the remote end of the LDP session. The remaining statement is explained separately.
Usage Guidelines	See “Configuring the TCP MD5 Signature for LDP Sessions” on page 363.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

strict-targeted-hellos

Syntax	strict-targeted-hellos;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Prevents LDP sessions from being established with remote neighbors that have not been specifically configured. LDP peers will not respond to targeted hellos coming from a source that is not one of the configured remote neighbors.
Usage Guidelines	See “Enabling Strict Targeted Hellos for LDP” on page 340.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

targeted-hello

Syntax	targeted-hello { hello-interval <i>seconds</i> ; hold-time <i>seconds</i> ; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced in JUNOS Release 9.5.
Description	Allows you to specify the LDP timer and LDP hold time for targeted hellos.
Options	The remaining statements are explained separately.
Usage Guidelines	See “Configuring the LDP Timer for Hello Messages” on page 338 and “Configuring the Delay Before LDP Neighbors are Considered Down” on page 339.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

traceoptions

Syntax	<pre> traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier> <disable>; } </pre>
Hierarchy Level	<pre> [edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp] </pre>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	LDP protocol-level trace options.
Default	The default LDP protocol-level trace options are inherited from the routing protocols <code>traceoptions</code> statement included at the <code>[edit routing-options]</code> hierarchy level.
Options	<p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as <code>all</code>.</p> <p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <code>/var/log</code>. We recommend that you place LDP tracing output in the file <code>ldp-log</code>.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named <i>trace-file</i> reaches its maximum size, it is renamed <i>trace-file.0</i>, then <i>trace-file.1</i>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>Range: 2 through 1000</p> <p>Default: 2 files</p> <p>If you specify a maximum number of files, you must also include the <code>size</code> statement to specify the maximum file size.</p> <p>flag <i>flag</i>—Tracing operation to perform. To specify more than one tracing operation, include multiple <code>flag</code> statements.</p> <ul style="list-style-type: none"> ■ address—Operation of address and address withdrawal messages ■ binding—Label-binding operations ■ error—Error conditions ■ event—Protocol events ■ initialization—Operation of initialization messages ■ label—Operation of label request, label map, label withdrawal, and label release messages

- **notification**—Operation of notification messages
- **packets**—Equivalent to setting **address**, **initialization**, **label**, **notification**, and **periodic** flags
- **path**—Label-switched path operations
- **periodic**—Operation of hello and keepalive messages
- **route**—Operation of route messages
- **state**—Protocol state transitions

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information.
- **disable**—Disable this trace flag.
- **filter**—Filter to apply to this flag. The **filter** flag modifier can be applied only to the **route**, **path**, and **binding** flags. This flag modifier has the following options:
 - **match-on**—Match on argument specified. The **match-on** option has the following suboption: **fec**, a filter based on the FEC associated with the traced object.
 - **policy *policy-name***—Specify the filter policy.
- **receive**—Packets being received.
- **send**—Packets being transmitted.

no-world-readable—(Optional) Disallow any user to read the log file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches this size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Syntax: **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 1 MB

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of files.

world-readable—(Optional) Allow any user to read the log file.

Usage Guidelines See “Tracing LDP Protocol Traffic” on page 358 and the *JUNOS Network Management Configuration Guide*.

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

track-igp-metric

Syntax	track-igp-metric;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Cause the IGP route metric to be used for the LDP routes instead of the default LDP route metric (the default LDP route metric is 1).
Usage Guidelines	See “Configuring LDP to Use the IGP Route Metric” on page 361.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

traffic-statistics

Syntax	<pre>traffic-statistics { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; interval <i>seconds</i>; no-penultimate-hop; }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]</pre>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	LDP traffic statistics display the amount of traffic passed through a router for a particular FEC.
Options	<p>file <i>filename</i>—Name of the file to receive the output of the LDP statistics operation. Enclose the name within quotation marks. All files are placed in the directory <code>/var/log</code>.</p> <p>files <i>number</i>—(Optional) Maximum number of LDP statistics files. When a statistics file named <i>ldp-stat</i> reaches its maximum size, it is renamed <i>ldp-stat.0</i>, then <i>ldp-stat.1</i>, and so on, until the maximum number of LDP statistics files is reached. Then the oldest file is overwritten. Range: 2 through 1000 Default: 2 files</p> <p>If you specify a maximum number of files, you also must include the size statement to specify the maximum file size.</p> <p>interval <i>seconds</i>—(Optional) Specifies the interval at which the statistics are polled and written to the file. Default: 300 seconds (5 minutes)</p> <p>no-penultimate-hop—(Optional) Do not collect traffic statistics on the penultimate hop router.</p> <p>no-world-readable—(Optional) Prevents all users from reading the log file.</p> <p>size <i>size</i>—(Optional) Maximum size of each statistics file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a statistics file named <i>ldp-stat</i> reaches this size, it is renamed <i>ldp-stat.0</i>. When <i>ldp-stat</i> again reaches this size, <i>ldp-stat.0</i> is renamed <i>ldp-stat.1</i> and <i>ldp-stat</i> is renamed <i>ldp-stat.0</i>. This renaming scheme continues until the maximum number of statistics files is reached. Then the oldest statistics file is overwritten. Syntax: <i>xk</i> to specify KB, <i>xm</i> to specify MB, or <i>xg</i> to specify GB Range: 10 KB through the maximum file size supported on your system Default: 1 MB</p>

If you specify a maximum file size, you also must also include the `files` statement to specify the maximum number of files.

`world-readable`—(Optional) Enables log file access for all users.

Usage Guidelines See “Collecting LDP Statistics” on page 355.

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

transport-address

Syntax `transport-address (interface | router-id);`

Hierarchy Level `[edit logical-systems logical-system-name protocols ldp],`
`[edit logical-systems logical-system-name routing-instances routing-instance-name protocols`
`ldp],`
`[edit protocols ldp interface interface-name],`
`[edit routing-instances routing-instance-name protocols ldp interface interface-name]`

Release Information Statement introduced before JUNOS Release 7.4.

Description Allow control of the transport address used by LDP.

Default `router-id`

Options `interface`—The first IP address on the interface is used as the transport address.
`router-id`—The router identifier is used as the transport address.

Usage Guidelines See “Specifying the Transport Address Used by LDP” on page 347.

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

Part 5

CCC and TCC

- CCC and TCC Overview on page 401
- CCC and TCC Configuration Guidelines on page 405
- Summary of CCC and TCC Configuration Statements on page 427

Chapter 17

CCC and TCC Overview

This chapter includes the following sections:

- CCC Overview on page 401
- Transmitting Nonstandard BPDUs on page 402
- TCC Overview on page 402
- CCC and TCC Graceful Restart on page 403

CCC Overview

Circuit cross-connect (CCC) allows you to configure transparent connections between two circuits, where a circuit can be a Frame Relay data-link connection identifier (DLCI), an Asynchronous Transfer Mode (ATM) virtual circuit (VC), a Point-to-Point Protocol (PPP) interface, a Cisco High-Level Data Link Control (HDLC) interface, or an MPLS label-switched path (LSP). Using CCC, packets from the source circuit are delivered to the destination circuit with, at most, the Layer 2 address being changed. No other processing—such as header checksums, time-to-live (TTL) decrementing, or protocol processing—is done.

CCC circuits fall into two categories: logical interfaces, which include DLCIs, VCs, virtual local area network (VLAN) IDs, PPP and Cisco HDLC interfaces, and LSPs. The two circuit categories provide three types of cross-connect:

- Layer 2 switching—Cross-connects between logical interfaces provide what is essentially Layer 2 switching. The interfaces that you connect must be of the same type.
- MPLS tunneling—Cross-connects between interfaces and LSPs allow you to connect two distant interface circuits of the same type by creating MPLS tunnels that use LSPs as the conduit.
- LSP stitching—Cross-connects between LSPs provide a way to “stitch” together two label-switched paths, including paths that fall in two different traffic engineering database areas.

For Layer 2 switching and MPLS tunneling, the cross-connect is bidirectional, so packets received on the first interface are transmitted out the second interface, and those received on the second interface are transmitted out the first. For LSP stitching, the cross-connect is unidirectional.

You can police (control) the amount of traffic flowing over CCC circuits. For more information, see the *JUNOS VPNs Configuration Guide*.

It is also possible to use the `ping` command to check the integrity of CCC LSPs. See “Pinging CCC LSPs” on page 171 for more information.

Transmitting Nonstandard BPDUs

CCC protocol (and Layer 2 Circuit and Layer 2 VPN) configurations can transmit nonstandard bridge protocol data units (BPDUs) generated by other vendors' equipment. This is the default behavior on all supported PICs and requires no additional configuration.


The following PICs are supported on M320 and T Series routers:

- 1-port Gigabit Ethernet PIC
- 2-port Gigabit Ethernet PIC
- 4-port Gigabit Ethernet PIC
- 10-port Gigabit Ethernet PIC

TCC Overview

Translational cross-connect (TCC) is a switching concept that allows you to establish interconnections between a variety of Layer 2 protocols or circuits. It is similar to CCC. However, while CCC requires the same Layer 2 encapsulations on each side of a Juniper Networks router (such as PPP-to-PPP or Frame Relay-to-Frame Relay), TCC lets you connect different types of Layer 2 protocols interchangeably. Using TCC, combinations such as PPP-to-ATM (see Figure 25 on page 402) and Ethernet-to-Frame Relay connections are possible.

Figure 25: TCC Example



The Layer 2 circuits and encapsulation types that can be interconnected by TCC are:

- Ethernet
- Extended VLANs
- PPP
- HDLC
- ATM
- Frame Relay

TCC works by removing the Layer 2 header when frames enter the router and adding a different Layer 2 header on the frames before they leave the router. In Figure 25 on page 402, the PPP encapsulation is stripped from the frames arriving at Router B, and the ATM encapsulation is added before the frames are sent to Router C.

Note that all control traffic is terminated at the interconnecting router (Router B). Examples of traffic controllers include the Link Control Protocol (LCP) and the Network Control Protocol (NCP) for PPP, keepalives for HDLC, and Local Management Interface (LMI) for Frame Relay.

TCC functionality is different from standard Layer 2 switching. TCC only swaps Layer 2 headers. No other processing, such as header checksums, TTL decrementing, or protocol handling is performed. TCC is supported for IPv4 only.

You can configure TCC for interface switching and for Layer 2 VPNs. For more information about using TCC for virtual private networks (VPNs), see the *JUNOS VPNs Configuration Guide*.


CCC and TCC Graceful Restart

CCC and TCC graceful restart allows Layer 2 connections between customer edge (CE) routers to restart gracefully. These Layer 2 connections are configured with the `remote-interface-switch` or `lsp-switch` statements. Because these CCC and TCC connections have an implicit dependency on RSVP LSPs, graceful restart for CCC and TCC uses the RSVP graceful restart capabilities.

RSVP graceful restart must be enabled on the PE routers and P routers to enable graceful restart for CCC and TCC. Also, because RSVP is used as the signaling protocol for signaling label information, the neighboring router must use helper mode to assist with the RSVP restart procedures.

Figure 26 on page 403 illustrates how graceful restart might work on a CCC connection between two CE routers.

Figure 26: Remote Interface Switch Connecting Two CE Routers Using CCC



PE Router A is the ingress for the transmit LSP from PE Router A to PE Router B and the egress for the receive LSP from PE Router B to PE Router A. With RSVP graceful restart enabled on all the PE and P routers, the following occurs when PE router A restarts:

- PE Router A preserves the forwarding state associated with the CCC routes (those from CCC to MPLS and from MPLS to CCC).
- Traffic flows without disruption from CE router to CE router.

- After the restart, PE Router A preserves the label for the LSP for which PE Router A is the egress (the receive LSP, for example). The transmit LSP from PE Router A to PE Router B can derive new label mappings, but should not cause any traffic disruption.

Chapter 18

CCC and TCC Configuration Guidelines

This chapter includes the following sections:

- Configuring Layer 2 Switching Cross-Connects Using CCC on page 405
- Configuring MPLS LSP Tunnel Cross-Connects Using CCC on page 413
- Configuring LSP Stitching Cross-Connects Using CCC on page 417
- Configuring TCC on page 418
- Configuring CCC and TCC Graceful Restart on page 424
- Configuring CCC Switching for Point-to-Multipoint LSPs on page 424


Configuring Layer 2 Switching Cross-Connects Using CCC

Layer 2 switching cross-connects join logical interfaces to form what is essentially Layer 2 switching. The interfaces that you connect must be of the same type.

Figure 27 on page 405 illustrates a Layer 2 switching cross-connect. In this topology, Router A and Router C have Frame Relay connections to Router B, which is a Juniper Networks router. Circuit cross-connect (CCC) allows you to configure Router B to act as a Frame Relay (Layer 2) switch.

To configure Router B to act as a Frame Relay switch, you configure a circuit from Router A to Router C that passes through Router B, effectively configuring Router B as a Frame Relay switch with respect to these routers. This configuration allows Router B to transparently switch packets (frames) between Router A and Router C without regard to the packets' contents or the Layer 3 protocols. The only processing that Router B performs is to translate DLCI 600 to 750.

Figure 27: Layer 2 Switching Cross-Connect



If the Router A-to-Router B and Router B-to-Router C circuits were PPP, for example, the Link Control Protocol and Network Control Protocol exchanges occur between Router A and Router C. These messages are handled transparently by Router B, allowing Router A and Router C to use various PPP options (such as header or address

compression and authentication) that Router B might not support. Similarly, Router A and Router C exchange keepalives, providing circuit-to-circuit connectivity status.

You can configure Layer 2 switching cross-connects on PPP, Cisco HDLC, Frame Relay, Ethernet, and ATM circuits. In a single cross-connect, only like interfaces can be connected.

To configure Layer 2 switching cross-connects, you must configure the following on the router that is acting as the switch (Router B in Figure 27 on page 405):

- Configuring the CCC Encapsulation for Layer 2 Switching Cross-Connects on page 406
- Configuring the CCC Connection for Layer 2 Switching Cross-Connects on page 410
- Configuring MPLS for Layer 2 Switching Cross-Connects on page 411
- Example: Configuring a Layer 2 Switching Cross-Connect on page 411

Configuring the CCC Encapsulation for Layer 2 Switching Cross-Connects

To configure Layer 2 switching cross-connects, configure the CCC encapsulation on the router that is acting as the switch (Router B in Figure 27 on page 405).



NOTE: You cannot configure families on CCC interfaces; that is, you cannot include the `family` statement at the `[edit interfaces interface-name unit logical-unit-number]` hierarchy level.

For instructions for configuring the encapsulation for Layer 2 switching cross-connects, see the following sections:

- Configuring ATM Encapsulation for Layer 2 Switching Cross-Connects on page 406
- Configuring Ethernet Encapsulation for Layer 2 Switching Cross-Connects on page 407
- Configuring Ethernet VLAN Encapsulation for Layer 2 Switching Cross-Connects on page 407
- Configuring Aggregated Ethernet Encapsulation for Layer 2 Switching Cross-Connects on page 408
- Configuring Frame Relay Encapsulation for Layer 2 Switching Cross-Connects on page 409
- Configuring PPP and Cisco HDLC Encapsulation for Layer 2 Switching Cross-Connects on page 410

Configuring ATM Encapsulation for Layer 2 Switching Cross-Connects

For ATM circuits, specify the encapsulation when configuring the virtual circuit (VC). Configure each VC as a circuit or a regular logical interface by including the following statements:

```
at-fpc/pic/port {
  atm-options {
```

```

        vpi vpi-identifier maximum-vcs maximum-vcs;
    }
    unit logical-unit-number {
        point-to-point; # Default interface type
        encapsulation encapsulation-type;
        vci vpi-identifier.vci-identifier;
    }
}

```

You can include these statements at the following hierarchy levels:

- [edit interfaces]
- [edit logical-systems *logical-system-name* interfaces]

Configuring Ethernet Encapsulation for Layer 2 Switching Cross-Connects

For Ethernet circuits, specify **ethernet-ccc** in the **encapsulation** statement. This statement configures the entire physical device. For these circuits to work, you must also configure a logical interface (unit 0).

Ethernet interfaces with standard Tag Protocol Identifier (TPID) tagging can use Ethernet CCC encapsulation. On M Series Multiservice Edge Routers, except the M320, one-port Gigabit Ethernet, two-port Gigabit Ethernet, four-port Gigabit Ethernet, and four-port Fast Ethernet PICs can use Ethernet CCC encapsulation. On T Series Core Routers and M320 routers, one-port Gigabit Ethernet and two-port Gigabit Ethernet PICs installed in FPC2 can use Ethernet CCC encapsulation. When you use this encapsulation type, you can configure the **ccc** family only.

```

fe-fpc/pic/port {
    encapsulation ethernet-ccc;
    unit 0;
}

```

You can include these statements at the following hierarchy levels:

- [edit interfaces]
- [edit logical-systems *logical-system-name* interfaces]

Configuring Ethernet VLAN Encapsulation for Layer 2 Switching Cross-Connects

An Ethernet virtual LAN (VLAN) circuit can be configured using either the **vlan-ccc** or **extended-vlan-ccc** encapsulation. If you configure the **extended-vlan-ccc** encapsulation on the physical interface, you cannot configure the **inet** family on the logical interfaces. Only the **ccc** family is allowed. If you configure the **vlan-ccc** encapsulation on the physical interface, both the **inet** and **ccc** families are supported on the logical interfaces. Ethernet interfaces in VLAN mode can have multiple logical interfaces.

For encapsulation type **vlan-ccc**, VLAN IDs from 512 through 4094 are reserved for CCC VLANs. For the **extended-vlan-ccc** encapsulation type, all VLAN IDs 1 and higher are valid. VLAN ID 0 is reserved for tagging the priority of frames.



NOTE: Some vendors use the proprietary TPIDs 0x9100 and 0x9901 to encapsulate a VLAN-tagged packet into a VLAN-CCC tunnel to interconnect a geographically separated metro Ethernet network. By configuring the **extended-vlan-ccc** encapsulation type, a Juniper Networks router can accept all three TPIDs (0x8100, 0x9100, and 0x9901).

Configure an Ethernet VLAN circuit with the **vlan-ccc** encapsulation as follows:

```
interfaces {
  type-fpc/pic/port {
    vlan-tagging;
    encapsulation vlan-ccc;
    unit logical-unit-number {
      encapsulation vlan-ccc;
      vlan-id vlan-id;
    }
  }
}
```

You can configure these statements at the following hierarchy levels:

- [edit interfaces]
- [edit logical-systems *logical-system-name* interfaces]

Configure an Ethernet VLAN circuit with the **extended-vlan-ccc** encapsulation statement as follows:

```
interfaces {
  type-fpc/pic/port {
    vlan-tagging;
    encapsulation extended-vlan-ccc;
    unit logical-unit-number {
      vlan-id vlan-id;
      family ccc;
    }
  }
}
```

You can configure these statements at the following hierarchy levels:

- [edit interfaces]
- [edit logical-systems *logical-system-name* interfaces]

Whether you configure the encapsulation as **vlan-ccc** or **extended-vlan-ccc**, you must enable VLAN tagging by including the **vlan-tagging** statement.

Configuring Aggregated Ethernet Encapsulation for Layer 2 Switching Cross-Connects

You can configure aggregated Ethernet interfaces for CCC connections and for Layer 2 virtual private networks (VPNs).

Aggregated Ethernet interfaces configured with VLAN tagging can be configured with multiple logical interfaces. The only encapsulation available for aggregated Ethernet logical interfaces is `vlan-ccc`. When you configure the `vlan-id` statement, you are limited to VLAN IDs 512 through 4094.

Aggregated Ethernet interfaces configured without VLAN tagging can be configured only with the `ethernet-ccc` encapsulation. All untagged Ethernet packets received are forwarded based on the CCC parameters.

To configure aggregated Ethernet interfaces for CCC connections, include the `ae0` statement at the `[edit interfaces]` hierarchy level:

```
[edit interfaces]
ae0 {
  encapsulation (ethernet-ccc | extended-vlan-ccc | vlan-ccc);
  vlan-tagging;
  aggregated-ether-options {
    minimum-links links;
    link-speed speed;
  }
  unit logical-unit-number {
    encapsulation vlan-ccc;
    vlan-id identifier;
    family ccc;
  }
}
```

Be aware of the following limitations when configuring CCC connections over aggregated Ethernet interfaces:

- If you configured load balancing between child links, be aware that a different hash key is used to distribute packets among the child links. Standard aggregated interfaces have family `inet` configured. An IP version 4 (IPv4) hash key (based on the Layer 3 information) is used to distribute packets among the child links. A CCC connection over an aggregated Ethernet interface has family `ccc` configured instead. Instead of an IPv4 hash key, an MPLS hash key (based on the destination media access control [MAC] address) is used to distributed packets among the child links.
- The `extended-vlan-ccc` encapsulation is not supported on the 12-port Fast Ethernet PIC and the 48-port Fast Ethernet PIC.
- The JUNOS Software does not support the Link Aggregation Control Protocol (LACP) when an aggregated interface is configured as a VLAN (with `vlan-ccc` encapsulation). LACP can be configured only when the aggregated interface is configured with the `ethernet-ccc` encapsulation.

For more information about how to configure aggregated Ethernet interfaces, see the *JUNOS Network Interfaces Configuration Guide*.

Configuring Frame Relay Encapsulation for Layer 2 Switching Cross-Connects

For Frame Relay circuits, specify the encapsulation when configuring the DLCI. Configure each DLCI as a circuit or a regular logical interface. The DLCI for regular

interfaces must be from 1 through 511. For CCC interfaces, it must be from 512 through 4094.

```

interfaces {
  type-fpc/pic/port {
    unit logical-unit-number {
      point-to-point; # Default interface type
      encapsulation encapsulation-type;
      dlci dlci-identifier;
    }
  }
}

```

You can configure these statements at the following hierarchy levels:

- [edit interfaces]
- [edit logical-systems *logical-system-name* interfaces]

Configuring PPP and Cisco HDLC Encapsulation for Layer 2 Switching Cross-Connects

For PPP and Cisco HDLC circuits, specify the encapsulation in the `encapsulation` statement. This statement configures the entire physical device. For these circuits to work, you must configure a logical interface (unit 0).

```

interfaces type-fpc/pic/port {
  encapsulation encapsulation-type;
  unit 0;
}

```

You can configure these statements at the following hierarchy levels:

- [edit interfaces *type-fpc/pic/port*]
- [edit logical-systems *logical-system-name* interfaces *type-fpc/pic/port*]

Configuring the CCC Connection for Layer 2 Switching Cross-Connects

To configure Layer 2 switching cross-connects, define the connection between the two circuits by including the `interface-switch` statement. You configure this connection on the router that is acting as the switch (Router B in Figure 27 on page 405). The connection joins the interface that comes from the circuit's source to the interface that leads to the circuit's destination. When you specify the interface names, include the logical portion of the name, which corresponds to the logical unit number. The cross-connect is bidirectional, so packets received on the first interface are transmitted out the second interface, and those received on the second interface are transmitted out the first.

```

interface-switch connection-name {
  interface interface-name.unit-number;
  interface interface-name.unit-number;
}

```


You can include this statement at the following hierarchy levels:

- [edit protocols connections]
- [edit logical-systems *logical-system-name* protocols connections]

Configuring MPLS for Layer 2 Switching Cross-Connects

For Layer 2 switching cross-connects to work, you must configure MPLS by including the interface statement. The following is a minimal MPLS configuration:

```
interface (interface-name | all);
```


You can include this statement at the following hierarchy levels:

- [edit protocols mpls]
- [edit logical-systems *logical-system-name* protocols mpls]

Example: Configuring a Layer 2 Switching Cross-Connect

Configure a full-duplex Layer 2 switching cross-connect between Router A and Router C, using a Juniper Networks router, Router B, as the virtual switch. See the topology in Figure 28 on page 411 and Figure 29 on page 412.

Figure 28: Topology of a Frame Relay Layer 2 Switching Cross-Connect




```
[edit]
interfaces {
  so-1/0/0 {
    encapsulation frame-relay-ccc;
    unit 1 {
      point-to-point;
      encapsulation frame-relay-ccc;
      dlci 600;
    }
  }
  so-2/0/0 {
    encapsulation frame-relay-ccc;
    unit 2 {
      point-to-point;
      encapsulation frame-relay-ccc;
      dlci 750;
    }
  }
}
protocols {
  connections {
```

```

        interface-switch router-a-to-router-c {
            interface so-1/0/0.1;
            interface so-2/0/0.2;
        }
    }
    mpls {
        interface all;
    }
}

```

Figure 29: Sample Topology of a VLAN Layer 2 Switching Cross-Connect



```

[edit]
interfaces {
    ge-2/1/0 {
        vlan-tagging;
        encapsulation vlan-ccc;
        unit 0 {
            encapsulation vlan-ccc;
            vlan-id 600;
        }
    }
    ge-2/2/0 {
        vlan-tagging;
        encapsulation vlan-ccc;
        unit 0 {
            encapsulation vlan-ccc;
            vlan-id 600;
        }
        unit 1 {
            family inet {
                vlan-id 1;
                address 10.9.200.1/24;
            }
        }
    }
}
protocols {
    mpls {
        interface all;
    }
    connections {
        interface-switch layer2-sw {
            interface ge-2/1/0.0;
            interface ge-2/2/0.0;
        }
    }
}

```

Configuring MPLS LSP Tunnel Cross-Connects Using CCC

MPLS tunnel cross-connects between interfaces and LSPs allow you to connect two distant interface circuits of the same type by creating MPLS tunnels that use LSPs as the conduit. The topology in Figure 30 on page 413 illustrates an MPLS LSP tunnel cross-connect. In this topology, two separate networks, in this case ATM access networks, are connected through an IP backbone. CCC allows you to establish an LSP tunnel between the two domains. With LSP tunneling, you tunnel the ATM traffic from one network across a SONET backbone to the second network by using an MPLS LSP.

Figure 30: MPLS LSP Tunnel Cross-Connect

ERROR: Unresolved graphic fileref = "g017093.eps" not found in "`\\cmsxml\default\main\TechPubsWorkInProgress\STAGING\images\`".

When traffic from Router A (VC 234) reaches Router B, it is encapsulated and placed into an LSP, which is sent through the backbone to Router C. At Router C, the label is removed, and the packets are placed onto the ATM permanent virtual circuit (PVC) (VC 591) and sent to Router D. Similarly, traffic from Router D (VC 591) is sent over an LSP to Router B, then placed on VC 234 to Router A.

You can configure LSP tunnel cross-connect on PPP, Cisco HDLC, Frame Relay, and ATM circuits. In a single cross-connect, only like interfaces can be connected.

When you use MPLS tunnel cross-connects to support IS-IS, you must ensure that the LSP's maximum transmission unit (MTU) can, at a minimum, accommodate a 1492-octet IS-IS protocol data unit (PDU) in addition to the link-level overhead associated with the technology being connected.

For the tunnel cross-connects to work, the IS-IS frame size on the edge routers (Routers A and D in Figure 30 on page 413) must be smaller than the LSP's MTU.



NOTE: Frame size values do not include the frame check sequence (FCS) or delimiting flags.

To determine the LSP MTU required to support IS-IS, use the following calculation:

IS-IS MTU (minimum 1492, default 1497) + frame overhead + 4 (MPLS shim header)
= Minimum LSP MTU

The framing overhead varies based on the encapsulation being used. The following lists the IS-IS encapsulation overhead values for various encapsulations:

- ATM
 - AAL5 multiplex—8 bytes (RFC 1483)
 - VC multiplex—0 bytes
- Frame Relay

- Multiprotocol—2 bytes (RFCs 1490 and 2427)
- VC multiplex—0 bytes
- HDLC—4 bytes
- PPP—4 bytes
- VLAN—21 bytes (802.3/LLC)

For IS-IS to work over VLAN-CCC, the LSP's MTU must be at least 1513 bytes (or 1518 for 1497-byte PDUs). If you increase the size of a Fast Ethernet MTU above the default of 1500 bytes, you might need to explicitly configure jumbo frames on intervening equipment.

To modify the MTU, include the `mtu` statement when configuring the logical interface family at the `[edit interfaces interface-name unit logical-unit-number encapsulation family]` hierarchy level. For more information about setting the MTU, see the *JUNOS Network Interfaces Configuration Guide*.

To configure an LSP tunnel cross-connect, you must configure the following on the interdomain router (Router B in Figure 30 on page 413):

- Configuring the CCC Encapsulation for LSP Tunnel Cross-Connects on page 414
- Configuring the CCC Connection for LSP Tunnel Cross-Connects on page 415
- Example: Configuring an LSP Tunnel Cross-Connect on page 416

Configuring the CCC Encapsulation for LSP Tunnel Cross-Connects

To configure LSP tunnel cross-connects, you must configure the CCC encapsulation on the ingress and egress routers (Router B and Router C, respectively, in Figure 30 on page 413).



NOTE: You cannot configure families on CCC interfaces; that is, you cannot include the `family` statement at the `[edit interfaces interface-name unit logical-unit-number]` hierarchy level.

For PPP or Cisco HDLC circuits, include the `encapsulation` statement to configure the entire physical device. For these circuits to work, you must configure logical unit 0 on the interface.

```
type-fpc/pic/port {
  encapsulation (ppp-ccc | cisco-hdlc-ccc);
  unit 0;
}
```

You can include these statements at the following hierarchy levels:

- `[edit interfaces]`
- `[edit logical-systems logical-system-name interfaces]`

For ATM circuits, specify the encapsulation when configuring the VC by including the following statements. For each VC, you configure whether it is a circuit or a regular logical interface.

```
at-fpc/pic/port {
  atm-options {
    vpi vpi-identifier maximum-vcs maximum-vcs;
  }
  unit logical-unit-number {
    point-to-point; # Default interface type
    encapsulation atm-ccc-vc-mux;
    vci vpi-identifier.vci-identifier;
  }
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces]
- [edit logical-systems *logical-system-name* interfaces]

For Frame Relay circuits, include the following statements to specify the encapsulation when configuring the DLCI. For each DLCI, you configure whether it is a circuit or a regular logical interface. The DLCI for regular interfaces must be in the range 1 through 511. For CCC interfaces, it must be in the range 512 through 1022.

```
type-fpc/pic/port {
  encapsulation frame-relay-ccc;
  unit logical-unit-number {
    point-to-point; # default interface type
    encapsulation frame-relay-ccc;
    dlci dlci-identifier;
  }
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces]
- [edit logical-systems *logical-system-name* interfaces]

For more information about the `encapsulation` statement, see the *JUNOS Network Interfaces Configuration Guide*.

Configuring the CCC Connection for LSP Tunnel Cross-Connects

To configure LSP tunnel cross-connects, include the `remote-interface-switch` statement to define the connection between the two circuits on the ingress and egress routers (Router B and Router C, respectively, in Figure 30 on page 413). The connection joins the interface or LSP that comes from the circuit's source to the interface or LSP that leads to the circuit's destination. When you specify the interface name, include the logical portion of the name, which corresponds to the logical unit number. For the cross-connect to be bidirectional, you must configure cross-connects on two routers.

```
remote-interface-switch connection-name {
```

```

interface interface-name.unit-number;
transmit-lsp label-switched-path;
receive-lsp label-switched-path;
}

```

You can include these statements at the following hierarchy levels:

- [edit protocols connections]
- [edit logical-systems *logical-system-name* protocols connections]

Example: Configuring an LSP Tunnel Cross-Connect

Configure a full-duplex MPLS LSP tunnel cross-connect from Router A to Router D, passing through Router B and Router C. See the topology in Figure 31 on page 416.

Figure 31: Example Topology of MPLS LSP Tunnel Cross-Connect

ERROR: Unresolved graphic fileref = "g017094.eps" not found in "\\cmsxml\default\main\TechPubsWorkInProgress\STAGING\images\".

On Router B:

```

[edit]
interfaces {
  at-7/1/1 {
    atm-options {
      vpi 1 maximum-vcs 600;
    }
    unit 1 {
      point-to-point; # default interface type
      encapsulation atm-ccc-vc-mux;
      vci 1.234;
    }
  }
}
protocols {
  connections {
    remote-interface-switch router-b-to-router-c {
      interface at-7/1/1.1;
      transmit-lsp lsp1;
      receive-lsp lsp2;
    }
  }
}

```

On Router C:

```

[edit]
interfaces {
  at-3/0/0 {
    atm-options {
      vpi 2 maximum-vcs 600;
    }
    unit 2 {

```

```


        point-to-point; # default interface type
        encapsulation atm-ccc-vc-mux;
        vci 2.591;
    }
}
}
protocols {
    connections {
        remote-interface-switch router-b-to-router-c {
            interface at-3/0/0.2;
            transmit-lsp lsp2;
            receive-lsp lsp1;
        }
    }
}

```

Configuring LSP Stitching Cross-Connects Using CCC

LSP stitching cross-connects “stitch” together LSPs to join two LSPs. For example, they stitch together LSPs that fall in two different traffic engineering database areas. The topology in Figure 32 on page 417 illustrates an LSP stitching cross-connect. In this topology, the network is divided into two traffic engineering domains. CCC allows you to establish an LSP between the two domains by stitching together LSPs from the two domains. For LSP stitching to work, the LSPs must be dynamic LSPs, not static.

Figure 32: LSP Stitching Cross-Connect



Without LSP stitching, a packet traveling from Router A to Router C is encapsulated on Router A (the ingress router for the first LSP), decapsulated on Router B (the egress router), and then reencapsulated on Router B (the ingress router for the second LSP). With LSP stitching, you connect LSP1 and LSP2 into a single, stitched LSP, which means that the packet is encapsulated once (on Router A) and decapsulated once (on Router C).

You can use LSP stitching to create a seamless LSP for LSPs carrying any kind of traffic.

To configure LSP stitching cross-connects, configure the two LSPs that you are stitching together on the two ingress routers. Then on the interdomain router (Router B in Figure 32 on page 417), you define the connection between the two LSPs. The connection joins the LSP that comes from the connection’s source to the LSP that leads to the connection’s destination.

```

protocols {
    connections {

```

```

    lsp-switch connection-name {
        transmit-lsp label-switched-path;
        receive-lsp label-switched-path;
    }
}

```


You can configure these statements at the following hierarchy levels:

- [edit protocols connections]
- [edit logical-systems *logical-system-name* protocols connections]

Example: Configuring an LSP Stitching Cross-Connect

Configure a full-duplex LSP stitching cross-connect between Router A and Router C. To do this, you configure Router B, which is the interdomain router. See the topology in Figure 33 on page 418.

Figure 33: Example Topology of LSP Stitching Cross-Connect



```

[edit]
protocols {
  connections interface-switch {
    lsp-switch router-a-to-router-c {
      transmit-lsp lsp2;
      receive-lsp lsp1;
    }
  }
  connections {
    lsp-switch router-c-to-router-a {
      receive-lsp lsp3;
      transmit-lsp lsp4;
    }
  }
}

```

Configuring TCC

This section describes how to configure translational cross-connect (TCC). Extensive examples on how to configure TCC for interface switching and for Layer 2.5 VPNs are available in the *JUNOS Feature Guide*.

To configure TCC, you must perform the following tasks on the router that is acting as the switch:

- Configuring the Encapsulation for Layer 2 Switching TCCs on page 419
- Configuring the Connection for Layer 2 Switching TCCs on page 422
- Configuring MPLS for Layer 2 Switching TCCs on page 423

Configuring the Encapsulation for Layer 2 Switching TCCs

To configure a Layer 2 switching TCC, specify the TCC encapsulation on the desired interfaces of the router that is acting as the switch.



NOTE: You cannot configure standard protocol families on TCC or CCC interfaces. Only the CCC family is allowed on CCC interfaces, and only the TCC family is allowed on TCC interfaces.

For Ethernet circuits and Ethernet extended VLAN circuits, you must also configure the Address Resolution Protocol (ARP). See “Configuring ARP for Ethernet and Ethernet Extended VLAN Encapsulations” on page 422.

-
- Configuring PPP and Cisco HDLC Encapsulation for Layer 2 Switching TCCs on page 419
 - Configuring ATM Encapsulation for Layer 2 Switching TCCs on page 420
 - Configuring Frame Relay Encapsulation for Layer 2 Switching TCCs on page 420
 - Configuring Ethernet Encapsulation for Layer 2 Switching TCCs on page 420
 - Configuring Ethernet Extended VLAN Encapsulation for Layer 2 Switching TCCs on page 421
 - Configuring ARP for Ethernet and Ethernet Extended VLAN Encapsulations on page 422

Configuring PPP and Cisco HDLC Encapsulation for Layer 2 Switching TCCs

For PPP and Cisco HDLC circuits, configure the encapsulation type for the entire physical device by specifying the appropriate value for the `encapsulation` statement. For these circuits to work, you must also configure the logical interface `unit 0`.

```
encapsulation (ppp-tcc | cisco-hdlc-tcc);
unit 0;
}
```

You can include these statements at the following hierarchy levels:

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

Configuring ATM Encapsulation for Layer 2 Switching TCCs

For ATM circuits, configure the encapsulation type by specifying the appropriate value for the `encapsulation` statement in the virtual circuit (VC) configuration. Specify whether each VC is a circuit or a regular logical interface.

```

atm-options {
    vpi vpi-identifier maximum-vcs maximum-vcs;
}
unit logical-unit-number {
    encapsulation (atm-tcc-vc-mux | atm-tcc-snap);
    point-to-point;
    vci vpi-identifier.vci-identifier;
}

```

You can include these statements at the following hierarchy levels:

- [edit interfaces *at-fpc/pic/port*]
- [edit logical-systems *logical-system-name* interfaces *at-fpc/pic/port*]

Configuring Frame Relay Encapsulation for Layer 2 Switching TCCs

For Frame Relay circuits, configure the encapsulation type by specifying the value `frame-relay-tcc` for the `encapsulation` statement when configuring the data-link connection identifier (DLCI). You configure each DLCI as a circuit or a regular logical interface. The DLCI for regular interfaces must be in the range from 1 through 511, but for TCC and CCC interfaces it must be in the range from 512 through 1022.

```

encapsulation frame-relay-tcc;
unit logical-unit-number {
    point-to-point;
    encapsulation frame-relay-tcc;
    dlci dlci-identifier;
}

```

You can include these statements at the following hierarchy levels:

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

Configuring Ethernet Encapsulation for Layer 2 Switching TCCs

For Ethernet TCC circuits, configuring the encapsulation type for the entire physical device by specifying the value `ethernet-tcc` for the `encapsulation` statement.

You must also specify static values for a remote address and a proxy address at the [edit interfaces *interface-name* unit *unit-number* family tcc] or [edit logical-systems *logical-system-name* interfaces *interface-name* unit *unit-number* family tcc] hierarchy level.

The remote address is associated with the TCC switching router's Ethernet neighbor; in the **remote** statement you must specify both the IP address and the media access control (MAC) address of the Ethernet neighbor. The proxy address is associated with the TCC router's other neighbor connected by the unlike link; in the **proxy** statement you must specify the IP address of the non-Ethernet neighbor.

You can configure Ethernet TCC encapsulation for the interfaces on 1-port Gigabit Ethernet, 2-port Gigabit Ethernet, 4-port Fast Ethernet, and 4-port Gigabit Ethernet PICs.

```
encapsulation ethernet-tcc;
unit logical-unit-number {
  family tcc {
    proxy {
      inet-address ip-address;
    }
    remote {
      inet-address ip-address;
      mac-address mac-address;
    }
  }
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces (fe | ge)-*fpc/pic/port*]
- [edit logical-systems *logical-system-name* interfaces (fe | ge)-*fpc/pic/port*]



NOTE: For Ethernet circuits, you must also configure the Address Resolution Protocol (ARP). See “Configuring ARP for Ethernet and Ethernet Extended VLAN Encapsulations” on page 422.

Configuring Ethernet Extended VLAN Encapsulation for Layer 2 Switching TCCs

For Ethernet extended VLAN circuits, configure the encapsulation type for the entire physical device by specifying the value **extended-vlan-tcc** for the **encapsulation** statement.

You must also enable VLAN tagging. Ethernet interfaces in VLAN mode can have multiple logical interfaces. With encapsulation type **extended-vlan-tcc**, all VLAN IDs from 0 through 4094 are valid, up to a maximum of 1024 VLANs. As with Ethernet circuits, you must also specify a proxy address and a remote address at the [edit interfaces *interface-name* unit *logical-unit-number* family tcc] or [edit logical-systems *logical-system-name* interfaces *interface-name* unit *unit-number* family tcc] hierarchy level (see “Configuring Ethernet Encapsulation for Layer 2 Switching TCCs” on page 420).

```
encapsulation extended-vlan-tcc;
vlan-tagging;
unit logical-unit-number {
```

```

vlan-id identifier;
family tcc;
proxy {
    inet-address ip-address;
}
remote {
    inet-address ip-address;
    mac-address mac-address;
}
}

```

You can configure these statements at the following hierarchy levels:

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



NOTE: For Ethernet extended VLAN circuits, you must also configure the Address Resolution Protocol (ARP). See “Configuring ARP for Ethernet and Ethernet Extended VLAN Encapsulations” on page 422.

Configuring ARP for Ethernet and Ethernet Extended VLAN Encapsulations

For Ethernet and Ethernet extended VLAN circuits with TCC encapsulation, you must also configure ARP. Because TCC simply removes one Layer 2 header and adds another, the default form of dynamic ARP is not supported; you must configure static ARP.

Because remote and proxy addresses are specified on the router performing TCC switching, you must apply the static ARP statement to the Ethernet-type interfaces of the routers that connect to the TCC-switched router. The **arp** statement must specify the IP address and the MAC address of the remotely connected neighbor by use of the unlike Layer 2 protocol on the far side of the TCC switching router.

```
arp ip-address mac mac-address;
```

You can include this statement at the following hierarchy levels:

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

Configuring the Connection for Layer 2 Switching TCCs

You must configure the connection between the two circuits of the Layer 2 switching TCC on the router acting as the switch. The connection joins the interface coming from the circuit’s source to the interface leading to the circuit’s destination. When you specify the interface names, include the logical portion of the name, which corresponds to the logical unit number. The cross-connect is bidirectional, so packets received on the first interface are transmitted from the second interface, and those received on the second interface are transmitted from the first.

To configure a connection for a local interface switch, include the following statements:

```
interface-switch connection-name {
    interface interface-name.unit-number;
}
lsp-switch connection-name {
    transmit-lsp lsp-number;
    receive-lsp lsp-number;
}
```

You can include these statements at the following hierarchy levels:

- [edit protocols connections]
- [edit logical-systems *logical-system-name* protocols connections]

To configure a connection for a remote interface switch, include the following statements:

```
remote-interface-switch connection-name {
    interface interface-name.unit-number;
    interface interface-name.unit-number;
    transmit-lsp lsp-number;
    receive-lsp lsp-number;
}
```

You can include these statements at the following hierarchy levels:

- [edit protocols connections]
- [edit logical-systems *logical-system-name* protocols connections]

Configuring MPLS for Layer 2 Switching TCCs

For a Layer 2 switching cross-connect to function, you need to configure MPLS.

To configure MPLS on an interface, configure a logical interface by including the `unit` statement:

```
unit logical-unit-number;
```

You can include this statement at the following hierarchy levels:

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

To enable MPLS, include the `interface` statement:

```
interface (interface-name | all);
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls]

- [edit logical-systems *logical-system-name* protocols mpls]



NOTE: LSP link protection does not support TCC.

Configuring CCC and TCC Graceful Restart

To enable CCC and TCC graceful restart, include the `graceful-restart` statement:

```
graceful-restart;
```

You can include this statement at the following hierarchy levels:

- [edit routing-options]
- [edit logical-systems *logical-system-name* routing-options]

The `graceful-restart` statement enables graceful restart for all protocols supporting this feature on the router. For more information about graceful restart, see the *JUNOS Routing Protocols Configuration Guide*.

CCC and TCC graceful restart depend on RSVP graceful restart. If you disable RSVP graceful restart, CCC and TCC graceful restart will not work. For more information about RSVP graceful restart, see “RSVP Graceful Restart” on page 268 and “Configuring RSVP Graceful Restart” on page 289.

Configuring CCC Switching for Point-to-Multipoint LSPs

You can configure CCC to switch traffic from interfaces to point-to-multipoint LSPs. This feature is useful for handling multicast or broadcast traffic (for example, a digital video stream).

To configure CCC switching for point-to-multipoint LSPs, you do the following:

- On the ingress provider edge (PE) router, you configure CCC to switch traffic from an incoming interface to a point-to-multipoint LSP.
- On the egress PE, you configure CCC to switch traffic from an incoming point-to-multipoint LSP to an outgoing interface.

The CCC connection for point-to-multipoint LSPs is unidirectional.

For more information on point-to-multipoint LSPs, see “Point-to-Multipoint LSPs” on page 48 and “Point-to-Multipoint LSP Configuration Guidelines” on page 141.

To configure a CCC connection for a point-to-multipoint LSP, complete the steps in the following sections:

- Configuring the Point-to-Multipoint LSP Switch on Ingress PE Routers on page 425
- Configuring the Point-to-Multipoint LSP Switch on Egress PE Routers on page 425

Configuring the Point-to-Multipoint LSP Switch on Ingress PE Routers

To configure the ingress PE router with a CCC switch for a point-to-multipoint LSP, include the `p2mp-transmit-switch` statement:

```
p2mp-transmit-switch switch-name {
    input-interface input-interface-name.unit-number;
    transmit-p2mp-lsp transmitting-lsp;
}
```

You can include the `p2mp-transmit-switch` statement at the following hierarchy levels:

- [edit protocols connections]
- [edit logical-systems *logical-system-name* protocols connections]

switch-name specifies the name of the ingress CCC switch.

input-interface *input-interface-name.unit-number* specifies the name of the ingress interface.

transmit-p2mp-lsp *transmitting-lsp* specifies the name of the transmitting point-to-multipoint LSP.

Configuring the Point-to-Multipoint LSP Switch on Egress PE Routers

To configure the CCC switch for a point-to-multipoint LSP on the egress PE router, include the `p2mp-receive-switch` statement.

```
p2mp-receive-switch switch-name {
    output-interface [ output-interface-name.unit-number ];
    receive-p2mp-lsp receptive-lsp;
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols connections]
- [edit logical-systems *logical-system-name* protocols connections]

switch-name specifies the name of the egress CCC switch.

output-interface [*output-interface-name.unit-number*] specifies the name or one or more egress interfaces.

receive-p2mp-lsp *receptive-lsp* specifies the name of the receptive point-to-multipoint LSP.

Chapter 19

Summary of CCC and TCC Configuration Statements

This chapter provides a reference for each circuit cross-connect (CCC) configuration statement. The statements are organized alphabetically.

connections

Syntax

```
connections {
  interface-switch connection-name {
    interface interface-name.unit-number;
  }
  lsp-switch connection-name {
    transmit-lsp label-switched-path;
    receive-lsp label-switched-path;
  }
  p2mp-receive-switch {
    output-interface [ interface-name.unit-number ];
    receive-p2mp-lsp receiving-point-to-multipoint-lsp;
  }
  p2mp-transmit-switch {
    input-interface interface-name.unit-number;
    transmit-p2mp-lsp transmitting-point-to-multipoint-lsp;
  }
  remote-interface-switch connection-name {
    interface interface-name.unit-number;
    receive-lsp label-switched-path;
    transmit-lsp label-switched-path;
  }
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols],
[edit protocols]

Release Information Statement introduced before JUNOS Release 7.4.

Description Define the connection between two circuits in a CCC connection.

Options The statements are explained separately.

Usage Guidelines See “CCC and TCC Configuration Guidelines” on page 405 and the *JUNOS Network Interfaces Configuration Guide*.

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

encapsulation

See the following sections:

- encapsulation (Logical Interface) on page 430
- encapsulation (Physical Interface) on page 433

encapsulation (Logical Interface)

Syntax	encapsulation (atm-ccc-cell-relay atm-ccc-vc-mux atm-tcc-vc-mux atm-cisco-nlpid atm-mlppp-llc atm-nlpid atm-ppp-llc atm-ppp-vc-mux atm-snap atm-tcc-snap atm-vc-mux ether-over-atm-llc ether-vpls-over-atm-llc frame-relay-ccc frame-relay-ppp frame-relay-tcc gre-fragmentation multilink-frame-relay-end-to-end multilink-ppp ppp-over-ether ppp-over-ether-over-atm-llc vlan-ccc vlan-ccc vlan-vpls);
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Logical link-layer encapsulation type.
Options	<p>atm-ccc-cell-relay—Use ATM cell-relay encapsulation.</p> <p>atm-ccc-vc-mux—Use ATM VC multiplex encapsulation on circuit cross-connect (CCC) circuits. When you use this encapsulation type, you can configure the family ccc only.</p> <p>atm-cisco-nlpid—Use Cisco ATM NLPID encapsulation. When you use this encapsulation type, you can configure the family inet only.</p> <p>atm-mlppp-llc—For ATM2 IQ interfaces only, use Multilink PPP over ATM adaptation layer 5 (AAL5) logical link control (LLC). For this encapsulation type, your router must be equipped with a Link Services PIC.</p> <p>atm-nlpid—Use ATM NLPID encapsulation. When you use this encapsulation type, you can configure the family inet only.</p> <p>atm-ppp-llc—For ATM2 IQ interfaces only, use PPP over ATM adaptation layer 5 (AAL5) logical link control (LLC) encapsulation.</p> <p>atm-ppp-vc-mux—For ATM2 IQ interfaces only, use PPP over ATM adaptation layer 5 (AAL5) multiplex encapsulation.</p> <p>atm-snap—Use ATM SNAP encapsulation.</p> <p>atm-tcc-snap—Use ATM SNAP encapsulation on translational cross-connect (TCC) circuits.</p> <p>atm-tcc-vc-mux—Use ATM VC multiplex encapsulation on translational cross-connect (TCC) circuits. When you use this encapsulation type, you can configure the family tcc only.</p> <p>atm-vc-mux—Use ATM VC multiplex encapsulation. When you use this encapsulation type, you can configure the family inet only.</p> <p>ether-over-atm-llc—For interfaces that carry IPv4 traffic, use Ethernet over ATM LLC encapsulation. When you use this encapsulation type, you cannot configure multipoint interfaces.</p>

ether-vpls-over-atm-llc—For ATM intelligent queuing interfaces only, use the Ethernet VPLS over ATM LLC encapsulation to bridge Ethernet interfaces and ATM interfaces over a VPLS routing instance (as described in RFC 2684). Packets from the ATM interfaces are converted to standard ENET2/802.3 encapsulated Ethernet frames with the FCS field removed.

frame-relay-ccc—Use Frame Relay encapsulation on CCC circuits. When you use this encapsulation type, you can configure the family **ccc** only.

frame-relay-ppp—Use Frame Relay encapsulation on PPP circuits.

frame-relay-tcc—Use Frame Relay encapsulation on TCC circuits for connecting unlike media. When you use this encapsulation type, you can configure the family **tcc** only.

gre-fragmentation—For adaptive services interfaces only, use GRE fragmentation encapsulation to enable fragmentation of IPv4 packets in GRE tunnels. This encapsulation clears the don't fragment (DF) bit in the packet header. If the packet's size exceeds the tunnel's maximum transmission unit (MTU) value, the packet is fragmented before encapsulation.

multilink-frame-relay-end-to-end—Use Multilink Frame Relay (MLFR) FRF.15 encapsulation. This encapsulation is used only on multilink and link services interfaces and their constituent T1 or E1 interfaces.

multilink-ppp—Use Multilink Point-to-Point Protocol (MLPPP) encapsulation. This encapsulation is used only on multilink and link services interfaces and their constituent T1 or E1 interfaces.

ppp-over-ether—For underlying Ethernet interfaces on J Series Services Routers only, use PPP over Ethernet encapsulation. When you use this encapsulation type, you cannot configure the interface address. Instead, configure the interface address on the PPP interface.

ppp-over-ether-over-atm-llc—For underlying ATM interfaces on J Series Services Routers only, use PPP over Ethernet over ATM LLC encapsulation. When you use this encapsulation type, you cannot configure the interface address. Instead, you configure the interface address on the PPP interface.

vlan-ccc—Use Ethernet virtual local area network (VLAN) encapsulation on CCC circuits. When you use this encapsulation type, you can configure the family **ccc** only.

vlan-tcc—Use Ethernet VLAN encapsulation on TCC circuits. When you use this encapsulation type, you can configure the family **tcc** only.

vlan-vpls—Use Ethernet VLAN encapsulation on virtual private LAN service (VPLS) circuits.

Usage Guidelines See “Configuring the CCC Encapsulation for Layer 2 Switching Cross-Connects” on page 406, “Configuring the CCC Encapsulation for LSP Tunnel Cross-Connects” on page 414, and “Configuring the Encapsulation for Layer 2 Switching TCCs” on page 419. For more information about how to configure interfaces, see the *JUNOS Network Interfaces Configuration Guide*.

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

encapsulation (Physical Interface)

Syntax	encapsulation (atm-ccc-cell-relay atm-pvc cisco-hdlc cisco-hdlc-ccc cisco-hdlc-tcc ethernet-ccc ethernet-over-atm ethernet-tcc ethernet-vpls extended-frame-relay-ccc extended-frame-relay-tcc extended-vlan-ccc extended-vlan-tcc extended-vlan-vpls flexible-ethernet-services flexible-frame-relay frame-relay frame-relay-ccc frame-relay-tcc frame-relay-port-ccc multilink-frame-relay-uni-nni ppp ppp-ccc ppp-tcc vlan-ccc vlan-vpls);
Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Physical link-layer encapsulation type.
Default	PPP encapsulation.
Options	<p>atm-ccc-cell-relay—Use ATM cell-relay encapsulation.</p> <p>atm-pvc—Use ATM PVC encapsulation.</p> <p>cisco-hdlc—Use Cisco-compatible HDLC framing.</p> <p>cisco-hdlc-ccc—Use Cisco-compatible HDLC framing on CCC circuits.</p> <p>cisco-hdlc-tcc—Use Cisco-compatible HDLC framing on TCC circuits for connecting unlike media.</p> <p>ethernet-ccc—Use Ethernet CCC encapsulation on Ethernet interfaces that must accept packets carrying standard TPID values.</p> <p>ethernet-over-atm—As defined in RFC 1483, this encapsulation type allows ATM interfaces to connect to devices that support only bridged-mode protocol data units (PDUs). The JUNOS Software does not completely support bridging, but accepts BPDU packets as a default gateway. If you use the router as an edge device, then the router acts as a default gateway. It accepts Ethernet LLC/SNAP frames with IP or ARP in the payload and drops the rest. For packets destined the Ethernet LAN, a route lookup is done by use of the destination IP address. If the route lookup yields a full address match, the packet is encapsulated with an LLC/SNAP and MAC header, and the packet is forwarded to the ATM interface.</p> <p>ethernet-tcc—For interfaces that carry IPv4 traffic, use Ethernet TCC encapsulation on interfaces that must accept packets carrying standard Tag Protocol Identifier (TPID) values. Ethernet TCC is not currently supported on Fast Ethernet 48-port PICs.</p> <p>ethernet-vpls—Use Ethernet VPLS encapsulation on Ethernet interfaces that have VPLS enabled and that must accept packets carrying standard TPID values.</p> <p>extended-frame-relay-ccc—Use Frame Relay encapsulation on CCC circuits. This encapsulation type allows you to dedicate DLCIs 1 through 1022 to CCC.</p>

extended-frame-relay-tcc—Use Frame Relay encapsulation on TCC circuits to connect unlike media. This encapsulation type allows you to dedicate DLCIs 1 through 1022 to TCC.

extended-vlan-ccc—Use extended VLAN encapsulation on CCC circuits with Gigabit Ethernet and 4-port Fast Ethernet interfaces that must accept packets carrying 802.1Q values.

extended-vlan-tcc—For interfaces that carry IPv4 traffic, use extended VLAN encapsulation on TCC circuits with Gigabit Ethernet interfaces on which you want to use 802.1Q tagging. Extended Ethernet TCC is not currently supported on Fast Ethernet 48-port PICs.

extended-vlan-vpls—Use extended VLAN VPLS encapsulation on Ethernet interfaces that have VLAN 802.1Q tagging and VPLS enabled and that must accept packets carrying TPIDs 0x8100, 0x9100, and 0x9901.

flexible-ethernet-services—For Gigabit Ethernet intelligent queuing interfaces only, use flexible Ethernet services encapsulation when you want to configure multiple per-unit Ethernet encapsulations. This encapsulation type allows you to configure any combination of routed, TCC, CCC, and VPLS encapsulations on a single physical port.

flexible-frame-relay—For intelligent queuing interfaces only, use flexible Frame Relay encapsulation when you want to configure multiple per-unit Frame Relay encapsulations. This encapsulation type allows you to configure any combination of TCC, CCC, or standard Frame Relay encapsulations on a single physical port. Also, each logical interface can have any DLCI value in the range 1 through 1022.

frame-relay—Use Frame Relay encapsulation.

frame-relay-ccc—Use plain Frame Relay encapsulation or Frame Relay encapsulation on circuit cross-connect (CCC) circuits.

frame-relay-port-ccc—Use Frame Relay port CCC encapsulation to transparently carry all the DLCIs between two CE routers without explicitly configuring each DLCI on the two PE routers with Frame Relay transport. When you use this encapsulation type, you can configure the **family ccc** only.

frame-relay-tcc—Use Frame Relay encapsulation on TCC circuits to connect unlike media.

multilink-frame-relay-uni-nni—Use MLFR user-to-network (UNI) network-to-network (NNI) encapsulation. This encapsulation is used only on link services interfaces functioning as FRF.16 bundles and their constituent T1 or E1 interfaces.

ppp—Use serial PPP encapsulation.

ppp-ccc—Use serial PPP encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **family ccc** only.

ppp-tcc—Use serial PPP encapsulation on TCC circuits for connecting unlike media. When you use this encapsulation type, you can configure the **family tcc** only.

vlan-ccc—Use Ethernet VLAN encapsulation on CCC circuits.

vlan-vpls—Use VLAN VPLS encapsulation on Ethernet interfaces with VLAN tagging and VPLS enabled. Interfaces with VLAN VPLS encapsulation accept packets carrying standard TPID values only.

Usage Guidelines See “Configuring the CCC Encapsulation for Layer 2 Switching Cross-Connects” on page 406, “Configuring the CCC Encapsulation for LSP Tunnel Cross-Connects” on page 414, and “Configuring the Encapsulation for Layer 2 Switching TCCs” on page 419. For more information about how to configure interfaces, see the *JUNOS Network Interfaces Configuration Guide*.

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

interface-switch

Syntax `interface-switch connection-name {
 interface interface-name.unit-number;
}`

Hierarchy Level [edit logical-systems *logical-system-name* protocols connections],
[edit protocols connections]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure Layer 2 switching cross-connects. The cross-connect is bidirectional, so packets received on the first interface are transmitted out the second interface, and those received on the second interface are transmitted out the first.

For Layer 2 switching cross-connects to work, you must also configure MPLS.

Options *connection-name*—Connection name.

interface interface-name.unit-number—Interface name. Include the logical portion of the name, which corresponds to the logical unit number.

Usage Guidelines See “Configuring the CCC Connection for Layer 2 Switching Cross-Connects” on page 410.

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

lsp-switch

Syntax	lsp-switch <i>connection-name</i> { transmit-lsp <i>label-switched-path</i> ; receive-lsp <i>label-switched-path</i> ; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols connections], [edit protocols connections]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure Layer 2 switching cross-connects.
Options	<p><i>connection-name</i>—Connection name.</p> <p><i>receive-lsp label-switched-path</i>—Name of the LSP from the connection's source.</p> <p><i>transmit-lsp label-switched-path</i>—Name of the LSP to the connection's destination.</p>
Usage Guidelines	See “CCC and TCC Configuration Guidelines” on page 405, “Configuring LSP Stitching Cross-Connects Using CCC” on page 417, and “Configuring the Connection for Layer 2 Switching TCCs” on page 422.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

p2mp-receive-switch

Syntax	<code>p2mp-receive-switch <i>point-to-multipoint-switch-name</i> { output-interface [<i>interface-name.unit-number</i>]; receive-p2mp-lsp <i>receiving-point-to-multipoint-lsp</i>; }</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols connections], [edit protocols connections]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure the CCC switch for a point-to-multipoint LSP on the egress PE router.
Options	<p><i>point-to-multipoint-switch-name</i>—Point-to-multipoint CCC receive switch name.</p> <p>output-interface <i>interface-name.unit-number</i>—Name of the egress interfaces for the point-to-multipoint LSP traffic. You can configure multiple output interfaces.</p> <p>receive-p2mp-lsp <i>receiving-point-to-multipoint-lsp</i>—Name of the point-to-multipoint LSP that is switched to the output interface.</p>
Usage Guidelines	See “Configuring the Point-to-Multipoint LSP Switch on Egress PE Routers” on page 425.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

p2mp-transmit-switch

Syntax	<pre>p2mp-transmit-switch <i>point-to-multipoint-transmit-switch-name</i> { input-interface <i>interface-name.unit-number</i>; transmit-p2mp-lsp <i>interface-name.unit-number</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols connections], [edit protocols connections]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure the CCC switch for the point-to-multipoint LSP on the ingress PE router.
Options	<p><i>point-to-multipoint-transmit-switch-name</i>—Point-to-multipoint CCC transmit switch name.</p> <p><i>input-interface input-interface-name.unit-number</i>—Specify the name of the interface carrying incoming traffic to be switched to the point-to-multipoint LSP.</p> <p><i>transmit-p2mp-lsp transmitting-point-to-multipoint-lsp</i>—Specify the name of the point-to-multipoint LSP carrying traffic to the CCC switch on the egress PE router.</p>
Usage Guidelines	See “Configuring the Point-to-Multipoint LSP Switch on Ingress PE Routers” on page 425.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

remote-interface-switch

Syntax	remote-interface-switch <i>connection-name</i> { interface <i>interface-name.unit-number</i> ; transmit-lsp <i>label-switched-path</i> ; receive-lsp <i>label-switched-path</i> ; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols connections], [edit protocols connections]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure MPLS LSP tunnel cross-connects.
Options	<p><i>connection-name</i>—Connection name.</p> <p>interface <i>interface-name.unit-number</i>—Interface name. Include the logical portion of the name, which corresponds to the logical unit number.</p> <p>receive-lsp <i>label-switched-path</i>—Name of the LSP from the connection's source.</p> <p>transmit-lsp <i>label-switched-path</i>—Name of the LSP to the connection's destination.</p>
Usage Guidelines	See “CCC and TCC Configuration Guidelines” on page 405 and “Configuring MPLS LSP Tunnel Cross-Connects Using CCC” on page 413.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

Part 6

GMPLS

- GMPLS Overview on page 443
- GMPLS Configuration Guidelines on page 449
- Hierarchy of RSVP LSPs Configuration Guidelines on page 465
- Summary of GMPLS Configuration Statements on page 471

Chapter 20

GMPLS Overview

This chapter includes the following topics:

- GMPLS Standards on page 443
- Terms and Acronyms on page 444
- Overview on page 445
- GMPLS Operation on page 446
- GMPLS and OSPF on page 447
- GMPLS and CSPF on page 447
- GMPLS Features on page 447

GMPLS Standards

The JUNOS Software supports the following Internet Engineering Task Force (IETF) RFC and Internet drafts related to GMPLS:

- RFC 3471, *Generalized Multi-Protocol Label Switching (GMPLS)-Signaling Functional Description* (generalized label request [bandwidth encoding only], generalized label [suggested label only], bidirectional LSPs [upstream label only], and control channel separation)
- RFC 3473, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions* (only Section 9, “Fault Handling”)
- RFC 4206, *Label Switched Paths (LSP) Hierarchy with Generalized Multi-Protocol Label Switching (GMPLS) Traffic Engineering (TE)*
- Internet draft draft-ietf-ccamp-gmpls-routing-09.txt, *Routing Extensions in Support of Generalized MPLS* (expires April 2004)
- Internet draft draft-ietf-ccamp-gmpls-rsvp-te-ason-02.txt, *Generalized MPLS (GMPLS) RSVP-TE Signaling in support of Automatically Switched Optical Network (ASON)* (expires January 2005)
- Internet draft draft-ietf-ccamp-gmpls-sonet-sdh-08.txt, *Generalized Multi-Protocol Label Switching Extensions for SONET and SDH Control* (expires August 2003), (SUKLM labels and SONET traffic parameters only)

- Internet draft draft-ietf-ccamp-imp-10.txt, *Link Management Protocol (LMP)* (expires April 2004)
- Internet draft draft-ietf-ccamp-ospf-gmpls-extensions-04.txt, *OSPF Extensions in Support of Generalized MPLS* (expires August 2002), (except Link Local/Remote Identifiers, Link Protection Type, Shared Risk Link Group (SRLG), and Implications of Graceful Restart)

The interface switching capability descriptor described in Section 6.4 is implemented; however, it is currently packet-switch capable only.

- Internet draft draft-ietf-mppls-bundle-04.txt, *Link Bundling in MPLS Traffic Engineering*

To access RFCs and drafts, go to the IETF website at <http://www.ietf.org/>.

Terms and Acronyms

F

Forwarding adjacency A forwarding path for sending data between GMPLS-enabled devices.

G

Generalized MPLS (GMPLS) An extension to MPLS that allows data from multiple layers to be switched over label-switched paths (LSPs). GMPLS LSP connections are possible between similar Layer 1, Layer 2, and Layer 3 devices.

GMPLS label Layer 3 identifiers, fiber port, time-division multiplexing (TDM) time slot, or dense wavelength-division multiplexing (DWDM) wavelength of a GMPLS-enabled device used as a next-hop identifier.

GMPLS LSP types The four types of GMPLS LSPs are:

- Fiber-switched capable (FSC)—LSPs are switched between two fiber-based devices, such as optical cross-connects (OXC) that operate at the level of individual fibers.
- Lambda-switched capable (LSC)—LSPs are switched between two DWDM devices, such as OXC that operate at the level of individual wavelengths.
- TDM-switched capable (TDM)—LSPs are switched between two TDM devices, such as SONET ADMs.
- Packet-switched capable (PSC)—LSPs are switched between two packet-based devices, such as routers or ATM switches.

L

Link Management Protocol

A protocol used to define a forwarding adjacency between peers and to maintain and allocate resources on the traffic engineering links.

T**Traffic engineering link**

A logical connection between GMPLS-enabled devices. Traffic engineering links can have addresses or IDs and are associated with certain resources or interfaces. They also have certain attributes (encoding-type, switching capability, bandwidth, and so on). The logical addresses can be routable, although this is not required because they are acting as link identifiers. Each traffic engineering link represents a forwarding adjacency between a pair of devices.

Overview

Traditional MPLS is designed to carry Layer 3 IP traffic using established IP-based paths and associating these paths with arbitrarily assigned labels. These labels can be configured explicitly by a network administrator, or can be dynamically assigned by means of a protocol such as LDP or RSVP.

GMPLS generalizes MPLS in that it defines labels for switching varying types of Layer 1, Layer 2, or Layer 3 traffic. GMPLS nodes can have links with one or more of the following switching capabilities:

- Fiber-switched capable (FSC)
- Lambda-switched capable (LSC)
- Time-division multiplexing (TDM) switched-capable (TSC)
- Packet-switched capable (PSC)

Label-switched paths (LSPs) must start and end on links with the same switching capability. For example, routers can establish packet-switched LSPs with other routers. The LSPs might be carried over a TDM-switched LSP between SONET add/drop multiplexers (ADMs), which in turn might be carried over a lambda-switched LSP.

The result of this extension of the MPLS protocol is an expansion in the number of devices that can participate in label switching. Lower-layer devices, such as OXCs and SONET ADMs, can now participate in GMPLS signaling and set up paths to transfer data. A router can participate in signaling optical paths across a transport network.

Two service models determine the visibility that a client node (a router, for example) has into the optical core or transport network. The first is through a user-to-network interface (UNI), which is often referred to as the overlay model. The second is known as the peer model. Juniper Networks supports both models.



NOTE: There is not necessarily a one-to-one correspondence between a physical interface and a GMPLS interface. If a GMPLS connection uses a nonchannelized physical connector, the GMPLS label can use the physical port ID. However, the label for channelized interfaces often is based on a channel or time slot. Consequently, it is best to refer to GMPLS labels as identifiers for a resource on a traffic engineering link.

To establish LSPs, GMPLS uses the following mechanisms:

- An out-of-band control channel and a data channel—RSVP messages for LSP setup are sent over an out-of-band control network. Once the LSP setup is complete and the path is provisioned, the data channel is up and can be used to carry traffic. The Link Management Protocol (LMP) is used to define and manage the data channels between a pair of nodes. You can optionally use LMP to establish and maintain LMP control channels between peers running the same JUNOS Software release.
- RSVP-TE extensions for GMPLS—RSVP-TE is already designed to signal the setup of packet LSPs. This has been extended for GMPLS to be able to request path setup for various kinds of LSPs (nonpacket) and request labels like wavelengths, time slots, and fibers as label objects.
- Bidirectional LSPs—Data can travel both ways between GMPLS devices over a single path, so nonpacket LSPs are signaled to be bidirectional.

GMPLS Operation

The basic functionality of GMPLS requires close interaction between RSVP and LMP. It works in the following sequence:

1. LMP notifies RSVP of the new entities:
 - Traffic engineering link (forwarding adjacency)
 - Resources available for the traffic engineering link
 - Control peer
2. GMPLS extracts the LSP attributes from the configuration and requests RSVP to signal one or more specific paths, which are specified by the traffic engineering link addresses.
3. RSVP determines the local traffic engineering link, corresponding control adjacency and active control channel, and transmission parameters (such as IP destination). It requests that LMP allocate a resource from the traffic engineering link with the specified attributes. If LMP finds a resource matching the attributes, label allocation succeeds. RSVP sends a PathMsg hop by hop until it reaches the target router.
4. When the target router receives the PathMsg, RSVP again requests that LMP allocate a resource based on the signaled parameters. If label allocation succeeds, the router sends back a ResvMsg.
5. If the signaling is successful, a bidirectional optical path is provisioned.

GMPLS and OSPF

You can configure OSPF for GMPLS. OSPF is an interior gateway protocol (IGP) that routes packets within a single autonomous system (AS). OSPF uses link-state information to make routing decisions.

GMPLS and CSPF

GMPLS introduces extra constraints for computing paths for GMPLS LSPs that use CSPF. These additional constraints affect the following link attributes:

- Signal type (minimum LSP bandwidth)
- Encoding type
- Switching type

These new constraints are populated in the traffic engineering database with the exchange of an interface-switching capability descriptor type, length, value (TLV) through an IGP.

The ignored constraints that are exchanged through the interface switching capability descriptor include:

- Maximum LSP bandwidth
- Maximum transmission unit (MTU)

The CSPF path computation is the same as in non-GMPLS environments, except that the links are also limited by GMPLS constraints.

Each link can have multiple interface-switching capability descriptors. All the descriptors are checked before a link is rejected.

The constraints are checked in the following order:

1. The signal type configured for the GMPLS LSP signifies the amount of bandwidth requested. If the desired bandwidth is less than the minimum LSP bandwidth, the interface-switching descriptor is rejected.
2. The encoding type of the link for the ingress and the egress interfaces should match. The encoding type is selected and stored at the ingress node after all the constraints are satisfied by the link and is used to select the link on the egress node.
3. The switching type of the links of the intermediate switches should match that of the GMPLS LSP specified in the configuration.

GMPLS Features

The JUNOS Software includes the following GMPLS functionality:

- An out-of-band control plane makes it possible to signal LSP path setup.
- RSVP-TE extensions support additional objects beyond Layer 3 packets, such as ports, time slots, and wavelengths.
- The LMP protocol creates and maintains a database of traffic engineering links and peer information. Only the static version of this protocol is supported in the JUNOS Software. You can optionally configure LMP to establish and maintain LMP control channels between peers running the same JUNOS Software release.
- Bidirectional LSPs are required between devices.
- Several GMPLS label types that are defined in RFC 3471, *Generalized MPLS —Signaling Functional Description*, such as MPLS, Generalized, SONET/SDH, Suggested, and Upstream, are supported. Generalized labels do not contain a type field, because the nodes should know from the context of their connection what type of label to expect.
- Traffic parameters facilitate GMPLS bandwidth encoding and SONET/SDH formatting.
- Other supported attributes include interface identification and errored interface identification, user-to-network (UNI)-style signaling, and secondary LSP paths.

Chapter 21

GMPLS Configuration Guidelines

This chapter describes how to configure GMPLS:

- LMP Configuration Overview on page 449
- Configuring LMP Traffic Engineering Links on page 450
- Configuring LMP Peers on page 452
- Configuring RSVP and OSPF for LMP Peer Interfaces on page 457
- Configuring MPLS Paths for GMPLS on page 458
- Tracing LMP Traffic on page 459
- Configuring MPLS LSPs for GMPLS on page 460
- Gracefully Tearing Down GMPLS LSPs on page 462

LMP Configuration Overview

You need to configure the Link Management Protocol (LMP) to define the data channel connection and the control channel connection between devices. Include the following statements at the [edit protocols link-management] hierarchy level:

```
[edit protocols link-management]
peer peer-name {
  address address;
  control-channel control-channel-name;
  lmp-control-channel control-channel-interface {
    remote-address ip-address;
  }
  lmp-protocol {
    hello-interval milliseconds;
    hello-dead-interval milliseconds;
    retransmission-interval milliseconds;
    retry-limit number;
    passive;
  }
  te-link te-link-name;
}
te-link te-link-name {
  disable;
  interface interface-name {
    disable;
    local-address ip-address;
    remote-address ip-address;
```

```

        remote-id id-number;
    }
    label-switched-path lsp-name;
    local-address ip-address;
    remote-address ip-address;
    remote-id id-number;
}
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}

```



NOTE: Although you can include GMPLS configuration statements at the [edit logical-systems *logical-system-name*] hierarchy level, GMPLS is not supported on logical systems.

For information about configuring LMP, see the following sections:

- Configuring LMP Traffic Engineering Links on page 450
- Configuring LMP Peers on page 452
- Configuring RSVP and OSPF for LMP Peer Interfaces on page 457
- Configuring MPLS Paths for GMPLS on page 458
- Tracing LMP Traffic on page 459

Configuring LMP Traffic Engineering Links

An LMP traffic engineering link acts as a data channel connection between GMPLS devices.

To configure a traffic engineering link, include the **te-link** statement at the [edit protocols link-management] hierarchy level:

```

[edit protocols link-management]
te-link te-link-name {
    disable;
    interface interface-name {
        local-address ip-address;
        remote-address ip-address;
        remote-id id-number;
    }
    label-switched-path lsp-name;
    local-address ip-address;
    remote-address ip-address;
    remote-id id-number;
}

```


Complete the procedures in the following sections to configure an LMP traffic engineering link:

- Configuring the Local IP Address for Traffic Engineering Links on page 451
- Configuring the Remote IP Address for Traffic Engineering Links on page 451
- Configuring the Remote ID for Traffic Engineering Links on page 452

When you configure a traffic engineering link that contains interfaces for an LMP peer, you must also configure a control channel. However, no control channel is required for a traffic engineering link that contains an LSP. For information on configuring control channels, see “Configuring LMP Peers” on page 452.

Configuring the Local IP Address for Traffic Engineering Links

Use the `local-address` statement to configure the local IP address associated with the traffic engineering link.

We recommend that you configure an IP address subnet for your traffic engineering link addresses that is different from the subnet configured for your physical interfaces. This configuration enables you to identify which addresses are physical and which addresses belong to the traffic engineering link.

To configure the local IP address for the traffic engineering link, include the `local-address` statement:

```
te-link te-link-name {
  interface interface-name {
    local-address ip-address;
  }
  local-address ip-address;
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Configuring the Remote IP Address for Traffic Engineering Links

You need to specify the address of the remote end of the data channel for each traffic engineering link. Use the `remote-address` statement to configure the remote IP address.

We recommend that you configure an IP address subnet for your traffic engineering link addresses that is different from the subnet configured for your physical interfaces. This enables you to identify which addresses are physical and which addresses belong to the traffic engineering link.

To configure the remote IP address for the traffic engineering link, include the `remote-address` statement:

```
te-link te-link-name {
  interface interface-name {
    remote-address ip-address;
  }
  remote-address ip-address;
}
```

```
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Configuring the Remote ID for Traffic Engineering Links

The local ID for the traffic engineering link is automatically assigned by LMP. The port identifier and labels for the interfaces (resources) in the traffic engineering link are also assigned automatically. However, you need to explicitly configure the remote ID for the traffic engineering link and the remote ID traffic engineering link interface. The remote ID for the interface must be based on the post-ID assignment of the peer node. The remote IDs are needed for static mapping of remote labels to local labels.

Before you can obtain the remote IDs for the traffic engineering link and traffic engineering link interface on the peer node, you must first configure the LMP peer, as described in “Configuring LMP Peers” on page 452. Once you have configured the LMP peer, you can obtain the traffic engineering link local ID and interface local ID by issuing the **show link-management te-link** command. Once you have these IDs, you can configure them as the remote IDs on the peer node.

To configure the remote ID for a traffic engineering link and for the traffic engineering link interface, include the **remote-id** statement:

```
te-link te-link-name {
  interface interface-name {
    remote-id id-number;
  }
  remote-id id-number;
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Configuring LMP Peers

You need to configure network peers for GMPLS. A peer is a network device that your router communicates with when setting up the control and data channels. The peer is often an optical cross-connect (OXC).

To configure an LMP peer name, include the **peer** statement at the [edit protocols link-management] hierarchy level:

```
peer peer-name {
  address ip-address;
  control-channel control-channel-interface;
  lmp-control-channel control-channel-interface {
    remote-address ip-address;
  }
  lmp-protocol {
    hello-interval milliseconds;
    hello-dead-interval milliseconds;
    retransmission-interval milliseconds;
  }
}
```

```

        retry-limit number;
        passive;
    }
    te-link te-link-name;
}

```

The following sections describe how to configure an LMP peer:

- Configuring the ID for LMP Peers on page 453
- Configuring the Interface for Control Channels Between LMP Peers on page 453
- Configuring the LMP Control Channel Interface for the Peer on page 454
- Configuring the Remote IP Address for LMP Control Channels on page 454
- Configuring Hello Message Intervals for LMP Control Channels on page 455
- Controlling Message Exchange for LMP Control Channels on page 456
- Preventing the Local Peer from Initiating LMP Negotiation on page 456
- Associating Traffic Engineering Links with LMP Peers on page 456
- Disabling the Traffic Engineering Link for LMP Peers on page 457

Configuring the ID for LMP Peers

To configure the LMP peer ID, include the **address** statement at the [edit protocols link-management peer *peer-name*] hierarchy level. The default value for the LMP peer ID is the loopback address.

```

[edit protocols link-management peer peer-name]
address ip-address;

```

Configuring the Interface for Control Channels Between LMP Peers

You must configure one or more control channels between the LMP peers. The control channels must travel across either a point-to-point link or a tunnel.

To configure the interface for the control channel, include the **control-channel** statement at the [edit protocols link-management peer *peer-name*] hierarchy level:

```

[edit protocols link-management peer peer-name]
control-channel [ interface-names ];

```

You can configure a generic routing encapsulation (GRE) interface for the control channel. This type of interface does not require a Tunnel PIC.



NOTE: You can configure GRE interfaces only for GMPLS control channels. GRE interfaces are not supported or configurable for other applications. For more information, see the *JUNOS Network Interfaces Configuration Guide*.

Configuring the LMP Control Channel Interface for the Peer

In an environment that uses LMP to establish and maintain an LMP control channel between peers, you can configure a number of attributes associated with LMP. To configure the interface to be associated with the LMP control channel for the peer, include the `lmp-control-channel` statement:

```
lmp-control-channel control-channel-interface;
```

You can configure this statement at the following hierarchy levels:

- [edit protocols link-management peer *peer-name*]
- [edit logical-systems *logical-system-name* protocols link-management peer *peer-name*]

You can configure a GRE interface for the LMP control channel. This type of interface does not require a Tunnel PIC.



NOTE: You can configure GRE interfaces only for GMPLS control channels. GRE interfaces are not supported or configurable for other applications. For more information, see the *JUNOS Network Interfaces Configuration Guide*.

When this LMP control channel interface comes up, the peers use LMP to negotiate channel parameters and configure the control channel.

The local peer repeatedly sends a Config message to the remote peer. The Config message contains the local control channel ID, the local peer's node ID, a message ID, and a CONFIG object that includes hello message attributes (the hello interval and the hello dead interval).

The channel is activated when the remote peer responds with a ConfigAck message. The remote peer does so only when its own configured hello interval and hello dead interval match the values in the received Config message or the default values. If these values do not match, the remote peer responds with a ConfigNack message. The local peer logs this event and resends the Config message until the message retry limit is reached. When the message retry limit is reached, the local peer logs that event and restarts the configuration process.

Configuring the Remote IP Address for LMP Control Channels

You need to specify the address of the remote end of the LMP control channel.

To configure the remote IP address for the LMP control channel, include the `remote-address` statement:

```
remote-address address;
```

You can configure this statement at the following hierarchy levels:

- [edit protocols link-management peer *peer-name* lmp-control-channel *control-channel-interface*]

- [edit logical-systems *logical-system-name* protocols link-management peer *peer-name* lmp-control-channel *control-channel-interface*]

Configuring Hello Message Intervals for LMP Control Channels

Hello messages are exchanged between LMP peers to maintain the control channel after LMP has activated the control channel. The LMP control channel is considered to be up only when the hello negotiation is successful. Successful negotiation consists of the local peer sending a hello message to the remote peer and receiving a hello message in response.

The LMP peers continue to exchange hello messages after the LMP control channel is up in order to maintain the channel.

The hello interval specifies the interval between periodic hello messages. The hello dead interval specifies how long the local peer waits for a hello response before it declares the LMP control channel to be down. When the channel goes down, the local peer restarts the LMP control channel negotiation and configuration process.

You can specify a hello interval from 150 through 300,000 milliseconds. The default hello interval is 150 milliseconds.

You can specify a hello dead interval from 500 through 300,000 milliseconds. The default hello dead interval is 500 milliseconds.

To configure the attributes for hello messages exchanged between LMP peers, include the `hello-interval` and `hello-dead-interval` statements:

```
hello-interval milliseconds;
hello-dead-interval milliseconds;
```

You can configure these statements at the following hierarchy levels:

- [edit protocols link-management peer *peer-name* lmp-protocol]
- [edit logical-systems *logical-system-name* protocols link-management peer *peer-name* lmp-protocol]

When an LMP control channel comes up after a successful exchange of hello messages between LMP peers, LMP uses link property correlation to verify the traffic engineering and data link information on both sides of a link. To do so, the local peer sends a LinkSummary message for each traffic engineering link governed by the LMP control channel. The LinkSummary message contains information that characterizes the traffic engineering link and each data link in the traffic engineering link.

The local peer continues sending a LinkSummary message for each link until the remote peer responds with a LinkSummaryAck message or until the message retry limit is reached. When the message retry limit is reached, the local peer logs that event and restarts the link property correlation process.

When the remote peer receives a LinkSummary message, it examines its own link information. If this information agrees with that in the LinkSummary message, the

remote peer responds with a LinkSummaryAck message. If the information is different, the remote peer responds with a LinkSummaryNack message.

Controlling Message Exchange for LMP Control Channels

You can configure message attributes that control the exchange of LMP Config and LinkSummary messages. The retransmission interval specifies the interval between resubmitted LMP messages. The retry limit specifies how many times LMP sends a message before restarting the process.

You can specify a retransmission interval from 500 through 300,000 milliseconds. The default retransmission interval is 500 milliseconds.

You can specify a retry limit from 3 through 1000 attempts. The default number of retry attempts is three.

To configure attributes governing the exchange of LMP messages between peers, include the `retransmission-interval` and `retry-limit` statements:

```
retransmission-interval milliseconds;  
retry-limit number;
```

You can configure these statements at the following hierarchy levels:

- [edit protocols link-management peer *peer-name* lmp-protocol]
- [edit logical-systems *logical-system-name* protocols link-management peer *peer-name* lmp-protocol]

Preventing the Local Peer from Initiating LMP Negotiation

You can specify that the local peer does not initiate LMP negotiation. Instead, the local peer waits for the remote peer to configure the LMP control channel.

To configure the local peer to wait for the remote peer to configure the LMP control channel, include the `passive` statement:

```
passive;
```

You can configure this statement at the following hierarchy levels:

- [edit protocols link-management peer *peer-name* lmp-protocol]
- [edit logical-systems *logical-system-name* protocols link-management peer *peer-name* lmp-protocol]

Associating Traffic Engineering Links with LMP Peers

To specify the name of a traffic engineering link to be associated with this peer, include the `te-link` statement at the [edit protocols link-management peer *peer-name*] hierarchy level:

```
[edit protocols link-management peer peer-name]
```

```
te-link te-link-name;
```

For information on how to configure a traffic engineering link, see “Configuring LMP Traffic Engineering Links” on page 450.

Disabling the Traffic Engineering Link for LMP Peers

To disable a specific traffic engineering link, include the **disable** statement:

```
disable;
```

You can configure this statement at the following hierarchy levels:

- [edit protocols link-management te-link *te-link-name*]
- [edit logical-systems *logical-system-name* protocols link-management te-link *te-link-name*]

Configuring RSVP and OSPF for LMP Peer Interfaces

After you have configured the LMP peers as described in “Configuring LMP Peers” on page 452, add the peer interfaces to RSVP and OSPF. The peer interface name must match the peer name configured in LMP. Once the peer interfaces are added to the protocols, the traffic engineering link local and remote addresses can be signaled and advertised to peers like any other interface enabled for RSVP and OSPF. These addresses act as virtual interfaces for GMPLS.



NOTE: When adding the virtual peer interfaces to RSVP and OSPF, do not configure the corresponding physical control channel interface in either protocol. If you include the **interface all** statement, you must disable RSVP and OSPF protocols manually on the control channel interface.

To configure peer interfaces in RSVP and OSPF, complete the procedures in the following sections:

- Configuring RSVP Signaling for LMP Peer Interfaces on page 457
- Configuring OSPF Routing for LMP Peer Interfaces on page 458
- Configuring the Hello Interval for LMP Peer Interfaces on page 458

Configuring RSVP Signaling for LMP Peer Interfaces

To configure RSVP signaling for LMP peers, configure the LMP peer interface by including the **peer-interface** statement at the [edit protocols **rsvp**] hierarchy level:

```
[edit protocols rsvp]
peer-interface peer-interface-name {
  (aggregate | no-aggregate);
  authentication-key key;
  disable;
  hello-interval seconds;
```

```
(reliable | no-reliable);
}
```

The statements configured at the [edit protocols rsvp peer-interface *peer-interface-name*] hierarchy level have the same functionality as the statements configured at the [edit protocols rsvp interface *interface-name*] hierarchy level.

Configuring OSPF Routing for LMP Peer Interfaces

To configure OSPF routing for LMP peers, configure the name of the LMP peer by including the peer-interface statement at the [edit protocols ospf area *area-number*] hierarchy level:

```
[edit protocols ospf area area-number]
peer-interface peer-interface-name {
  dead-interval seconds;
  disable;
  hello-interval seconds;
  retransmit-interval seconds;
  transit-delay seconds;
}
```

For information on how to configure OSPF statements, see the *JUNOS Routing Protocols Configuration Guide*.

Configuring the Hello Interval for LMP Peer Interfaces

Hello packets are used to indicate to neighboring routers that the peer interface is still up and running. The hello interval must be the same for all routers on a shared logical IP network. You can specify a hello interval from 1 through 255 seconds. The default hello interval is normally 10 seconds. For nonbroadcast networks, the default hello interval is 120 seconds.

To specify how often the router sends hello packets out the peer interface, configure the hello-interval statement:

```
hello-interval seconds;
```

You can configure this statement at the following hierarchy levels:

- [edit protocols ospf area *area-number* peer-interface *peer-interface-name*]
- [edit logical-systems *logical-system-name* protocols ospf area *area-number* peer-interface *peer-interface-name*]

Configuring MPLS Paths for GMPLS

As part of the configuration for GMPLS, you need to establish an MPLS path for each unique device connected through GMPLS. Configure the traffic engineering link remote address as the address at the [edit protocols mpls path *path-name*] hierarchy level. Constrained Shortest Path First (CSPF) is supported so you can choose either the strict or loose option with the address.

See “LMP Configuration Overview” on page 449 for information about how to obtain a traffic engineering link remote address.

To configure the MPLS path, include the `path` statement at the `[edit protocols mpls]` hierarchy level:

```
[edit protocols mpls]
path path-name {
    next-hop-address (strict | loose);
}
```

For information about how to configure MPLS paths, see “Creating Named Paths” on page 52.

Tracing LMP Traffic

To trace LMP protocol traffic, include the `traceoptions` statement at the `[edit protocols link-management]` hierarchy level:

```
[edit protocols link-management]
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

Use the `file` statement to specify the name of the file that receives the output of the tracing operation. All files are placed in the directory `/var/log`.

The following trace flags display the operations associated with the sending and receiving of various LMP messages:

- `all`—Trace all available operations
- `hello-packets`—Trace hello packets on any LMP control channel
- `init`—Output from the initialization messages
- `packets`—Trace all packets other than hello packets on any LMP control channel
- `parse`—Operation of the parser
- `process`—Operation of the general configuration
- `route-socket`—Operation of route socket events
- `routing`—Operation of the routing protocols
- `server`—Server processing operations
- `show`—Servicing operations for `show` commands
- `state`—Trace state transitions of the LMP control channels and traffic engineering links

Each flag can carry one or more of the following flag modifiers:

- `detail`—Provide detailed trace information
- `receive`—Packets being received

- **send**—Packets being transmitted

Configuring MPLS LSPs for GMPLS

To enable the proper GMPLS switching parameters, configure the label-switched path (LSP) attributes that are appropriate for your network connection. The default value for **switching-type** is **psc-1**, which is also appropriate for standard MPLS.

To configure the LSP attributes, include the **lsp-attributes** statement at the **[edit protocols mpls label-switched-path *lsp-name*]** hierarchy level:

```
[edit protocols mpls label-switched-path lsp-name]
lsp-attributes {
  encoding-type type;
  gp-id gp-id;
  signal-bandwidth type;
  switching-type type;
}
```

If you include the **no-cspf** statement in the label-switched path configuration, you must also configure primary and secondary paths, or the configuration cannot be committed.

The following sections describe how to configure each of the LSP attributes for a GMPLS LSP:

- Configuring the Encoding Type on page 460
- Configuring the GPID on page 461
- Configuring the Signal Bandwidth Type on page 461
- Configuring GMPLS Bidirectional LSPs on page 461
- Allowing Non-Packet GMPLS LSPs to Establish Paths Through Routers Running the JUNOS Software on page 462

Configuring the Encoding Type

You need to specify the encoding type of the payload carried by the LSP. It can be any of the following:

- **ethernet**—Ethernet
- **packet**—Packet
- **pdh**—Plesiochronous digital hierarchy (PDH)
- **sonet-sdh**—SONET/SDH

The default value is **packet**.

To configure the encoding type, include the **encoding-type** statement at the **[edit protocols mpls label-switched-path *lsp-name* lsp-attributes]** hierarchy level:

```
[edit protocols mpls label-switched-path lsp-name lsp-attributes]
```

```
encoding-type type;
```

Configuring the GPID

You need to specify the type of payload carried by the LSP. The payload is the type of packet underneath the MPLS label. The payload is specified by the generalized payload identifier (GPID).

You can specify the GPID with any of the following values:

- `hdlc`—High-Level Data Link Control (HDLC)
- `ethernet`—Ethernet
- `ipv4`—IP version 4 (default)
- `pos-scrambling-crc-16`—For interoperability with other vendors' equipment
- `pos-no-scrambling-crc-16`—For interoperability with other vendors' equipment
- `pos-scrambling-crc-32`—For interoperability with other vendors' equipment
- `pos-no-scrambling-crc-32`—For interoperability with other vendors' equipment
- `ppp`—Point-to-Point Protocol (PPP)

To configure the GPID, include the `gpid` statement at the `[edit protocols mpls label-switched-path lsp-name lsp-attributes]` hierarchy level:

```
[edit protocols mpls label-switched-path lsp-name lsp-attributes]
gpid gpid;
```

Configuring the Signal Bandwidth Type

The signal bandwidth type is the encoding used for path computation and admission control. To configure the signal bandwidth type, include the `signal-bandwidth` statement at the `[edit protocols mpls label-switched-path lsp-name lsp-attributes]` hierarchy level:

```
[edit protocols mpls label-switched-path lsp-name lsp-attributes]
signal-bandwidth type;
```

Configuring GMPLS Bidirectional LSPs

Because MPLS and GMPLS use the same configuration hierarchy for LSPs, it is helpful to know which LSP attributes control LSP functionality. Standard MPLS packet-switched LSPs are unidirectional, whereas GMPLS nonpacket LSPs are bidirectional.

If you use the default packet-switching type of `psc-1`, your LSP becomes unidirectional. To enable a GMPLS bidirectional LSP, you must select a non-packet-switching type option, such as `lambda`, `fiber`, or `ethernet`. Include the `switching-type` statement at the `[edit protocols mpls label-switched-path lsp-name lsp-attributes]` hierarchy level:

```
[edit protocols mpls label-switched-path lsp-name lsp-attributes]
switching-type (lambda | fiber | ethernet);
```

Allowing Non-Packet GMPLS LSPs to Establish Paths Through Routers Running the JUNOS Software

By setting the A-bit in the Admin Status object, you can enable nonpacket GMPLS LSPs to establish paths through routers that run the JUNOS Software. When an ingress router sends an RSVP PATH message with the Admin Status A-bit set, an external device (not a router running the JUNOS Software) can either perform a Layer 1 path setup test or help bring up an optical cross-connect.

When set, the A-bit in the Admin Status object indicates the administrative down status for an LSP. This feature is used specifically by nonpacket GMPLS LSPs. It does not affect control path setup or data forwarding for packet LSPs.

The JUNOS Software does not distinguish between the control path setup and data path setup. Other nodes along the network path use RSVP PATH signaling using the A-bit in a meaningful way.

To configure the Admin Status object for a GMPLS LSP, include the `admin-down` statement:

```
admin-down;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls label-switched-path *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*]

Gracefully Tearing Down GMPLS LSPs

You can gracefully tear down nonpacket GMPLS LSPs. An LSP that is torn down abruptly, a common process in a packet-switched network, can cause stability problems in nonpacket-switched networks. To maintain the stability of nonpacket-switched networks, it might be necessary to tear down LSPs gracefully.

The following sections describe how to tear down GMPLS LSPs gracefully:

- Temporarily Deleting GMPLS LSPs on page 462
- Permanently Deleting GMPLS LSPs on page 463
- Configuring the Graceful Deletion Timeout Interval on page 463

Temporarily Deleting GMPLS LSPs

You can gracefully tear down a GMPLS LSP using the `clear rsvp session gracefully` command.

This command gracefully tears down 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 endpoint of the LSP. During the second pass, the LSP is taken down. Using this command, the LSP is taken down temporarily. After the appropriate interval, the GMPLS LSP is resignaled and then reestablished.

The `clear rsvp session gracefully` command has the following properties:

- It only works on the ingress and egress routers of an RSVP session. If used on a transit router, it has the same behavior as the `clear rsvp session` command.
- It only works for nonpacket LSPs. If used with packet LSPs, it has the same behavior as the `clear rsvp session` command.

For more information, see the *JUNOS Routing Protocols and Policies Command Reference*.

Permanently Deleting GMPLS LSPs

When you disable an LSP in the configuration, the LSP is permanently deleted. By configuring the `disable` statement, you can disable a GMPLS LSP permanently. If the LSP being disabled is a nonpacket LSP, then the graceful LSP tear-down procedures that use the Admin Status object are used. If the LSP being disabled is a packet LSP, then the regular signaling procedures for LSP deletion are used.

To disable a GMPLS LSP, include the `disable` statement at any of the following hierarchy levels:

- `[edit protocols mpls label-switched-path lsp-name]`—Disables the LSP.
- `[edit protocols link-management te-link te-link-name]`—Disables a traffic engineering link.
- `[edit protocols link-management te-link te-link-name interface interface-name]`—Disables an interface used by a traffic engineering link.

Configuring the Graceful Deletion Timeout Interval

The router that initiates the graceful deletion procedure for an RSVP session waits for the graceful deletion timeout interval to ensure that all routers along the path (especially the ingress and egress routers) have prepared for the LSP to be taken down.

The ingress router initiates the graceful deletion procedure by sending the Admin Status object in the path message with the `D` bit set. The ingress router expects to receive an Resv message with the `D` bit set from the egress router. If the ingress router does not receive this message within the time specified by the graceful deletion timeout interval, it initiates a forced tear-down of the LSP by sending a PathTear message.

To configure the graceful deletion timeout interval, include the `graceful-deletion-timeout` statement at the `[edit protocols rsvp]` hierarchy level. You can configure a time from between 1 through 300 seconds. The default value is 30 seconds.

```
graceful-deletion-timeout seconds;
```

You can configure this statement at the following hierarchy levels:

- `[edit protocols rsvp]`
- `[edit logical-systems logical-system-name protocols rsvp]`

You can use the `show rsvp version` command to determine the current value configured for the graceful deletion timeout.

Chapter 22

Hierarchy of RSVP LSPs Configuration Guidelines

- Hierarchy of RSVP LSPs Standard on page 465
- Hierarchy of RSVP LSPs Terminology on page 465
- Hierarchy of RSVP LSPs Overview on page 465
- Hierarchy of RSVP LSPs on page 466
- Advertising the Forwarding Adjacency with OSPF on page 466
- Configuring a Hierarchy of RSVP LSPs on page 466

Hierarchy of RSVP LSPs Standard

For more information on how a hierarchy of RSVP LSPs functions, see RFC 4206, *Label Switched Paths (LSP) Hierarchy with Generalized Multi-Protocol Label Switching (GMPLS) Traffic Engineering (TE)*.

Hierarchy of RSVP LSPs Terminology

F

Forwarding adjacency	A traffic engineering link created by a forwarding adjacency LSP. You can create a forwarding adjacency between two routers in a network by configuring a forwarding adjacency LSP. Forwarding adjacencies can only be statically configured. However, you can configure OSPF to advertise the forwarding adjacency to other routers. When an RSVP LSP traverses a forwarding adjacency, existing MPLS features such as fast reroute continue to function.
Forwarding adjacency LSP	An RSVP LSP used to tunnel other RSVP LSPs; forms the basis for a forwarding adjacency.

Hierarchy of RSVP LSPs Overview

This chapter provides overview information and configuration instructions for hierarchies of RSVP label-switched paths (LSPs), which enable you to tunnel multiple RSVP LSPs over a single RSVP LSP.

The following sections provide an overview of how a hierarchy of RSVP LSPs functions:

- Hierarchy of RSVP LSPs on page 466
- Advertising the Forwarding Adjacency with OSPF on page 466

Hierarchy of RSVP LSPs

Forwarding adjacencies are configured and managed as point-to-point traffic engineering links by including statements at the `[edit protocols link-management]` hierarchy level. For the forwarding adjacency to function properly, you also need to make RSVP aware of the forwarding adjacency by configuring the corresponding peer interface at the `[edit protocols rsvp]` hierarchy level.

Although forwarding adjacency LSPs are configured and managed as traffic engineering links on the local router, it is not necessary to advertise these traffic engineering links to other routers in the network. However, if you want to automatically forward MPLS traffic over the forwarding adjacency or want other routers to compute paths over the forwarding adjacency, you must configure OSPF to advertise the forwarding adjacency to the other routers in the network and add the forwarding adjacency to the traffic engineering database. OSPF is the only supported interior gateway protocol (IGP).

Advertising the Forwarding Adjacency with OSPF

Once a forwarding adjacency LSP and the corresponding traffic engineering link you have configured, you can configure OSPF to advertise the forwarding adjacency. Unlike regular traffic engineering links, OSPF hellos are not exchanged between the forwarding adjacency LSP endpoints and therefore no routing adjacency is created between the forwarding adjacency endpoints. If you issue a `show ospf neighbor` command on an ingress forwarding adjacency, the command displays the egress router of the forwarding adjacency LSP as a neighbor. However, no real OSPF adjacency is established (no OSPF hellos are exchanged) between the ingress and egress routers. For display purposes only, OSPF creates a pseudo-neighbor corresponding to the peer.

You can configure forwarding adjacencies over existing MPLS networks. A forwarding adjacency LSP is signaled as a regular MPLS LSP without generalized MPLS (GMPLS) extensions. When the forwarding adjacency LSP is advertised as a traffic engineering link in OSPF, the corresponding traffic engineering link in OSPF is also advertised as a regular MPLS traffic engineering link without GMPLS extensions.

Configuring a Hierarchy of RSVP LSPs

The following sections describe how to configure a hierarchy of RSVP LSPs:

- Configuring an RSVP LSP on Ingress Routers on page 467
- Configuring Forwarding Adjacencies on page 467
- Configuring RSVP for Forwarding Adjacencies on page 468
- Advertising Forwarding Adjacencies Using OSPF on page 469

Configuring an RSVP LSP on Ingress Routers

To configure a standard RSVP LSP on the ingress router to be used as the forwarding adjacency LSP, see “LSP Configuration Overview” on page 70. This LSP requires no special configuration to function as a forwarding adjacency LSP.

Configuring Forwarding Adjacencies

A forwarding adjacency is a type of GMPLS traffic engineering link. It requires that you configure local and remote addresses to identify the link. A forwarding adjacency is associated with a specific peer router. You could configure multiple forwarding adjacencies to the same peer router.

To configure a forwarding adjacency, you need to configure the **te-link** statement at the `[edit protocols link-management]` hierarchy level:

```
[edit protocols link-management]
te-link te-link-name {
    label-switched-path lsp-name;
    local-address ip-address;
    remote-address ip-address;
}
```

For more information on how to configure GMPLS traffic engineering links, see “Configuring LMP Traffic Engineering Links” on page 450.



NOTE: Do not configure the control channel for a forwarding adjacency peer router. Configuring a control channel causes the commit to fail.

The following sections describe how to configure the **te-link** statement for a forwarding adjacency:

- Configuring the Local IP Address for Forwarding Adjacencies on page 467
- Configuring the Remote IP Address for Forwarding Adjacencies on page 468
- Configuring the LSP for Forwarding Adjacencies on page 468

Configuring the Local IP Address for Forwarding Adjacencies

To configure the local IP address for the forwarding adjacency, include the **local-address** statement:

```
local-address ip-address;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Configuring the Remote IP Address for Forwarding Adjacencies

The address of the peer router is the node ID for the forwarding adjacency LSP egress node. You configure this node ID for the forwarding adjacency using the `remote-address` statement:

```
remote-address ip-address;
```

You can include this statement at the following hierarchy levels:

- [edit protocols link-management te-link *te-link-name*],
- [edit logical-systems *logical-system-name* protocols link-management te-link *te-link-name*]

Configuring the LSP for Forwarding Adjacencies

To configure a router to function as a forwarding adjacency, the `label-switched-path` statement and specify the LSP configured in “Configuring an RSVP LSP on Ingress Routers” on page 467:

```
label-switched-path label-switched-path-name;
```

You can include this statement at the following hierarchy levels:

- [edit protocols link-management te-link *te-link-name*]
- [edit logical-systems *logical-system-name* protocols link-management te-link *te-link-name*]

Configuring RSVP for Forwarding Adjacencies

For the forwarding adjacency to function properly, RSVP must be made aware of it. Do this by specifying the name of the peer interface corresponding to the link-management peer associated with the forwarding adjacency. Including the `peer-interface` statement at the [edit protocols rsvp] hierarchy level enables RSVP to use all of the traffic engineering links configured for that peer. You can also configure RSVP control-plane parameters such as the hello interval and refresh reduction.

To configure RSVP to recognize a forwarding adjacency, include the `peer-interface` statement:

```
peer-interface peer-interface-name {
  disable;
  (aggregate | no-aggregate);
  authentication-key key;
  hello-interval seconds;
  (reliable | no-reliable);
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp]

- [edit logical-systems *logical-system-name* protocols rsvp]

For more information on how to configure the **peer-interface** statement, see “Configuring RSVP and OSPF for LMP Peer Interfaces” on page 457.

Advertising Forwarding Adjacencies Using OSPF

You can allow other routers to dynamically signal paths over a forwarding adjacency LSP by configuring OSPF. This configuration is optional.

If you configure OSPF to advertise a forwarding adjacency LSP, the LSP is added to the traffic engineering database on each router in the traffic engineering domain. Because the forwarding adjacency LSP is unidirectional, the corresponding traffic engineering link (forwarding adjacency) is also unidirectional. The forwarding adjacency LSP appears as a standard traffic engineering database half-link to all routers in the traffic engineering domain.

The Constrained Shortest Path First (CSPF) performs a bidirectional link check to ensure that traffic can flow in both directions. CSPF checks for a reverse link, either the exact reverse forwarding adjacency or another reverse link. If there is no reverse link from the forwarding adjacency LSP egress router to the forwarding adjacency LSP ingress router, the CSPF check fails.

CSPF might find another parallel reverse link. However, the LSP cannot function properly over the forwarding adjacency unless you have explicitly configured a corresponding forwarding adjacency LSP to handle the traffic flowing in the opposite direction on the forwarding adjacency LSP egress router.

To advertise the traffic engineering properties of a forwarding adjacency to a specific peer router, include the **peer-interface** statement:

```
peer-interface peer-interface-name {
  dead-interval seconds;
  disable;
  hello-interval seconds;
  retransmit-interval seconds;
  transit-delay seconds;
}
```

You can configure this statement at the following hierarchy levels:

- [edit protocols ospf area *area-name*]
- [edit logical-systems *logical-system-name* protocols ospf area *area-name*]

For more information on how to configure the **peer-interface** statement, see “Configuring RSVP and OSPF for LMP Peer Interfaces” on page 457.

Chapter 23

Summary of GMPLS Configuration Statements

This chapter provides a reference for each Generalized MPLS (GMPLS) configuration statement. The statements are organized alphabetically.

address

Syntax	address <i>ip-address</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols link-management peer <i>peer-name</i>], [edit protocols link-management peer <i>peer-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the ID of the peer.
Default	The loopback address is advertised.
Options	<i>ip-address</i> —IP address of the peer.
Usage Guidelines	See “Configuring the ID for LMP Peers” on page 453.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

admin-down

See admin-down

control-channel

Syntax	<code>control-channel control-channel-interface;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols link-management peer <i>peer-name</i>], [edit protocols link-management peer <i>peer-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the control channel interface for the peer.
Options	<i>control-channel-interface</i> —Name of the control channel interface.
Usage Guidelines	See “Configuring LMP Peers” on page 452.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

dead-interval

Syntax	<code>dead-interval seconds;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-number</i> peer-interface <i>peer-interface-name</i>], [edit protocols ospf area <i>area-number</i> peer-interface <i>peer-interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify how long OSPF and OSPF version 3 (OSPFv3) wait before declaring that a neighboring router is unavailable. This is an interval during which the router receives no hello packets from the neighbor.
Options	<i>seconds</i> —Interval to wait. Range: 1 through 65,535 Default: 40 seconds (four times the hello interval)
Usage Guidelines	See “Configuring RSVP and OSPF for LMP Peer Interfaces” on page 457.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Topics	hello-interval

disable

See the following sections:

- [disable \(GMPLS\) on page 473](#)
- [disable \(OSPF Peer Interface\) on page 473](#)

disable (GMPLS)

Syntax	disable;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols link-management te-link <i>te-link-name</i>], [edit protocols link-management te-link <i>te-link-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Disable a traffic engineering link.
Default	The configured object is enabled (operational) unless explicitly disabled.
Usage Guidelines	See “Disabling the Traffic Engineering Link for LMP Peers” on page 457.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

disable (OSPF Peer Interface)

Syntax	disable;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-number</i> peer-interface <i>peer-interface-name</i>], [edit protocols ospf area <i>area-number</i> peer-interface <i>peer-interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Disable an OSPF peer interface.
Default	The configured object is enabled (operational) unless explicitly disabled.
Usage Guidelines	See “Configuring RSVP and OSPF for LMP Peer Interfaces” on page 457.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

hello-dead-interval

Syntax	hello-dead-interval <i>milliseconds</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols link-management peer <i>peer-name</i> imp-protocol], [edit protocols link-management peer <i>peer-name</i> imp-protocol]
Release Information	Statement introduced in JUNOS Release 8.0.
Description	Specify how long the Link Management Protocol (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.
Options	<i>milliseconds</i> —Interval to wait before declaring the control channel to be dead. Range: 500 through 300,000 Default: 500 milliseconds (three times the hello interval)
Usage Guidelines	See “Configuring Hello Message Intervals for LMP Control Channels” on page 455.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Topics	hello-interval

hello-interval

See the following sections:

- hello-interval (LMP) on page 475
- hello-interval (OSPF) on page 475

hello-interval (LMP)

Syntax	hello-interval <i>milliseconds</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols link-management peer <i>peer-name</i> lmp-protocol], [edit protocols link-management peer <i>peer-name</i> lmp-protocol]
Release Information	Statement introduced in JUNOS Release 8.1.
Description	Specify how often the router sends Link Management Protocol (LMP) hello packets.
Options	<i>milliseconds</i> —Length of time between hello packets. Range: 150 through 300,000 Default: 150 milliseconds
Usage Guidelines	See “Configuring Hello Message Intervals for LMP Control Channels” on page 455.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Topics	hello-dead-interval

hello-interval (OSPF)

Syntax	hello-interval <i>seconds</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-number</i> peer-interface <i>peer-interface-name</i>], [edit protocols ospf area <i>area-number</i> peer-interface <i>peer-interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify how often the router sends hello packets out the peer interface. The hello interval must be the same for all routers on a shared logical IP network.
Options	<i>seconds</i> —Length of time between hello packets. Range: 1 through 255 Default: 10 seconds; 120 seconds (nonbroadcast networks)
Usage Guidelines	See “Configuring RSVP and OSPF for LMP Peer Interfaces” on page 457.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Topics	dead-interval.

interface

Syntax	<code>interface <i>interface-name</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols link-management te-link <i>te-link-name</i>], [edit protocols link-management te-link <i>te-link-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the egress router interface.
Options	<i>interface-name</i> —Name of the interface to the egress router.
Usage Guidelines	See “LMP Configuration Overview” on page 449.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

label-switched-path

Syntax	<code>label-switched-path <i>lsp-name</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols link-management te-link <i>te-link-name</i>], [edit protocols link-management te-link <i>te-link-name</i>]
Release Information	Statement introduced in JUNOS Release 7.4.
Description	Specify the label-switched path (LSP) to be used by the forwarding adjacency.
Options	<i>lsp-name</i> —Name of the LSP.
Usage Guidelines	See “Configuring Forwarding Adjacencies” on page 467.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

link-management

Syntax	link-management { ... }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols], [edit protocols]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Enable Link Management Protocol (LMP) on the router.
Usage Guidelines	See “LMP Configuration Overview” on page 449.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

Imp-control-channel

Syntax	Imp-control-channel <i>control-channel-interface</i> { remote-address <i>ip-address</i> ; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols link-management peer <i>peer-name</i>], [edit protocols link-management peer <i>peer-name</i>]
Release Information	Statement introduced in JUNOS Release 8.1.
Description	Specify the Link Management Protocol (LMP) control channel interface for the peer.
Options	<i>control-channel-interface</i> —Name of the control channel interface. The remaining statement is described separately in this chapter.
Usage Guidelines	See “Configuring the LMP Control Channel Interface for the Peer” on page 454.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

Imp-protocol

Syntax	<pre> imp-protocol { hello-interval <i>milliseconds</i>; hello-dead-interval <i>milliseconds</i>; passive; retransmission-interval <i>milliseconds</i>; retry-limit <i>number</i>; } </pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols link-management peer <i>peer-name</i>], [edit protocols link-management peer <i>peer-name</i>]
Release Information	Statement introduced in JUNOS Release 8.1.
Description	Configure attributes of Link Management Protocol (LMP) to establish and maintain the LMP control channel for the peer.
Options	The statements are described separately in this chapter.
Usage Guidelines	See “Configuring LMP Peers” on page 452.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

local-address

Syntax	local-address <i>ip-address</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols link-management te-link <i>te-link-name</i>], [edit logical-systems <i>logical-system-name</i> protocols link-management te-link <i>te-link-name</i> interface <i>interface-name</i>], [edit protocols link-management te-link <i>te-link-name</i>], [edit protocols link-management te-link <i>te-link-name</i> interface <i>interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the local IP address associated with the traffic engineering link.
Options	<i>local-address</i> —Local IP address of the traffic engineering link.
Usage Guidelines	See “Configuring the Local IP Address for Traffic Engineering Links” on page 451 and “Configuring the Local IP Address for Forwarding Adjacencies” on page 467.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

passive

Syntax	<code>passive;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols link-management peer <i>peer-name</i> <code>imp-protocol</code>], [edit protocols link-management peer <i>peer-name</i> <code>imp-protocol</code>]
Release Information	Statement introduced in JUNOS Release 8.1.
Description	Specify the router to not configure the Link Management Protocol (LMP) control channels, but to wait for the remote peer to configure the LMP control channels.
Usage Guidelines	See “Preventing the Local Peer from Initiating LMP Negotiation” on page 456.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

peer

Syntax	<pre> peer <i>peer-name</i> { address <i>ip-address</i>; control-channel <i>control-channel-interface</i>; imp-control-channel <i>control-channel-interface</i>; imp-protocol { hello-interval <i>milliseconds</i>; hello-dead-interval <i>milliseconds</i>; passive; retransmission-interval <i>milliseconds</i>; retry-limit <i>number</i>; } te-link <i>te-link-name</i>; } </pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols link-management], [edit protocols link-management]
Release Information	Statement introduced before JUNOS 7.4. For JUNOS 8.1, the <code>imp-protocol</code> statement and sub-statements were added.
Description	Configure a network peer.
Options	<p><i>peer-name</i>—Name of the network peer.</p> <p>The remaining statements are described separately in this chapter.</p>
Usage Guidelines	See “Configuring LMP Peers” on page 452.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

peer-interface

See the following sections:

- [peer-interface \(OSPF\) on page 480](#)
- [peer-interface \(for RSVP\) on page 480](#)

peer-interface (OSPF)

Syntax `peer-interface peer-interface-name {
 dead-interval seconds;
 disable;
 hello-interval seconds;
 retransmit-interval seconds;
 transit-delay seconds;
 }`

Hierarchy Level [edit logical-systems *logical-system-name* protocols ospf area *area-id*],
 [edit protocols ospf area *area-id*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the control channel. The peer interface name is the same as the peer interface name configured under LMP.

Options The options are explained separately.

Usage Guidelines See “Configuring RSVP and OSPF for LMP Peer Interfaces” on page 457 and “Advertising Forwarding Adjacencies Using OSPF” on page 469.

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

Related Topics *JUNOS Routing Protocols Configuration Guide*

peer-interface (for RSVP)

See [peer-interface](#)

remote-address

See the following sections:

- remote-address (for LMP Control Channel) on page 481
- remote-address (for LMP Traffic Engineering) on page 481

remote-address (for LMP Control Channel)

Syntax	remote-address <i>ip-address</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols link-management peer <i>peer-name</i> lmp-control-channel <i>control-channel-interface</i>], [edit protocols link-management peer <i>peer-name</i> lmp-control-channel <i>control-channel-interface</i>]
Release Information	Statement introduced in JUNOS Release 8.1.
Description	Specify the remote IP address for the Link Management Protocol (LMP) control channel interface.
Options	<i>ip-address</i> —Remote IP address mapped to the LMP control channel interface.
Usage Guidelines	See “Configuring the Remote IP Address for LMP Control Channels” on page 454.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

remote-address (for LMP Traffic Engineering)

Syntax	remote-address <i>ip-address</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols link-management te-link <i>te-link-name</i>], [edit logical-systems <i>logical-system-name</i> protocols link-management te-link <i>te-link-name</i> interface <i>interface-name</i>], [edit protocols link-management te-link <i>te-link-name</i>], [edit protocols link-management te-link <i>te-link-name</i> interface <i>interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the remote IP address for the traffic engineering link.
Options	<i>ip-address</i> —Remote IP address mapped to the traffic engineering link.
Usage Guidelines	See “Configuring the Remote IP Address for Traffic Engineering Links” on page 451 and “Configuring the Remote IP Address for Forwarding Adjacencies” on page 468.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

remote-id

Syntax	<code>remote-id <i>id-number</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols link-management te-link <i>te-link-name</i>], [edit logical-systems <i>logical-system-name</i> protocols link-management te-link <i>te-link-name</i> interface <i>interface-name</i>], [edit protocols link-management te-link <i>te-link-name</i>], [edit protocols link-management te-link <i>te-link-name</i> interface <i>interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the ID assigned to a traffic engineering link or an interface (resource) on the peer node.
Options	<i>id-number</i> —ID number for the remote device.
Usage Guidelines	See “Configuring the Remote ID for Traffic Engineering Links” on page 452.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

retransmission-interval

Syntax	<code>retransmission-interval <i>milliseconds</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols link-management peer <i>peer-name</i> lmp-protocol], [edit protocols link-management peer <i>peer-name</i> lmp-protocol]
Release Information	Statement introduced in JUNOS Release 8.1.
Description	Specify how often Link Management Protocol (LMP) sends Config and LinkSummary messages on the LMP control channel.
Options	<i>milliseconds</i> —Length of time between Config messages. Range: 500 through 300,000 Default: 500 milliseconds
Usage Guidelines	See “Controlling Message Exchange for LMP Control Channels” on page 456.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Topics	retry-limit

retransmit-interval

Syntax	retransmit-interval <i>seconds</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-number</i> peer-interface <i>peer-interface-name</i>], [edit protocols ospf area <i>area-number</i> peer-interface <i>peer-interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify how long the router waits to receive a link-state acknowledgment packet before retransmitting link-state advertisements to a peer interface's neighbors.
Options	<i>seconds</i> —Interval to wait for a link-state acknowledgment packet. Range: 1 through 65,535 Default: 5 seconds
Usage Guidelines	See “Configuring RSVP and OSPF for LMP Peer Interfaces” on page 457.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

retry-limit

Syntax	retry-limit <i>number</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols link-management peer <i>peer-name</i> lmp-protocol], [edit protocols link-management peer <i>peer-name</i> lmp-protocol]
Release Information	Statement introduced in JUNOS Release 8.1.
Description	Specify how many times the Link Management Protocol (LMP) sends Config and LinkSummary messages on the LMP control channel without receiving an appropriate acknowledgment before it logs a message and restarts the LMP control channel configuration process.
Options	<i>number</i> —Maximum number of times messages are sent without receiving an acknowledgment. Range: 3 through 1000 Default: 3
Usage Guidelines	See “Controlling Message Exchange for LMP Control Channels” on page 456.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Topics	retransmission-interval

te-link

Syntax	<pre> te-link <i>te-link-name</i> { disable; interface <i>interface-name</i> { disable; local-address <i>ip-address</i>; remote-address <i>ip-address</i>; remote-id <i>id-number</i>; } local-address <i>ip-address</i>; remote-address <i>ip-address</i>; remote-id <i>id-number</i>; } </pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols link-management], [edit logical-systems <i>logical-system-name</i> protocols link-management peer <i>peer-name</i>], [edit protocols link-management], [edit protocols link-management peer <i>peer-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Represent a collection of physical ports or time slots. Assign a traffic engineering link to the specified network peer.
Options	<p><i>te-link-name</i>—Name of the collection of physical ports or the name of the time slots.</p> <p><i>disable</i>—Disable the traffic engineering link or an interface to a traffic engineering link.</p> <p>The other statements are described separately in this chapter.</p>
Usage Guidelines	See “Configuring LMP Traffic Engineering Links” on page 450.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

traceoptions

Syntax	<pre> traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier> <disable>; } </pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols link-management], [edit protocols link-management]
Release Information	Statement introduced before JUNOS Release 7.4. Support for the hello-packets, packets, and state flags added in JUNOS Release 8.1.
Description	Trace options for the LMP protocol.
Options	<p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>filename—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <code>/var/log</code>.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named <i>trace-file</i> reaches its maximum size, it is renamed <i>trace-file.0</i>, then <i>trace-file.1</i>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>Range: 2 through 1000</p> <p>Default: 2 files</p> <p>If you specify a maximum number of files, you must also include the size statement to specify the maximum file size.</p> <p>flag—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements.</p> <ul style="list-style-type: none"> ■ all—Trace all available operations ■ hello-packets—Trace hello packets on any LMP control channel ■ init—Output from the initialization messages ■ packets—Trace all packets other than hello packets on any LMP control channel ■ parse—Operation of the parser ■ process—Operation of the general configuration ■ route-socket—Operation of route socket events ■ routing—Operation of the routing protocols ■ server—Server processing operations ■ show—show command servicing operations

- **state**—Trace state transitions of the LMP control channels and traffic engineering links

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

no-world-readable—(Optional) Prevents all users from reading the log file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *trace-file* reaches this size, it is renamed *trace-file.0*. When the *trace-file* again reaches this size, *trace-file.0* is renamed *trace-file.1* and *trace-file* is renamed *trace-file.0*. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Syntax: *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 1 MB

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of files.

world-readable—(Optional) Enables log file access for all users.

Usage Guidelines See “Tracing LMP Traffic” on page 459 and the *JUNOS Network Management Configuration Guide*.

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

transit-delay

Syntax	transit-delay <i>seconds</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-number</i> peer-interface <i>peer-interface-name</i>], [edit protocols ospf area <i>area-number</i> peer-interface <i>peer-interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Set the estimated time required to transmit a link-state update on the peer interface. When calculating this time, you should account for transmission and propagation delays.
Options	<i>seconds</i> —Estimated time for transmitting the link-state update. Range: 1 through 65,535 Default: 1 second
Usage Guidelines	See “Configuring RSVP and OSPF for LMP Peer Interfaces” on page 457.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

Part 7

Indexes

- Index on page 491
- Index of Statements and Commands on page 507

Index

Symbols

#, comments in configuration statements.....	xxxvi
(), in syntax descriptions.....	xxxvi
< >, in syntax descriptions.....	xxxv
[], in configuration statements.....	xxxvi
{ }, in configuration statements.....	xxxvi
(pipe), in syntax descriptions.....	xxxvi

A

adaptive rerouting.....	101, 173
adaptive statement.....	173
usage guidelines.....	101
address (tracing flag).....	394
address statement	
LMP.....	471
usage guidelines.....	453
addresses	
associating with LSPs.....	80, 204
egress router address.....	74, 247
ingress router.....	73, 198
adjust-interval statement.....	174
usage guidelines.....	90
adjust-threshold statement.....	174
usage guidelines.....	91
adjust-threshold-overflow-limit statement.....	175
usage guidelines.....	91
Admin Status object, GMPLS.....	462
admin-down statement.....	175
configuration guidelines.....	462
admin-group statement	
bypass LSPs.....	299
configuration	
bypass LSPs.....	284
LSPs.....	177
MPLS interfaces.....	176
admin-groups statement.....	178
usage guidelines.....	96
administrative groups	
admin-groups statement.....	178
configuration.....	96
exclude statement.....	192
fast reroute.....	78

include-all statement.....	201
include-any statement.....	203
advertisement messages, LDP.....	334
advertisement-hold-time statement.....	178
usage guidelines.....	107
aggregate statement	
RSVP.....	300
usage guidelines.....	275
aggregated Ethernet interfaces.....	408
aggregated interfaces.....	38
aggregation, RSVP.....	300
all (tracing flag).....	248
LMP.....	485
RSVP.....	324
allocation of labels.....	26
allow-fragmentation statement.....	179
usage guidelines.....	293
allow-subnet-mismatch statement.....	367
usage guidelines.....	366
ARP configuration.....	422
associate-backup-pe-groups statement.....	179
usage guidelines.....	147
ATM	
circuits.....	406, 414
ATM encapsulation	
Layer 2 TCC.....	420
authentication	
RSVP.....	278, 301
authentication-key statement	
LDP.....	368
usage guidelines.....	363
RSVP.....	301
usage guidelines.....	278
auto-bandwidth statement.....	180
usage guidelines.....	89
auto-policing statement.....	181
usage guidelines.....	162
automatic bandwidth allocation.....	48, 180
bandwidth monitoring.....	93
LSPs.....	88
manually trigger.....	94
threshold.....	91
automatic policers	
LSP bandwidth, changing.....	163
LSPs.....	163

overview.....	162	explicit paths.....	288
point-to-multipoint LSPs.....	164	hop limit.....	285
B		maximum number.....	286
backup paths.....	33	multiple.....	267
backup-pe-group statement.....	182	node protection, disabling.....	287
bandwidth		optimization interval.....	287
automatic allocation, LSPs.....	88	priority and preemption.....	289
LSP paths.....	106	types.....	267
RSVP reservations.....	323	bypass statement.....	303
bandwidth model.....	116	usage guidelines.....	283
bandwidth oversubscription		C	
overview.....	119	CCC	
bandwidth statement		aggregated Ethernet.....	408
fast reroute.....	183	BPDUs, nonstandard.....	402
usage guidelines.....	78	encapsulation	
link protection.....	302	Ethernet CCC.....	407
usage guidelines.....	284	example configurations.....	411, 416, 418
LSPs		graceful restart	
usage guidelines.....	127	configuration.....	424
multiclass LSPs.....	183	overview.....	403
usage guidelines.....	129	Layer 2 switching cross-connects	
RSVP.....	302	configuration.....	405
signaled LSPs.....	183	overview.....	401
usage guidelines.....	106	LSP stitching cross-connects.....	401, 417
bandwidth update threshold.....	279	MPLS tunneling	
bandwidth-model statement.....	184	cross-connects.....	401, 413, 415, 439
usage guidelines.....	116	overview.....	401
bandwidth-percent statement.....	185	ping CCC LSPs.....	171
usage guidelines.....	128, 131	point-to-multipoint LSPs.....	424, 437, 438
BFD		traffic policing.....	401
ECMP paths.....	353	Cisco HDLC circuits.....	406
fast reroute.....	167	Cisco HDLC encapsulation	
LDP LSPs.....	351, 353	Layer 2 switching cross-connect.....	419
revert timer.....	234	class types	
RSVP LSPs.....	167, 169	bandwidth subscription.....	124
bfd-liveness-detection statement		class-of-service statement	
LDP LSPs.....	369	bypass LSPs.....	304
usage guidelines.....	351	usage guidelines.....	285
RSVP LSPs.....	186	ingress routers.....	187
usage guidelines.....	167	usage guidelines.....	133
BGP		signaled LSPs.....	187
destinations.....	39	usage guidelines.....	99
binding (tracing flag).....	394	static LSPs.....	187
braces, in configuration statements.....	xxxvi	usage guidelines.....	136
brackets		colored links.....	78, 96, 178
angle, in syntax descriptions.....	xxxv	comments, in configuration statements.....	xxxvi
square, in configuration statements.....	xxxvi	connection (tracing flag).....	248
branch LSPs.....	142	connection-detail (tracing flag).....	248
bypass LSPs.....	283	connections statement.....	428
administrative groups.....	284	complete hierarchy under.....	10
bandwidth.....	284	TCC	
bandwidth subscription.....	288	usage guidelines.....	422
class-of-service.....	285	Constrained Shortest Path First algorithm <i>See</i> CSPF	
CSPF, disabling.....	287	algorithm	

constrained-path LSPs	
computation	
CSPF algorithm	29
disabling	95, 216
overview	29
example configurations	57
overview	29
scope	29
control-channel statement	472
usage guidelines	453
conventions	
text and syntax	xxxv
CoS	116
Differentiated Services	125
CoS requests using RSVP	256
CoS values	99
cross-connect, circuit <i>See</i> CCC	
cspf (tracing flag)	248
CSPF algorithm	
fate-sharing	56
offline path computation	6, 32
online path computation	29
disabling	95, 216
overview	6
cspf-link (tracing flag)	248
cspf-node (tracing flag)	248
curly braces, in configuration statements	xxxvi
customer support	xxxvii
contacting JTAC	xxxvii

D

damping	
LSP transitions	107
dead-interval statement	472
usage guidelines	458
deaggregate statement	370
usage guidelines	349
default route	
MPLS	60
default-route statement	188
usage guidelines	60
description statement	
MPLS	188
usage guidelines	78
detail (tracing flag modifier)	
LDP	395
LMP	486
RSVP	325
detours <i>See</i> fast reroute	
Differentiated Services	
bandwidth model	116
extended MAM	116
interface bandwidth constraints	125
LSPs	113

MAM	116
RDM	116
diffserv-te statement	189
usage guidelines	115
disable (tracing flag modifier)	395
disable option to traceoptions statement	
LDP	394
LMP	485
RSVP	324
disable statement	
GMPLS	473
usage guidelines	457
LDP	371
usage guidelines	338
link protection	305
usage guidelines	281
MPLS	190
usage guidelines	95
OSPF	473
usage guidelines	458
RSVP	305
usage guidelines	273
RSVP graceful restart	305
usage guidelines	290
discard statement	
MPLS	190
usage guidelines	136
discovery messages, LDP	333
distinct reservations	261
documentation set	
comments on	xxxvi
double-push statement	191
usage guidelines	133
DSCP	
MPLS-tagged packets	165
dynamic LSP metric	82

E

ECMP paths	
BFD	353
ecmp statement	372
usage guidelines	353
egress policy, loopback address	348
egress routers	
example configuration	138
overview	28
signaled LSPs	74
static LSPs	136, 205
egress-policy statement	372
usage guidelines	348
empty paths	224
encapsulation	
CCC	405
TCC	419

encapsulation statement	
Layer 2 switching cross-connect.....	429
usage guidelines.....	406
LSP tunnel cross-connect.....	429
usage guidelines.....	414
TCC.....	429
usage guidelines.....	419
encoding-type statement.....	191
usage guidelines.....	460
error (tracing flag)	
LDP.....	394
MPLS.....	248
RSVP.....	324
Ethernet extended VLAN TCC, ARP	
configuration.....	422
Ethernet TCC	
ARP configuration.....	422
ethernet-ccc encapsulation type.....	407
event (tracing flag)	
LDP.....	394
RSVP.....	324
exclude statement.....	192
administrative groups	
usage guidelines.....	96
fast reroute	
usage guidelines.....	78
EXP and IP precedence bits.....	166
EXP bits.....	26, 99, 100, 158
DSCP values.....	165
rewrite.....	166
EXP rewrite rule.....	100
expand-loose-hop statement.....	193
usage guidelines.....	152
experimental bits <i>See</i> EXP bits	
Explicit Null label.....	25
Explicit Route object.....	7
explicit routes.....	6
explicit senders, RSVP.....	261
explicit-null statement	
LDP.....	373
usage guidelines.....	362
MPLS.....	194
usage guidelines.....	149
RSVP.....	194
usage guidelines.....	295
explicit-path LSPs	
computation, disabling.....	95, 216
configuring.....	140
example configurations.....	57
overview.....	29
scope.....	29
export statement.....	373
usage guidelines.....	345
extended MAM.....	116, 184

F

facility backup.....	266
failed LSPs	
fast reroute.....	43, 78, 183, 196
standby secondary paths.....	43
failure-action statement	
LDP LSPs.....	374
usage guidelines.....	353
RSVP LSPs.....	195
usage guidelines.....	169
fast reroute.....	266
BFD.....	167
configuring.....	196
multiclass LSPs.....	131
overview.....	43, 78
path optimization.....	79
path optimization overview.....	47
PFE fast reroute.....	78, 291
soft preemption.....	87
traffic-engineered LSPs.....	128
fast-reroute optimize-timer statement.....	306
usage guidelines.....	79
fast-reroute statement.....	196
fate-sharing	
CSPF algorithm.....	56
example configuration.....	56
overview.....	33
signaled LSPs.....	54, 197
fate-sharing statement.....	197
usage guidelines.....	54
FECs.....	329
FF (reservation style).....	261
filtering received labels.....	343, 380
fixed-filter reservation style.....	261
font conventions.....	xxxv
forwarding <i>See</i> MPLS	
forwarding adjacency	
configuration.....	467
LSP.....	468
OSPF configuration.....	469
peer router address.....	468
RSVP configuration.....	468
forwarding equivalence classes <i>See</i> FECs	
forwarding next hop.....	41
Frame Relay circuits.....	409, 414
Frame Relay encapsulation	
Layer 2 TCC.....	420
from statement	
MPLS.....	198
usage guidelines.....	73

G

Generalized MPLS *See* GMPLS

- GMPLS
 - Admin Status object.....462
 - configuration.....449
 - configuration statements.....471
 - graceful deletion timeout interval.....463
 - graceful LSP teardown.....462
 - non-packet LSPs.....462
 - overview.....443
 - permanent LSP deletion.....463
 - RSVP LSP hierarchy.....465
 - temporary LSP deletion.....462
- gpid statement.....199
 - usage guidelines.....460
- graceful deletion timeout interval.....463
- graceful restart
 - LDP.....375
 - point-to-multipoint LSPs.....144
 - RSVP.....307
- graceful teardown, GMPLS LSPs.....462
- graceful-deletion-timeout statement.....306
 - usage guidelines.....463
- graceful-restart (tracing flag).....248
- graceful-restart statement
 - LDP.....375
 - usage guidelines.....342
 - RSVP.....307
 - usage guidelines.....290
- GRE tunnels.....62
- groups
 - administrative.....78, 96, 178, 192, 201, 203
- H**
- headers, MPLS and IPv4.....166
- hello interval
 - LDP.....338, 376
 - RSVP.....278, 308
- hello messages.....333
- hello packets
 - RSVP.....259
- hello-dead-interval statement.....474
 - usage guidelines.....455
- hello-interval statement
 - LDP.....376
 - usage guidelines.....338
 - LMP.....475
 - usage guidelines.....455
 - OSPF.....475
 - usage guidelines.....458
 - RSVP.....308
 - usage guidelines.....278
- hello-packets (tracing flag)
 - LMP.....485
- helper-disable statement
 - LDP.....377
 - usage guidelines.....342
 - RSVP
 - usage guidelines.....290
- hold priority.....102
- hold time
 - LDP.....339, 378
 - signaled LSPs.....107, 178
- hold-time statement
 - LDP.....378
 - usage guidelines.....339
- holddown-interval statement.....377
 - usage guidelines.....354
- hop-limit statement.....200, 308
 - usage guidelines.....105, 285
- host routes.....39, 80
- hot-standby state.....106
- I**
- icmp-tunneling statement.....200
 - usage guidelines.....67
- icons defined, notice.....xxxiv
- IEEE 802.p rewrite rule.....100
- ignore-lsp-metrics statement.....379
 - usage guidelines.....363
- IGP
 - advertising LSPs.....37
 - destinations.....41
 - shortcuts
 - enabling.....35
 - LSP metrics.....82
 - overview.....34
 - qualified LSPs.....36
 - routing tables.....37
 - uses of.....36
- igp-synchronization statement.....379
 - usage guidelines.....365
- Implicit Null label.....25
- import statement
 - LDP.....380
 - usage guidelines.....343
- include statement
 - fast reroute
 - usage guidelines.....78
- include-all statement.....201
 - administrative groups
 - usage guidelines.....96
- include-any statement.....203
 - administrative groups
 - usage guidelines.....96
- inet option to static-path statement.....133, 241
- inet.0 routing table
 - IGP shortcuts.....37
 - MPLS.....41

inet.3 routing table	
IGP shortcuts.....	37
MPLS.....	41
routes, installing.....	80
information distribution, traffic engineering.....	5
ingress routers	
address configuration.....	73, 198
configuring for static LSPs.....	133
example configurations.....	57, 135
overview.....	28
path connection retry information.....	81, 233
ingress-policy statement.....	380
usage guidelines.....	354
init (tracing flag)	
LMP.....	485
initialization (tracing flag).....	394
install statement	
MPLS.....	204
usage guidelines.....	80
Integrity object.....	257
inter-area traffic engineering.....	152
interface (from operator, LDP).....	343
interface statement	
LDP.....	381
usage guidelines.....	338
LMP.....	476
usage guidelines.....	449
RSVP.....	309
usage guidelines.....	273
static LSPs.....	205
usage guidelines.....	136
interface-switch statement.....	435
Layer 2 switching cross-connects	
usage guidelines.....	410
usage guidelines.....	410
interfaces	
aggregated.....	38
interior gateway protocol <i>See</i> IGP	
intermediate routers	
configuring for static LSPs.....	136, 205
example configurations.....	137
intraregion LSPs.....	36
IP packets over aggregated interfaces.....	38
IPv4 Explicit Null label.....	25
IPv6	
Implicit Null label.....	25
tunneling over MPLS.....	64
VPNs.....	67
ipv6-tunneling statement.....	206
usage guidelines.....	67

K

keep multiplier, RSVP.....	292, 310
keep-multiplier statement.....	310
usage guidelines.....	292

keepalive-interval statement.....	382
usage guidelines.....	341
keepalive-timeout statement.....	382
usage guidelines.....	341
keepalives	
interval.....	341, 382
timeout.....	341, 382

L

l2-smart-policy statement.....	383
usage guidelines.....	350
label (tracing flag).....	394
Label Distribution Protocol <i>See</i> LDP	
label filtering.....	343, 380
Label object.....	7
Label Request object.....	7
label-map statement.....	206
usage guidelines.....	136
label-switched paths <i>See</i> LSPs	
label-switched-path statement	
GMPLS.....	476
MPLS.....	207
usage guidelines.....	70
MPLS with RSVP.....	207
usage guidelines.....	274
usage guidelines.....	468
label-withdrawal-delay statement.....	383
usage guidelines.....	365
labels	
allocation.....	26
numerical ranges.....	25
operations.....	27, 332
overview.....	21, 22
properties.....	133
reserved labels.....	25
stacks.....	26
swapping.....	4
values.....	25
Layer 2 switching	
MPLS.....	423
TCC.....	422
Layer 2 switching cross-connect	
CCC connections.....	410
CCC encapsulation.....	406
configuration.....	405
configuring MPLS.....	411
example configuration.....	411
overview.....	401
TCC encapsulation.....	419
Layer 2 VPNs	
aggregated Ethernet.....	408
LDP	
BFD.....	353
carrier-of-carriers VPNs.....	361
configuration statements.....	367

- configuring.....381, 384
- disabling.....338, 371
- egress policy.....348
- enabling.....338
- example configuration
 - received label filtering.....345
 - tracing.....359
- Explicit Null label.....295, 362
- FEC policers.....349
- graceful restart.....334, 342, 375
- hello interval.....338, 376
- hello messages.....333
- hold time.....339, 378
- Implicit Null label.....295, 362
- Internet drafts.....330
- JUNOS implementation.....330
- keepalive
 - interval.....341, 382
 - timeout.....341, 382
- label operations.....332
- message types.....333
- metrics.....361
- minimum configuration.....338
- multiple instances.....361
- operations.....331
- overview.....329
- policy filters.....380
- received label filtering.....343, 380
- RFC.....330
- route preferences.....341, 391
- standards.....330
- synchronization with the IGP.....364
- targeted hello messages.....333
- timer.....338, 376
- tracing operation of.....358, 394
- tunneling through RSVP LSPs.....209, 331, 362
- ultimate-hop popping.....295, 362, 373
- ldp statement.....384
 - complete hierarchy under.....10
 - usage guidelines.....338
- ldp-synchronization statement.....384
 - usage guidelines.....364
- ldp-tunneling statement.....209
 - usage guidelines.....362
- least-fill statement.....209
 - usage guidelines.....83
- least-fill tie-breaking rule.....32, 83, 231
- link attributes considered by CSPF algorithm.....29
- link coloring.....78, 96, 178
- link hello messages, LDP.....376
- link protection.....266, 281
 - bypass LSPs
 - administrative groups.....284
 - multiple bypass LSPs.....286
 - RSVP.....265
 - soft preemption.....87
- link-layer protocols.....24
- link-management statement.....477
 - complete hierarchy under.....12
 - usage guidelines.....449
- link-protection statement
 - MPLS.....210
 - usage guidelines.....144, 281
 - RSVP.....311
 - usage guidelines.....282
- LMP
 - peer network device configuration.....452
 - tracing protocol operations.....485
 - tracing protocol traffic.....459
 - traffic engineering links.....449
- lmp (tracing flag).....324
- lmp-control-channel statement.....477
 - usage guidelines.....454
- lmp-protocol statement.....478
- load balancing
 - IP header.....50
 - LSPs.....84
 - MPLS labels.....50
 - without CSPF.....86
- load-balance statement.....312
 - usage guidelines.....291
- local-address statement
 - link management.....478
 - usage guidelines.....451
 - usage guidelines.....467
- log-updown statement
 - LDP.....385
 - usage guidelines.....364
 - MPLS.....211
 - usage guidelines.....157
- logical-router *See* logical-system
- logical-routers *See* logical-systems
- loopback address, egress policy.....348
- loose explicit routes.....6, 140
- LSP graceful teardown.....462
- LSP metric.....152
- lsp-attributes statement.....212
 - usage guidelines.....460
- lsp-ping-interval statement
 - LDP LSPs.....387
 - RSVP LSPs.....220
- lsp-switch statement.....436
 - usage guidelines.....417
- lsping (tracing flag).....249
- LSPs
 - adaptive rerouting.....101, 173
 - administrative groups
 - admin-group statement.....176
 - admin-groups statement.....178
 - configuring.....96
 - exclude statement.....192
 - fast reroute.....78

- advertising in IGP.....37
- associating addresses.....80, 204
- attributes considered by CSPF algorithm.....29
- automatic bandwidth allocation.....180
- automatic policers.....163
- bandwidth
 - maximum bounds.....90
 - minimum bounds.....90
- BFD configuration.....167
- bypass.....281
- configuration statements.....70, 207
- configuring.....127
- constrained-path *See* constrained-path LSPs
- CoS values.....99
- creating.....70
- damping LSP transitions.....107
- description, textual.....78
- differentiated service aware.....113
- egress routers.....74, 136, 138, 205
- example configurations.....57
- explicit-path *See* explicit-path LSPs
- failure of.....43
- fast reroute.....43, 78, 183, 196
- fate-sharing.....33, 54, 197
- forwarding next hops
 - selecting.....41
- hold time.....107, 178
- host routes.....39
- IGP shortcuts.....34
- ingress routers.....73, 198
- intermediate routers.....136, 205
- intraregion LSPs.....36
- load balancing.....84
 - without CSPF.....86
- metrics.....82, 213
- MPLS routers, configuring.....60
- multiple bypass.....267
- named paths.....52, 224
- OAM configuration.....167
- overview.....4, 24
- packet traversal.....5, 28
- path
 - bandwidth.....106
 - calculation.....3
 - connection retry information.....81, 233
 - length.....105, 200
- pings.....170
 - ping interval, LDP.....352
 - ping interval, RSVP.....168
- policing.....160
- preemption.....102, 229
- preference levels.....98, 227
- primary.....75, 228
- priorities.....102, 229
- recording routes.....98
- reoptimization.....103, 221, 222

- router functions.....28
- routing options.....7
- RSVP-signaled.....29
- scope of.....29
- secondary.....75, 237
- signaled *See* signaled LSPs
- soft preemption.....87
- standby secondary paths.....43
- standby state.....106, 240
- static *See* static LSPs
- stitching cross-connects.....401, 417
- text description.....188
- tie-breaking rules.....32, 83, 231
- traffic engineering, configuring.....150
- TTL decrementing, disabling.....86, 217, 218
- tunnel cross-connects, MTU.....413
- tunneling through RSVP LSPs.....209, 331, 362

M

- MAM.....116, 184
- manuals
 - comments on.....xxvii
- max-bypasses statement.....313
 - usage guidelines.....286
- maximum-bandwidth statement.....212
 - usage guidelines.....90
- maximum-neighbor-recovery-time statement.....385
 - usage guidelines.....343
- maximum-recovery-time statement.....385
- MD5 authentication.....278
- messages
 - LDP message types.....333
 - MPLS system log.....157, 211
 - Resv, RSVP.....260
 - ResvConfirm, RSVP.....261
 - ResvErr, RSVP.....261
 - ResvTear, RSVP.....260
 - RSVP message types.....259
 - RSVP PathErr.....60
 - RSVP refresh.....292
- metric statement
 - MPLS.....213
 - usage guidelines.....82
- metrics
 - dynamic LSP metric.....82
 - LDP tracking IGP.....361
 - static LSP metric.....82, 213
- minimum-bandwidth statement.....213
 - usage guidelines.....90
- monitor-bandwidth statement.....214
 - usage guidelines.....93
- most-fill statement.....214
 - usage guidelines.....83
- most-fill tie-breaking rule.....32, 83, 214

- MPLS.....3
 - aggregated interfaces.....38
 - automatic bandwidth allocation.....180
 - backbones, packet traversal.....5, 28
 - BFD.....167, 169
 - BGP destinations.....39
 - configuring.....51
 - CoS values.....99
 - default route.....60
 - DSCP and EXP values.....165
 - exception monitoring.....60
 - EXP bits.....26, 99, 100, 158
 - Explicit Null label.....149
 - fast reroute.....43, 78, 183, 196
 - firewall filter.....158
 - GRE tunnels.....62
 - IGP and BGP destinations.....41
 - Implicit Null label.....149
 - IPv4 packet headers.....166
 - IPv6.....64
 - label range.....25
 - Layer 2 switching TCC.....423
 - link-layer protocols supported.....24
 - load balancing.....50
 - LSP tunnel cross-connects
 - MTU.....413
 - LSPs *See* LSPs
 - OAM.....167
 - overview.....21, 22
 - packets over aggregated interfaces.....38
 - ping
 - Layer 3 VPNs.....171
 - LSP end points.....171
 - LSPs.....170
 - routing tables.....41
 - RSVP *See* RSVP
 - signaled LSPs *See* signaled LSPs
 - SNMP traps.....157, 211
 - soft preemption.....87
 - standards supported.....22
 - standby secondary paths.....43
 - static.....133, 205
 - static LSPs *See* static LSPs
 - statistics.....89
 - statistics output.....156
 - system log messages.....157, 211
 - tracing protocol operations.....171, 248
 - traffic engineering.....152
 - overview.....24
 - traffic protection.....43
 - traffic statistics.....156, 242
 - tunneling
 - CCC connection.....415, 439
 - CCC encapsulation.....414
 - example configurations.....416
 - IPv6.....64
 - overview.....401, 413
 - ultimate-hop popping.....149, 194
 - See also* LDP, LSPs, RSVP, traffic engineering
 - database
 - mpls statement
 - Layer 2 switching cross-connect.....214
 - usage guidelines.....411
 - MPLS.....214
 - complete hierarchy under.....13
 - usage guidelines.....51
 - mpls.0 routing table.....41
 - MTU signaling, in RSVP.....263
 - mtu-signaling statement.....215
 - usage guidelines.....293
 - multicast
 - RPF check policy.....145
 - multiclass LSPs
 - bandwidth subscription.....124
 - configuring.....129
 - fast reroute.....131
 - multiple bypass LSPs.....267, 283, 284, 286
 - Multiple Push (label operation).....27
- N**
- named paths
 - empty paths.....224
 - example configuration.....54
 - overview.....52
- neighbor (from operator, LDP).....343
- next hop (from operator, LDP).....343
- next hops
 - selecting.....41
- next-hop bypass LSP.....267
- next-hop statement.....215
- next-next-hop bypass LSP.....267
- no-adjacency-down-notification statement.....313
 - configuration guidelines.....295
- no-aggregate statement.....300
 - usage guidelines.....275
- no-cspf statement.....216, 314
 - usage guidelines.....95, 287
- no-decrement-ttl statement.....217
- no-forwarding statement.....386
 - usage guidelines.....361
- no-install-to-address statement.....218
 - usage guidelines.....74
- no-neighbor-down-notification statement.....314
 - usage guidelines.....295
- no-node-id-subobject statement.....315
 - usage guidelines.....281
- no-node-protection statement
 - usage guidelines.....287
- no-p2mp-sublsp statement.....315
 - usage guidelines.....147

no-propagate-ttl statement.....	218
no-record statement.....	232
usage guidelines.....	98
no-reliable statement.....	321
usage guidelines.....	275
no-trap statement.....	219
usage guidelines.....	157
no-world-readable option to traceoptions statement	
LDP.....	395
LMP.....	486
MPLS.....	249
RSVP.....	325
node protection.....	266, 267
soft preemption.....	87
node-link-protection statement.....	316
usage guidelines.....	281
notice icons defined.....	xxxiv
notification (tracing flag).....	395
notification messages	
LDP.....	334
O	
oam statement	
LDP LSPs.....	387
usage guidelines.....	351
RSVP LSPs.....	220
usage guidelines.....	167
OAM, MPLS.....	167
offline path calculation.....	6, 32
one-to-one backup.....	266
operations on labels.....	27
optimize-aggressive statement.....	221
usage guidelines.....	103
optimize-timer statement	
bypass LSPs.....	316
usage guidelines.....	287
MPLS.....	222
usage guidelines.....	103
optimizing LSPs.....	103, 221, 222
OSPF	
hello interval.....	475
inter-area traffic engineering.....	152
link-state	
advertisements.....	483
LSP metric advertisement.....	152
router dead interval.....	472
outgoing MTU value in RSVP	
determining.....	265
P	
p2mp statement.....	223
usage guidelines.....	141, 142
p2mp-lsp-next-hop statement.....	223
usage guidelines.....	138
p2mp-receive-switch statement.....	437
usage guidelines.....	425
p2mp-transmit-switch statement.....	438
usage guidelines.....	425
packet forwarding component	
traffic engineering.....	4
packet headers, MPLS and IPv4.....	166
packet traversal on LSPs.....	5, 28
packets (tracing flag)	
LDP.....	395
LMP.....	485
RSVP.....	324
parentheses, in syntax descriptions.....	xxxvi
parser (tracing flag)	
LMP.....	485
passive statement.....	479
usage guidelines.....	456
path	
bandwidth, LSP.....	106
calculation	
constrained-path computation.....	95, 216
CSPF algorithm.....	6, 29
offline path computation.....	6, 32
routing options.....	7
tie-breaking rules.....	32, 83, 231
connection retry information.....	81, 233
length, LSP.....	105, 200
selection component, traffic engineering.....	6
path (tracing flag)	
LDP.....	395
RSVP.....	324
path messages, RSVP.....	260
path optimization	
fast reroute.....	79
path selection.....	77
path statement	
MPLS.....	224
RSVP.....	317
usage guidelines.....	288
path-mtu statement.....	225, 293
PathErr messages.....	60, 261
pathtear (tracing flag).....	324
PathTear messages, RSVP.....	260
peer network device configuration.....	452
peer statement	
LMP.....	479
usage guidelines.....	452
peer-interface statement	
OSPF.....	480
usage guidelines.....	457
RSVP.....	318
usage guidelines.....	457
usage guidelines.....	468, 469
periodic (tracing flag).....	395
periodic-traceroute statement.....	389
usage guidelines.....	353, 354

- permanent GMPLS LSP deletion.....463
 - PFE fast reroute.....78, 291
 - ping
 - Layer 3 VPNs.....171
 - LSP end point.....171
 - LSPs.....170
 - P2MP LSP.....170
 - PLP bit.....99
 - point-to-multipoint LSPs
 - automatic policers.....164
 - branch LSPs.....142
 - dynamic.....142
 - static.....143
 - CCC.....424
 - configuration.....141
 - graceful restart.....144
 - link protection.....144
 - overview.....48
 - RPF check policy.....145
 - static routes.....138
 - ultimate-hop popping.....296
 - policers
 - LDP FECs.....349
 - policing.....160, 162, 401
 - policing filter statement
 - usage guidelines.....160
 - policing statement.....225, 390
 - usage guidelines.....162, 349
 - policy filters, LDP.....380
 - Pop (label operation).....27
 - pop statement
 - MPLS.....226
 - usage guidelines.....136
 - PPP circuits
 - Layer 2 switching cross-connects.....406
 - preemption
 - LSPs.....87
 - RSVP sessions.....293
 - signaled LSPs.....102, 229
 - preemption statement.....319
 - usage guidelines.....293
 - preference levels
 - LDP routes.....341, 391
 - signaled LSPs.....98, 227
 - static LSPs.....133, 136
 - preference statement
 - LDP.....391
 - usage guidelines.....341
 - signaled LSPs.....227
 - usage guidelines.....98
 - static LSPs.....227
 - usage guidelines (egress router).....136
 - usage guidelines (ingress router).....133
 - primary LSPs.....75, 228
 - primary paths
 - revert timer.....76
 - revert timer, BFD.....234
 - selection.....77
 - primary statement
 - MPLS.....228
 - usage guidelines.....75
 - priorities
 - signaled LSPs.....102, 229
 - priority statement
 - MPLS.....229
 - usage guidelines.....102
 - RSVP.....320
 - usage guidelines.....289
 - process (tracing flag).....485
 - Push (label operation).....27
 - push statement
 - MPLS.....230
- ## R
- random statement.....231
 - usage guidelines.....83
 - random tie-breaking rule.....32, 83, 231
 - RDM.....116, 184
 - receive (tracing flag modifier)
 - LDP.....395
 - LMP.....486
 - RSVP.....325
 - received label filtering.....380
 - Record Route object.....98
 - record statement.....232
 - usage guidelines.....98
 - recording routes.....98
 - recovery-time statement.....391
 - usage guidelines.....343
 - refresh messages, RSVP.....292
 - refresh reduction, RSVP.....262
 - refresh time, RSVP.....292
 - refresh-time statement.....321
 - usage guidelines.....292
 - reject statement.....232
 - usage guidelines.....136
 - reliable statement.....321
 - usage guidelines.....275
 - remote-address statement
 - control channel management
 - usage guidelines.....454
 - LMP control channel.....481
 - LMP traffic engineering.....481
 - usage guidelines.....451, 468
 - remote-id statement
 - link management.....482
 - usage guidelines.....452
 - remote-interface-switch statement.....439
 - usage guidelines.....415

reoptimizing LSPs.....	103, 221, 222
requests, CoS.....	256
rerouting LSPs	
adaptive rerouting.....	101, 173
fast reroute.....	43, 78, 183, 196
reservation styles.....	261
reserved labels.....	25
reserving network resources <i>See</i> RSVP	
resource classes.....	78, 96
Resource Reservation Protocol <i>See</i> RSVP	
resv (tracing flag).....	324
Resv messages, RSVP.....	260
ResvConfirm messages, RSVP.....	261
ResvErr messages, RSVP.....	261
resvtear (tracing flag).....	324
ResvTear messages, RSVP.....	260
retransmission-interval statement.....	482
usage guidelines.....	456
retransmit-interval statement.....	483
usage guidelines.....	458
retry information.....	81, 233
retry-limit statement.....	233, 483
usage guidelines.....	81, 456
retry-timer statement.....	233
usage guidelines.....	81
revert-timer statement.....	234
usage guidelines.....	76
rewrite rules.....	100
IEEE 802.p and MPLS CoS.....	100
MPLS and VPNs.....	166
route (tracing flag).....	325
LDP.....	395
route preferences	
LDP.....	341, 391
signaled LSPs.....	98, 227
route-socket (tracing flag)	
LMP.....	485
Router Alert label.....	25
routers	
egress <i>See</i> egress routers	
ingress <i>See</i> ingress routers	
label operations.....	27
LSP functions.....	28
transit.....	28
routes	
recording.....	98
route preferences.....	98, 227, 341, 391
routing (tracing flag).....	485
routing options, traffic engineering.....	7
routing tables	
host routes, installing.....	80, 204
IGP shortcuts.....	37
inet.0.....	37, 41
inet.3.....	37, 41, 80
MPLS.....	41
mpls.0.....	41
rpf-check-policy statement.....	235
configuration guidelines.....	145
RRO node ID sub-object, disabling.....	281
RSVP.....	3
aggregation.....	300
authentication.....	278, 301
bandwidth	
reserving.....	323
update threshold.....	279
BFD.....	167, 169
configuration, minimum.....	273
Differentiated Services.....	125
disabling.....	273
ECMP-aware BFD.....	353
enabling.....	273
example configurations.....	274, 298
Explicit Null label.....	295
graceful restart.....	268, 307
hello interval.....	278, 308
hello packets.....	259
IGP hello packets.....	259
Implicit Null label.....	295
JUNOS implementation.....	257
keep multiplier.....	310
link protection.....	281
load balancing.....	291
message types.....	259
MPLS, configuring with RSVP.....	274
MTU signaling in.....	263
overview.....	256
PathErr messages.....	60
preemption.....	293
reservation styles.....	261
RFC draft documents.....	256
sessions.....	258, 274
signaled LSPs.....	29
signaling extensions.....	7
timers.....	292, 321
timers, hello packets.....	259
tracing protocol traffic.....	297, 324
tunneling LDP LSPs through RSVP	
LSPs.....	209, 331, 362
ultimate-hop popping.....	295, 296
unnumbered interfaces.....	280
<i>See also</i> LDP	
RSVP LSP hierarchy.....	465
configuration.....	466
overview.....	465
RSVP refresh reduction	
configuration.....	275
overview.....	262
rsvp statement.....	322
complete hierarchy under.....	17
usage guidelines.....	273
rsvp-error-hold-time statement.....	236
usage guidelines.....	60

S

- scope of LSPs.....29
- SE (reservation style).....261
- secondary
 - LSPs.....75, 106, 237
 - paths.....43
 - revert timer.....76
 - selection.....77
- secondary statement.....237
 - usage guidelines.....75
- select statement.....238
 - usage guidelines.....77
- send (tracing flag modifier)
 - LDP.....395
 - LMP.....486
 - RSVP.....325
- server (tracing flag).....485
- session messages, LDP.....334
- session statement.....392
 - usage guidelines.....363
- sessions, RSVP.....258, 274
- setup priority, signaled LSPs.....102
- shared explicit reservation style.....261
- shared reservations.....261
- show (tracing flag)
 - LMP.....485
- signal-bandwidth statement.....238
 - usage guidelines.....460
- signaled LSPs.....224
 - adaptive rerouting.....101, 173
 - administrative groups
 - admin-group statement.....176
 - admin-groups statement.....178
 - configuring.....96
 - exclude statement.....192
 - fast reroute.....78
 - associating addresses.....80, 204
 - configuration statements.....70, 207
 - constrained-path computation
 - disabling.....95, 216
 - CoS values.....99
 - creating.....70
 - damping LSP transitions.....107
 - egress router address.....74, 247
 - example configurations.....57
 - fast reroute.....78, 183, 196
 - fate-sharing.....54, 197
 - hold time.....107, 178
 - ingress router address.....73, 198
 - load balancing without CSPF.....86
 - metrics.....82, 213
 - MPLS routers, configuring.....60
 - named paths.....52
 - overview.....51, 69
 - path
 - bandwidth.....106
 - connection retry information.....81, 233
 - length.....105, 200
 - preemption.....102, 229
 - preference levels.....98, 227
 - primary.....75, 228
 - priorities.....102, 229
 - recording routes.....98
 - reoptimization.....103, 221, 222
 - RSVP *See* RSVP
 - secondary.....75, 237
 - standby state.....106, 240
 - tie-breaking rules.....32, 85, 231
 - TTL decrementing.....86, 217, 218
- signaling component, traffic engineering.....7
- signaling extensions, RSVP.....7
- size option to traceoptions statement
 - LMP.....486
- smart-optimize-timer statement.....239
 - usage guidelines.....105
- SNMP traps
 - MPLS.....157, 211
- soft-preemption statement
 - MPLS.....239
 - usage guidelines.....87
 - RSVP.....322
 - usage guidelines.....87
- special labels.....25
- stacked labels.....26
- standby secondary paths.....43
- standby state, signaled LSPs.....106, 240
- standby statement.....240
 - usage guidelines.....106
- state (tracing flag)
 - LDP.....395
 - LMP.....486
 - MPLS.....249
 - RSVP.....325
- static LSPs
 - configuring.....133
 - egress routers.....136, 138, 205
 - ingress routers.....133
 - intermediate routers.....136, 205
 - overview.....29
 - static LSP metric.....82, 213
- static MPLS.....133
- static routes
 - point-to-multipoint LSPs.....138
- static-path statement.....241
 - usage guidelines.....133
- statistics
 - MPLS traffic.....156, 242
 - output file.....156
- statistics statement.....242
 - usage guidelines.....156

strict explicit routes.....	6, 140
strict-targeted-hellos statement.....	392
usage guidelines.....	340
subscribing to bandwidth.....	323
subscription statement.....	323
usage guidelines.....	124
summary LSA.....	152
support, technical <i>See</i> technical support	
Swap (label operation).....	27
Swap and Push (label operation).....	27
swap statement	
MPLS.....	243
usage guidelines.....	136
swap-push statement	
MPLS.....	244
switching-type statement.....	245
usage guidelines.....	460
syntax conventions.....	xxxv
system log messages	
MPLS.....	157, 211

T

targeted hello messages.....	333
targeted hello messages, LDP.....	376
targeted-hello statement.....	393
usage guidelines.....	339, 340
TCC	
configuration.....	418
connections.....	422
encapsulation.....	419
graceful restart	
configuration.....	424
overview.....	403
Layer 2 switching.....	419
overview.....	402
te-class-matrix statement.....	246
usage guidelines.....	117
te-link statement.....	484
LMP traffic engineering link	
usage guidelines.....	450
traffic engineering link associated with peer	
usage guidelines.....	456
usage guidelines.....	467
technical support	
contacting JTAC.....	xxxvii
temporary GMPLS LSP deletion.....	462
tie-breaking rules, path calculation.....	32, 83, 231
timer (tracing flag)	
MPLS.....	249
timer, LDP.....	338, 376
timers	
RSVP.....	292, 321
to statement	
MPLS.....	247
usage guidelines.....	74

traceoptions statement	
LDP.....	394
usage guidelines.....	358
LMP.....	485
usage guidelines.....	459
MPLS.....	248
usage guidelines.....	171
RSVP.....	324
usage guidelines.....	297
tracing flag modifiers	
detail	
LDP.....	395
LMP.....	486
RSVP.....	325
disable.....	395
receive	
LDP.....	395
LMP.....	486
RSVP.....	325
send	
LDP.....	395
LMP.....	486
RSVP.....	325
tracing flags	
address.....	394
all.....	248
LMP.....	485
RSVP.....	324
binding.....	394
connection.....	248
connection-detail.....	248
cspf.....	248
cspf-link.....	248
cspf-node.....	248
error	
LDP.....	394
MPLS.....	248
RSVP.....	324
event	
LDP.....	394
RSVP.....	324
graceful-restart.....	248
hello-packets	
LMP.....	485
init	
LMP.....	485
initialization.....	394
label.....	394
lmp.....	324
lsping.....	249
notification.....	395
packets	
LDP.....	395
LMP.....	485
RSVP.....	324

- parse
 - LMP.....485
 - path
 - LDP.....395
 - RSVP.....324
 - pathtear.....324
 - periodic395
 - process.....485
 - resv.....324
 - resvtear.....324
 - route.....325
 - LDP.....395
 - route-socket
 - LMP.....485
 - routing.....485
 - server.....485
 - show
 - LMP.....485
 - state
 - LDP.....395
 - LMP.....486
 - MPLS.....249
 - RSVP.....325
 - timer
 - MPLS.....249
 - tracing operations
 - LDP.....358, 394
 - LMP.....459, 485
 - MPLS.....171, 248
 - RSVP.....297, 324
 - track-igp-metric statement.....396
 - usage guidelines.....361
 - traffic
 - policing.....401
 - protection, MPLS.....43
 - statistics.....156, 242
 - traffic engineering
 - BGP destinations.....39
 - fate-sharing.....33
 - IGP and BGP destinations.....41
 - IGP shortcuts.....34
 - information distribution component.....5
 - inter-area, OSPF.....152
 - links.....449
 - LSP metric advertisement.....152
 - overview.....3, 24
 - packet-forwarding component.....4
 - path-selection component.....6
 - routing options.....7
 - signaling component.....7
 - traffic engineering database accuracy.....60
 - traffic engineering database.....29
 - accuracy.....60
 - bandwidth update threshold.....279
 - traffic-engineered LSPs
 - fast reroute.....128
 - traffic-engineering statement
 - bgp-igp option.....150
 - bgp-igp-both-ribs option.....151
 - MPLS.....250
 - usage guidelines.....150
 - mpls-forwarding option.....151
 - traffic-statistics statement.....397
 - usage guidelines.....355
 - transit routers.....28
 - transit-delay statement.....487
 - usage guidelines.....458
 - transitions
 - advertising.....107, 178
 - damping.....107
 - translational cross-connect *See* TCC
 - transport-address statement.....398
 - usage guidelines.....347
 - traps, SNMP *See* SNMP traps
 - triple-push statement.....251
 - usage guidelines.....133
 - TTL decrementing
 - disabling.....86, 217, 218
 - tunnel-services statement.....326
 - usage guidelines.....296
 - tunneling, MPLS
 - CCC encapsulation.....414
 - example configurations.....416
 - overview.....401, 413
 - RSVP LSPs.....209, 331, 362
 - RSVP LSPs, heterogeneous networks.....363
- ## U
- ultimate-hop popping.....149
 - point-to-multipoint LSPs.....296
 - unnumbered interfaces, RSVP.....280
 - unstable LSPs
 - fate-sharing *See* fate-sharing
 - update-threshold statement.....326
 - usage guidelines.....279
- ## W
- wildcard filter (WF) reservation style.....261
 - wildcard senders, RSVP.....261
 - world-readable option to statistics statement
 - MPLS.....242
 - world-readable option to traceoptions statement
 - LDP.....395
 - LMP.....486
 - MPLS.....249
 - RSVP.....325

Index of Statements and Commands

A

adaptive statement.....	173
address statement	
LMP.....	471
adjust-interval statement.....	174
adjust-threshold statement.....	174
adjust-threshold-overflow-limit statement.....	175
admin-down statement.....	175
admin-group statement	
bypass LSPs.....	299
LSPs.....	177
MPLS interfaces.....	176
admin-groups statement.....	178
advertisement-hold-time statement.....	178
aggregate statement	
RSVP.....	300
allow-fragmentation statement.....	179
allow-subnet-mismatch statement.....	367
associate-backup-pe-groups statement.....	179
authentication-key statement	
LDP.....	368
RSVP.....	301
auto-bandwidth statement.....	180
auto-policing statement.....	181

B

backup-pe-group statement.....	182
bandwidth statement	
fast reroute.....	183
link protection.....	302
multiclass LSPs.....	183
RSVP.....	302
signaled LSPs.....	183
bandwidth-model statement.....	184
bandwidth-percent statement.....	185
bfd-liveness-detection statement	
LDP LSPs.....	369
RSVP LSPs.....	186
bypass statement.....	303

C

class-of-service statement	
bypass LSPs.....	304
ingress routers.....	187
signaled LSPs.....	187
static LSPs.....	187
connections statement.....	428
control-channel statement.....	472

D

dead-interval statement.....	472
deaggregate statement.....	370
default-route statement.....	188
description statement	
MPLS.....	188
diffserv-te statement.....	189
disable statement	
GMPLS.....	473
LDP.....	371
link protection.....	305
MPLS.....	190
OSPF.....	473
RSVP.....	305
RSVP graceful restart.....	305
discard statement	
MPLS.....	190
double-push statement.....	191

E

ecmp statement.....	372
egress-policy statement.....	372
encapsulation statement	
Layer 2 switching cross-connect.....	429
LSP tunnel cross-connect.....	429
TCC.....	429
encoding-type statement.....	191
exclude statement.....	192
expand-loose-hop statement.....	193
explicit-null statement	
LDP.....	373
MPLS.....	194
RSVP.....	194

export statement.....373

F

failure-action statement
 LDP LSPs.....374
 RSVP LSPs.....195
 fast-reroute optimize-timer statement.....306
 fast-reroute statement.....196
 fate-sharing statement.....197
 from statement
 MPLS.....198

G

gpid statement.....199
 graceful-deletion-timeout statement.....306
 graceful-restart statement
 LDP.....375
 RSVP.....307

H

hello-dead-interval statement.....474
 hello-interval statement
 LDP.....376
 LMP.....475
 OSPF.....475
 RSVP.....308
 helper-disable statement
 LDP.....377
 hold-time statement
 LDP.....378
 holddown-interval statement.....377
 hop-limit statement.....200, 308

I

icmp-tunneling statement.....200
 ignore-lsp-metrics statement.....379
 igp-synchronization statement.....379
 import statement
 LDP.....380
 include-all statement.....201
 include-any statement.....203
 ingress-policy statement.....380
 install statement
 MPLS.....204
 interface statement
 LDP.....381
 LMP.....476
 RSVP.....309
 static LSPs.....205
 interface-switch statement.....435
 ipv6-tunneling statement.....206

K

keep-multiplier statement.....310
 keepalive-interval statement.....382
 keepalive-timeout statement.....382

L

l2-smart-policy statement.....383
 label-map statement.....206
 label-switched-path statement
 GMPLS.....476
 MPLS.....207
 label-withdrawal-delay statement.....383
 ldp statement.....384
 ldp-synchronization statement.....384
 ldp-tunneling statement.....209
 least-fill statement.....209
 link-management statement.....477
 link-protection statement
 MPLS.....210
 RSVP.....311
 lmp-control-channel statement.....477
 lmp-protocol statement.....478
 load-balance statement.....312
 local-address statement
 link management.....478
 log-updown statement
 LDP.....385
 MPLS.....211
 lsp-attributes statement.....212
 lsp-ping-interval statement
 LDP LSPs.....387
 RSVP LSPs.....220
 lsp-switch statement.....436

M

max-bypasses statement.....313
 maximum-bandwidth statement.....212
 maximum-neighbor-recovery-time statement.....385
 maximum-recovery-time statement.....385
 metric statement
 MPLS.....213
 minimum-bandwidth statement.....213
 monitor-bandwidth statement.....214
 most-fill statement.....214
 mpls statement
 Layer 2 switching cross-connect.....214
 MPLS.....214
 mtu-signaling statement.....215

N

next-hop statement.....215
 no-adjacency-down-notification statement.....313
 no-cspf statement.....216, 314

no-decrement-ttl statement.....	217
no-forwarding statement.....	386
no-install-to-address statement.....	218
no-neighbor-down-notification statement.....	314
no-node-id-subobject statement.....	315
no-p2mp-sublsp statement.....	315
no-propagate-ttl statement.....	218
no-trap statement.....	219
node-link-protection statement.....	316

O

oam statement	
LDP LSPs.....	387
RSVP LSPs.....	220
optimize-aggressive statement.....	221
optimize-timer statement	
bypass LSPs.....	316
MPLS.....	222

P

p2mp statement.....	223
p2mp-lsp-next-hop statement.....	223
p2mp-receive-switch statement.....	437
p2mp-transmit-switch statement.....	438
passive statement.....	479
path statement	
MPLS.....	224
RSVP.....	317
path-mtu statement.....	225, 293
peer statement	
LMP.....	479
peer-interface statement	
OSPF.....	480
RSVP.....	318
periodic-traceroute statement.....	389
policing statement.....	225, 390
pop statement	
MPLS.....	226
preemption statement.....	319
preference statement	
LDP.....	391
signaled LSPs.....	227
static LSPs.....	227
primary statement	
MPLS.....	228
priority statement	
MPLS.....	229
RSVP.....	320
push statement	
MPLS.....	230

R

random statement.....	231
-----------------------	-----

record statement.....	232
recovery-time statement.....	391
refresh-time statement.....	321
reject statement.....	232
reliable statement.....	321
remote-address statement	
LMP control channel.....	481
LMP traffic engineering.....	481
remote-id statement	
link management.....	482
remote-interface-switch statement.....	439
retransmission-interval statement.....	482
retransmit-interval statement.....	483
retry-limit statement.....	233, 483
retry-timer statement.....	233
revert-timer statement.....	234
rpf-check-policy statement.....	235
rsvp statement.....	322
rsvp-error-hold-time statement.....	236

S

secondary statement.....	237
select statement.....	238
session statement.....	392
signal-bandwidth statement.....	238
smart-optimize-timer statement.....	239
soft-preemption statement	
MPLS.....	239
RSVP.....	322
standby statement.....	240
static-path statement.....	241
statistics statement.....	242
strict-targeted-hellos statement.....	392
subscription statement.....	323
swap statement	
MPLS.....	243
swap-push statement	
MPLS.....	244
switching-type statement.....	245

T

targeted-hello statement.....	393
te-class-matrix statement.....	246
te-link statement.....	484
to statement	
MPLS.....	247
traceoptions statement	
LDP.....	394
LMP.....	485
MPLS.....	248
RSVP.....	324
track-igp-metric statement.....	396
traffic-engineering statement	
MPLS.....	250

traffic-statistics statement.....	397
transit-delay statement.....	487
transport-address statement.....	398
triple-push statement.....	251
tunnel-services statement.....	326

U

update-threshold statement.....	326
---------------------------------	-----