



JUNOS® Software for EX Series Ethernet Switches, Release 10.0: MPLS

Juniper Networks, Inc.

1194 North Mathilda Avenue
Sunnyvale, California 94089

USA

408-745-2000

www.juniper.net

Revision 1
Published: 2009-11-04

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 for EX Series Ethernet Switches, Release 10.0: MPLS

Copyright © 2009, Juniper Networks, Inc.

All rights reserved. Printed in USA.

Writing:

Editing:

Illustration:

Cover Design:

Revision History

4 November 2009—Revision 1

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

END USER LICENSE AGREEMENT

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

1. **The Parties.** The parties to this Agreement are (i) Juniper Networks, Inc. (if the Customer’s principal office is located in the Americas) or Juniper Networks (Cayman) Limited (if the Customer’s principal office is located outside the Americas) (such applicable entity being referred to herein as “Juniper”), and (ii) the person or organization that originally purchased from Juniper or an authorized Juniper reseller the applicable license(s) for use of the Software (“Customer”) (collectively, the “Parties”).
2. **The Software.** In this Agreement, “Software” means the program modules and features of the Juniper or Juniper-supplied software, for which Customer has paid the applicable license or support fees to Juniper or an authorized Juniper reseller, or which was embedded by Juniper in equipment which Customer purchased from Juniper or an authorized Juniper reseller. “Software” also includes updates, upgrades and new releases of such software. “Embedded Software” means Software which Juniper has embedded in or loaded onto the Juniper equipment and any updates, upgrades, additions or replacements which are subsequently embedded in or loaded onto the equipment.
3. **License Grant.** Subject to payment of the applicable fees and the limitations and restrictions set forth herein, Juniper grants to Customer a non-exclusive and non-transferable license, without right to sublicense, to use the Software, in executable form only, subject to the following use restrictions:
 - a. Customer shall use Embedded Software solely as embedded in, and for execution on, Juniper equipment originally purchased by Customer from Juniper or an authorized Juniper reseller.
 - b. Customer shall use the Software on a single hardware chassis having a single processing unit, or as many chassis or processing units for which Customer has paid the applicable license fees; provided, however, with respect to the Steel-Belted Radius or Odyssey Access Client software only, Customer shall use such Software on a single computer containing a single physical random access memory space and containing any number of processors. Use of the Steel-Belted Radius or IMS AAA software on multiple computers or virtual machines (e.g., Solaris zones) requires multiple licenses, regardless of whether such computers or virtualizations are physically contained on a single chassis.
 - c. Product purchase documents, paper or electronic user documentation, and/or the particular licenses purchased by Customer may specify limits to Customer’s use of the Software. Such limits may restrict use to a maximum number of seats, registered endpoints, concurrent users, sessions, calls, connections, subscribers, clusters, nodes, realms, devices, links, ports or transactions, or require the purchase of separate licenses to use particular features, functionalities, services, applications, operations, or capabilities, or provide throughput, performance, configuration, bandwidth, interface, processing, temporal, or geographical limits. In addition, such limits may restrict the use of the Software to managing certain kinds of networks or require the Software to be used only in conjunction with other specific Software. Customer’s use of the Software shall be subject to all such limitations and purchase of all applicable licenses.
 - d. For any trial copy of the Software, Customer’s right to use the Software expires 30 days after download, installation or use of the Software. Customer may operate the Software after the 30-day trial period only if Customer pays for a license to do so. Customer may not extend or create an additional trial period by re-installing the Software after the 30-day trial period.
 - e. The Global Enterprise Edition of the Steel-Belted Radius software may be used by Customer only to manage access to Customer’s enterprise network. Specifically, service provider customers are expressly prohibited from using the Global Enterprise Edition of the Steel-Belted Radius software to support any commercial network access services.

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

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

6. **Confidentiality.** The Parties agree that aspects of the Software and associated documentation are the confidential property of Juniper. As such, Customer shall exercise all reasonable commercial efforts to maintain the Software and associated documentation in confidence, which at a minimum includes restricting access to the Software to Customer employees and contractors having a need to use the Software for Customer's internal business purposes.
7. **Ownership.** Juniper and Juniper's licensors, respectively, retain ownership of all right, title, and interest (including copyright) in and to the Software, associated documentation, and all copies of the Software. Nothing in this Agreement constitutes a transfer or conveyance of any right, title, or interest in the Software or associated documentation, or a sale of the Software, associated documentation, or copies of the Software.
8. **Warranty, Limitation of Liability, Disclaimer of Warranty.** The warranty applicable to the Software shall be as set forth in the warranty statement that accompanies the Software (the "Warranty Statement"). Nothing in this Agreement shall give rise to any obligation to support the Software. Support services may be purchased separately. Any such support shall be governed by a separate, written support services agreement. TO THE MAXIMUM EXTENT PERMITTED BY LAW, JUNIPER SHALL NOT BE LIABLE FOR ANY LOST PROFITS, LOSS OF DATA, OR COSTS OR PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES, OR FOR ANY SPECIAL, INDIRECT, OR CONSEQUENTIAL DAMAGES ARISING OUT OF THIS AGREEMENT, THE SOFTWARE, OR ANY JUNIPER OR JUNIPER-SUPPLIED SOFTWARE. IN NO EVENT SHALL JUNIPER BE LIABLE FOR DAMAGES ARISING FROM UNAUTHORIZED OR IMPROPER USE OF ANY JUNIPER OR JUNIPER-SUPPLIED SOFTWARE. EXCEPT AS EXPRESSLY PROVIDED IN THE WARRANTY STATEMENT TO THE EXTENT PERMITTED BY LAW, JUNIPER DISCLAIMS ANY AND ALL WARRANTIES IN AND TO THE SOFTWARE (WHETHER EXPRESS, IMPLIED, STATUTORY, OR OTHERWISE), INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NONINFRINGEMENT. IN NO EVENT DOES JUNIPER WARRANT THAT THE SOFTWARE, OR ANY EQUIPMENT OR NETWORK RUNNING THE SOFTWARE, WILL OPERATE WITHOUT ERROR OR INTERRUPTION, OR WILL BE FREE OF VULNERABILITY TO INTRUSION OR ATTACK. In no event shall Juniper's or its suppliers' or licensors' liability to Customer, whether in contract, tort (including negligence), breach of warranty, or otherwise, exceed the price paid by Customer for the Software that gave rise to the claim, or if the Software is embedded in another Juniper product, the price paid by Customer for such other product. Customer acknowledges and agrees that Juniper has set its prices and entered into this Agreement in reliance upon the disclaimers of warranty and the limitations of liability set forth herein, that the same reflect an allocation of risk between the Parties (including the risk that a contract remedy may fail of its essential purpose and cause consequential loss), and that the same form an essential basis of the bargain between the Parties.
9. **Termination.** Any breach of this Agreement or failure by Customer to pay any applicable fees due shall result in automatic termination of the license granted herein. Upon such termination, Customer shall destroy or return to Juniper all copies of the Software and related documentation in Customer's possession or control.
10. **Taxes.** All license fees payable under this agreement are exclusive of tax. Customer shall be responsible for paying Taxes arising from the purchase of the license, or importation or use of the Software. If applicable, valid exemption documentation for each taxing jurisdiction shall be provided to Juniper prior to invoicing, and Customer shall promptly notify Juniper if their exemption is revoked or modified. All payments made by Customer shall be net of any applicable withholding tax. Customer will provide reasonable assistance to Juniper in connection with such withholding taxes by promptly: providing Juniper with valid tax receipts and other required documentation showing Customer's payment of any withholding taxes; completing appropriate applications that would reduce the amount of withholding tax to be paid; and notifying and assisting Juniper in any audit or tax proceeding related to transactions hereunder. Customer shall comply with all applicable tax laws and regulations, and Customer will promptly pay or reimburse Juniper for all costs and damages related to any liability incurred by Juniper as a result of Customer's non-compliance or delay with its responsibilities herein. Customer's obligations under this Section shall survive termination or expiration of this Agreement.
11. **Export.** Customer agrees to comply with all applicable export laws and restrictions and regulations of any United States and any applicable foreign agency or authority, and not to export or re-export the Software or any direct product thereof in violation of any such restrictions, laws or regulations, or without all necessary approvals. Customer shall be liable for any such violations. The version of the Software supplied to Customer may contain encryption or other capabilities restricting Customer's ability to export the Software without an export license.
12. **Commercial Computer Software.** The Software is "commercial computer software" and is provided with restricted rights. Use, duplication, or disclosure by the United States government is subject to restrictions set forth in this Agreement and as provided in DFARS 227.7201 through 227.7202-4, FAR 12.212, FAR 27.405(b)(2), FAR 52.227-19, or FAR 52.227-14(ALT III) as applicable.
13. **Interface Information.** To the extent required by applicable law, and at Customer's written request, Juniper shall provide Customer with the interface information needed to achieve interoperability between the Software and another independently created program, on payment of applicable fee, if any. Customer shall observe strict obligations of confidentiality with respect to such information and shall use such information in compliance with any applicable terms and conditions upon which Juniper makes such information available.
14. **Third Party Software.** Any licensor of Juniper whose software is embedded in the Software and any supplier of Juniper whose products or technology are embedded in (or services are accessed by) the Software shall be a third party beneficiary with respect to this Agreement, and such licensor or vendor shall have the right to enforce this Agreement in its own name as if it were Juniper. In addition, certain third party software may be provided with the Software and is subject to the accompanying license(s), if any, of its respective owner(s). To the extent portions of the Software are distributed under and subject to open source licenses obligating Juniper to make the source code for such portions publicly available (such as the GNU General Public License ("GPL") or the GNU Library General Public License ("LGPL")), Juniper will make such source code portions (including Juniper modifications, as appropriate) available upon request for a period of up to three years from the date of distribution. Such request can be made in writing to Juniper Networks, Inc., 1194 N. Mathilda Ave., Sunnyvale, CA 94089, ATTN: General Counsel. You may obtain a copy of the GPL at <http://www.gnu.org/licenses/gpl.html>, and a copy of the LGPL at <http://www.gnu.org/licenses/lgpl.html>.
15. **Miscellaneous.** This Agreement shall be governed by the laws of the State of California without reference to its conflicts of laws principles. The provisions of the U.N. Convention for the International Sale of Goods shall not apply to this Agreement. For any disputes arising under this Agreement, the Parties hereby consent to the personal and exclusive jurisdiction of, and venue in, the state and federal courts within Santa Clara County, California. This Agreement constitutes the entire and sole agreement between Juniper and the Customer with respect to the Software, and supersedes all prior and contemporaneous

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

Table of Contents

About This Topic Collection xi

How to Use This Guide	xi
List of EX Series Guides for JUNOS Release 10.0	xi
Downloading Software	xii
Documentation Symbols Key	xiii
Documentation Feedback	xiv
Requesting Technical Support	xv
Self-Help Online Tools and Resources	xv
Opening a Case with JTAC	xv

Part 1

MPLS

Chapter 1

MPLS—Overview 3

JUNOS MPLS for EX Series Switches Overview	3
Benefits of MPLS	3
Additional Benefits of MPLS and Traffic Engineering	4
MPLS Label Switched Paths and MPLS Labels on EX Series Switches	4
MPLS Label Operations on EX Series Switches	5
Understanding JUNOS MPLS Components for EX Series Switches	6
Provider Edge Switches	7
MPLS Protocol and Label Switched Paths	7
Circuit Cross-Connect	7
Provider Switch	8
Components Required for All Switches in the MPLS Network	8
OSPF or IS-IS as a Routing Protocol	8
Traffic Engineering	8
MPLS Protocol	9
RSVP	9
Family MPLS	9
Understanding MPLS and Path Protection on EX Series Switches	10

Chapter 2

Example of MPLS Configuration 11

Example: Configuring MPLS on EX Series Switches	11
---	----

Chapter 3	Configuring MPLS	27
	Configuring Path Protection in an MPLS Network (CLI Procedure)	27
	Configuring the Primary Path	28
	Configuring the Secondary Path	28
	Configuring the Revert Timer	29
	Configuring MPLS on Provider Edge Switches (CLI Procedure)	30
	Enabling the OSPF or the IS-IS Routing Protocol on the Loopback and Core Interfaces	31
	Enabling Traffic Engineering for the Routing Protocol	31
	Configuring IP Addresses for the Loopback and Core Interfaces	31
	Enabling MPLS, Defining the Label Switched Path, and Applying MPLS to the Core Interfaces	32
	Enabling RSVP and Applying It to the Loopback and Core Interfaces	32
	Enabling Family MPLS on the Core Interfaces	33
	Configuring a Circuit Cross-Connect on a Customer Edge Interface	33
	Configuring MPLS on Provider Switches (CLI Procedure)	35
	Enabling the OSPF or the IS-IS Routing Protocol on the Loopback and Core Interfaces	36
	Enabling Traffic Engineering for the Routing Protocol	36
	Enabling MPLS and Applying MPLS to the Core Interfaces	36
	Enabling RSVP and Applying It to the Loopback and Core Interfaces	37
	Configuring IP Addresses for the Loopback and Core Interfaces	37
	Enabling Family MPLS on the Core Interfaces	37
Chapter 4	Verifying MPLS	39
	Verifying That MPLS Is Working Correctly	39
	Verifying the Physical Layer on the Switches	39
	Verifying the Routing Protocol	40
	Verifying the Core Interfaces Being Used for the MPLS Traffic	40
	Verifying RSVP	40
	Verifying the Assignment of Interfaces for MPLS Label Operations	41
	Verifying the Status of the CCC	41
Chapter 5	Configuration Statements for MPLS	43
	[edit protocols] Configuration Statement Hierarchy	43
	connections	49
	family ccc	49
	family mpls	50
	interface	50
	label-switched-path	51
	mpls	52
	path	53
	primary	54
	remote-interface-switch	55
	revert-timer	56

rsvp	56
secondary	57
standby	57
traffic-engineering	58

Chapter 6**Operational Mode Commands for MPLS** **59**

show connections	60
show route forwarding-table	64
show mpls interface	71
show rsvp session	72

About This Topic Collection

- How to Use This Guide on page xi
- List of EX Series Guides for JUNOS Release 10.0 on page xi
- Downloading Software on page xii
- Documentation Symbols Key on page xiii
- Documentation Feedback on page xiv
- Requesting Technical Support on page xv

How to Use This Guide

Complete documentation for the EX Series product family is provided on webpages at http://www.juniper.net/techpubs/en_US/release-independent/information-products/pathway-pages/ex-series/product/index.html. We have selected content from these webpages and created a number of EX Series guides that collect related topics into a book-like format so that the information is easy to print and easy to download to your local computer.

The release notes are at

http://www.juniper.net/techpubs/en_US/junos10.0/information-products/topic-collections/release-notes/10.0/junos-release-notes-10.0.pdf.

List of EX Series Guides for JUNOS Release 10.0





Title	Description
<i>Complete Hardware Guide for EX3200 and EX4200 Switches</i>	Component descriptions, site preparation, installation, replacement, and safety and compliance information for EX3200 and EX4200 switches
<i>Complete Hardware Guide for EX8208 Switches</i>	Component descriptions, site preparation, installation, replacement, and safety and compliance information for EX8208 switches
<i>Complete Hardware Guide for EX8216 Switches</i>	Component descriptions, site preparation, installation, replacement, and safety and compliance information for EX8216 switches
<i>Complete Software Guide for JUNOS® Software for EX Series Switches, Release 10.0</i>	Software feature descriptions, configuration examples, and tasks for JUNOS Software for EX Series switches

Title	Description
Software Topic Collections	Software feature descriptions, configuration examples and tasks, and reference pages for configuration statements and operational commands (This information also appears in the <i>Complete Software Guide</i> .)
<i>JUNOS® Software for EX Series Switches, Release 10.0: Access Control</i>	
<i>JUNOS® Software for EX Series Switches, Release 10.0: Alarms and System Log Messages</i>	
<i>JUNOS® Software for EX Series Switches, Release 10.0: Configuration and File Management</i>	
<i>JUNOS® Software for EX Series Switches, Release 10.0: Class of Service</i>	
<i>JUNOS® Software for EX Series Switches, Release 10.0: Device Security</i>	
<i>JUNOS® Software for EX Series Switches, Release 10.0: Ethernet Switching</i>	
<i>JUNOS® Software for EX Series Switches, Release 10.0: Interfaces</i>	
<i>JUNOS® Software for EX Series Switches, Release 10.0: Layer 3 Protocols</i>	
<i>JUNOS® Software for EX Series Switches, Release 10.0: MPLS</i>	
<i>JUNOS® Software for EX Series Switches, Release 10.0: Multicast</i>	
<i>JUNOS® Software for EX Series Switches, Release 10.0: Network Management and Monitoring</i>	
<i>JUNOS® Software for EX Series Switches, Release 10.0: Port Security</i>	
<i>JUNOS® Software for EX Series Switches, Release 10.0: Routing Policy and Packet Filtering</i>	
<i>JUNOS® Software for EX Series Switches, Release 10.0: Spanning-Tree Protocols</i>	
<i>JUNOS® Software for EX Series Switches, Release 10.0: System Setup</i>	
<i>JUNOS® Software for EX Series Switches, Release 10.0: User and Access Management</i>	
<i>JUNOS® Software for EX Series Switches, Release 10.0: Virtual Systems</i>	

Downloading Software

You can download JUNOS Software for EX Series switches from the Download Software area at <http://www.juniper.net/customers/support/>. To download the software, you must have a Juniper Networks user account. For information about obtaining an account, see <http://www.juniper.net/entitlement/setupAccountInfo.do>.

Documentation Symbols Key

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.

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 <code>stub</code> statement at the [edit protocols ospf area area-id] hierarchy level. ■ The console port is labeled CONSOLE.

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

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. Send e-mail to techpubs-comments@juniper.net with the following:

- Document URL or title
- Page number if applicable
- Software version
- Your name and company

Requesting Technical Support

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

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the JTAC User Guide located at <http://www.juniper.net/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: <https://tools.juniper.net/SerialNumberEntitlementSearch/>

Opening a Case with JTAC

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

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

For international or direct-dial options in countries without toll-free numbers, see [http://www.juniper.net/support/requesting support.html](http://www.juniper.net/support/requesting_support.html) .

Part 1

MPLS

- MPLS—Overview on page 3
- Example of MPLS Configuration on page 11
- Configuring MPLS on page 27
- Verifying MPLS on page 39
- Configuration Statements for MPLS on page 43
- Operational Mode Commands for MPLS on page 59

Chapter 1

MPLS—Overview

- JUNOS MPLS for EX Series Switches Overview on page 3
- Understanding JUNOS MPLS Components for EX Series Switches on page 6
- Understanding MPLS and Path Protection on EX Series Switches on page 10

JUNOS MPLS for EX Series Switches Overview

JUNOS MPLS for Juniper Networks EX Series Ethernet Switches supports Layer 2 protocols and Layer 2 virtual private networks (VPNs). You can configure MPLS on your switches to increase transport efficiency in your network. MPLS services can be used to connect various sites to a backbone network or to ensure better performance for low-latency applications such as VoIP and other business-critical functions.

JUNOS MPLS for EX Series switches supports RSVP-based label switched paths (LSPs) and MPLS-based circuit cross-connects (CCCs).



NOTE: MPLS configurations on EX Series switches are compatible with configurations on other Juniper Networks devices that support MPLS and CCC.

- Benefits of MPLS on page 3
- Additional Benefits of MPLS and Traffic Engineering on page 4
- MPLS Label Switched Paths and MPLS Labels on EX Series Switches on page 4
- MPLS Label Operations on EX Series Switches on page 5

Benefits of MPLS

MPLS has the following advantages over conventional packet forwarding:

- Packets arriving on different ports may be assigned different labels.
- A packet arriving at a particular provider edge switch may be assigned a different label than the same packet entering the network at a different provider edge switch. As a result, forwarding decisions that depend on the ingress provider edge switch can be easily made.
- Sometimes it is desirable to force a packet to follow a particular route that is explicitly chosen at or before the time the packet enters the network, rather than

being chosen by the normal dynamic routing algorithm as the packet travels through the network. In MPLS, a label can be used to represent the route so that the identity of the explicit route need not be carried with the packet.



NOTE: MPLS configurations on EX Series switches do not support:

- LDP-based MPLS
 - Routed VLAN interfaces (RVIs)
-

Additional Benefits of MPLS and Traffic Engineering

MPLS is the packet-forwarding component of the JUNOS traffic engineering architecture. 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 multiservice Internet.

For additional information on MPLS traffic protection on EX Series switches, see “Understanding MPLS and Path Protection on EX Series Switches” on page 10.

MPLS Label Switched Paths and MPLS Labels on EX Series Switches

When a packet enters the MPLS network, it is assigned to a label switched path (LSP). Each LSP 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. Because no additional parsing or lookup is done on the encapsulated packet, MPLS supports the transmission of any other protocols within the packet payload.



NOTE: MPLS for EX Series switches supports only single-label packets.

MPLS Label Operations on EX Series Switches

In the traditional packet-forwarding paradigm, as a packet travels from one switch to the next, an independent forwarding decision is made at each hop. The IP network 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 made only once, when a packet enters the MPLS tunnel (that is, the path used for MPLS traffic).

When an IP packet enters an LSP, the ingress provider edge switch 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 provider switch in the LSP. This switch and all subsequent switches 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 switch in the path. When the packet reaches the egress provider edge switch, the label is removed, and the packet again becomes a native IP packet and is again forwarded based on its IP routing information.

EX Series switches support the following label operations:

- Push
- Pop
- Swap

The push operation affixes a new label to the top of the IP packet. For IPv4 packets, the new label is the first label. The time to live (TTL) field value in the packet header is derived from the IP packet header. The push operation cannot be applied to a packet that has an existing MPLS label.

The pop operation removes a 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.

The swap operation removes an existing MPLS label from an IP packet and replaces it with a new MPLS label, based on the following:

- Incoming interface
- Label
- Label forwarding table

Figure 1 on page 6 shows an IP packet without a label arriving on the customer-edge interface (`ge-0/0/1`) of the provider edge ingress switch. The provider edge ingress switch examines the packet and identifies that packet's destination is the provider edge egress switch. The provider edge ingress switch applies label 100 to the packet and sends the MPLS packet to its outgoing MPLS core interface (`ge-0/0/5`). The MPLS packet is transmitted on the MPLS tunnel through the provider switch, where it arrives at interface `ge-0/0/5` with label 100. The provider switch swaps label 100 to label

200 and forwards the MPLS packet through its core interface (`ge-0/0/7`) to the next hop on the tunnel, which is the provider edge egress switch. The provider edge egress switch receives the MPLS packet through its core interface (`ge-0/0/7`), removes the MPLS label and sends the IP packet out of its customer-edge interface (`ge-0/0/1`) to a destination that is beyond the scope of the tunnel.

Figure 1: MPLS Label Swapping




Figure 1 on page 6 shows the path of a packet as it passes in one direction from the provider edge ingress switch to the provider edge egress switch. However, the MPLS configuration also allows traffic to travel in the reverse direction. Thus, each provider edge switch operates as both an ingress switch and an egress switch.

- Related Topics**
- Understanding JUNOS MPLS Components for EX Series Switches on page 6
 - Example: Configuring MPLS on EX Series Switches on page 11
 - Configuring MPLS on Provider Edge Switches (CLI Procedure) on page 30
 - Configuring MPLS on Provider Switches (CLI Procedure) on page 35
 - Configuring Path Protection in an MPLS Network (CLI Procedure) on page 27
 - *JUNOS Software MPLS Applications Configuration Guide* at <http://www.juniper.net/techpubs/software/junos/junos95/index.html>
 - *JUNOS Software VPNs Configuration Guide* at <http://www.juniper.net/techpubs/software/junos/junos95/index.html>

Understanding JUNOS MPLS Components for EX Series Switches

JUNOS MPLS for Juniper Networks EX Series Ethernet Switches supports Layer 2 protocols and Layer 2 virtual private networks (VPNs). You can configure MPLS on your switches to increase transport efficiency in your network. MPLS services can be used to connect various sites to a backbone network or to ensure better performance for low-latency applications such as VoIP and other business-critical functions.

- Provider Edge Switches on page 7
- Provider Switch on page 8
- Components Required for All Switches in the MPLS Network on page 8
- Family MPLS on page 9

Provider Edge Switches

To implement MPLS on EX Series switches, you must configure two provider edge switches—that is, an ingress provider edge switch and an egress provider edge switch.

The ingress switch (the entry point to the MPLS tunnel) receives an IP packet, analyzes it, and pushes an MPLS label onto it, which places it into a forwarding equivalence class (FEC) and determines its handling and destination through the MPLS tunnel. The egress provider edge switch (the exit point from the MPLS tunnel) pops the MPLS label off of the outgoing packet.

MPLS traffic is bidirectional. So each provider edge switch is both an ingress switch and an egress switch, depending on the direction of the traffic.

EX Series switches can handle only single-label MPLS packets. If the packet has an existing MPLS label, the provider edge switch removes the label and swaps it for another MPLS label.

MPLS Protocol and Label Switched Paths

Each provider edge switch must be configured to support the MPLS protocol, and the MPLS stanza must include the configuration of a label switched path (LSP) that specifies the address of the remote provider edge switch.

JUNOS MPLS for EX Series switches supports RSVP-based LSPs.

Circuit Cross-Connect

You must configure the customer-edge interfaces of the provider edge switches as a circuit cross-connect (CCC), creating a transparent connection between two circuits. When you configure an interface as a CCC, the interface no longer belongs to a default VLAN. The interface becomes an MPLS tunnel, used exclusively for MPLS packets. You can create different CCCs for different customers or for segregating different traffic streams over different MPLS tunnels.

Using CCC, you can connect the following types of circuits:

- Local interface with remote interface or VLAN
- Local VLAN with remote interface or VLAN



NOTE: To configure a VLAN circuit as a CCC, you must enable VLAN tagging and specify a VLAN ID.

MPLS on EX Series switches does not support the following types of CCC configurations:

- Aggregated Ethernet interface (LAG)
- Q-in-Q tunneling

- Routed VLAN interface (RVI)
- Beginning and end of the CCC on the same switch

Provider Switch

You must configure one or more provider switches as transit switches within the network to support the forwarding of MPLS packets. You can add provider switches without changing the configuration of the provider edge switches.

A provider switch does not analyze the packets. It refers to an MPLS label forwarding table and swaps one label for another. The new label determines the next hop along the MPLS tunnel. A provider switch cannot perform the push or pop operations.

Components Required for All Switches in the MPLS Network

You must configure the following components on both the provider edge and the provider switches:

- OSPF or IS-IS as a Routing Protocol on page 8
- Traffic Engineering on page 8
- MPLS Protocol on page 9
- RSVP on page 9

OSPF or IS-IS as a Routing Protocol

MPLS works in coordination with the interior gateway protocol (IGP). Therefore, you must configure either OSPF or IS-IS as a routing protocol on the loopback interface and core interfaces of both the provider edge and provider switches.

These core interfaces can be either Gigabit Ethernet or 10-Gigabit Ethernet interfaces, and they can be configured as either individual interfaces or aggregated Ethernet interfaces.



NOTE: These core interfaces cannot be configured with VLAN tagging or a VLAN ID. When you configure them to belong to **family mpls**, they are removed from the default VLAN. They operate as an exclusive tunnel for MPLS traffic.

Traffic Engineering

Traffic engineering maps traffic flows onto an existing physical topology and provides the ability to move traffic flow away from the shortest path selected by the IGP and onto a potentially less congested physical path across a network.

Traffic engineering enables the selection of specific end-to-end paths to send given types of traffic through your network. In order for MPLS to work properly, you must enable traffic engineering for the specified routing protocol (either OSPF or IS-IS).

MPLS Protocol

You must enable the MPLS protocol on all switches that participate in the MPLS network and apply it to the core interface addresses of both the provider edge and provider switches. You do not need to apply it to the loopback address, because the MPLS protocol uses the framework established by the RSVP session to create LSPs. On the provider edge switches, the configuration of the MPLS protocol must also include the definition of an LSP.

RSVP

Resource Reservation Protocol (RSVP) is a signaling protocol that allocates and distributes labels throughout an MPLS network. RSVP sets up unidirectional paths between the ingress provider edge switch and the egress provider edge switch. RSVP makes the LSPs dynamic; it can detect topology changes and outages and establish new LSPs to move around a failure.

You must enable RSVP and apply it to the loopback address and the core interface addresses of both the provider edge and provider switches. The path message contains the configured information about the resources required for the LSP to be established.

When the egress switch receives the path message, it sends a reservation message back to the ingress switch. This reservation message is passed along from switch to switch along the same path as the original path message. Once the ingress switch receives this reservation message, an RSVP path is established.

The established LSP stays active as long as the RSVP session remains active. RSVP continues activity through the transmissions and responses to RSVP path and reservation messages. If the messages stop for three minutes, the RSVP session terminates and the LSP is lost.

RSVP runs as a separate software process in the Juniper Networks JUNOS Software and is not in the packet forwarding path.

Family MPLS

You must also configure the core interface addresses used for MPLS traffic to belong to `family mpls`.



NOTE: You can enable `family mpls` on either individual interfaces or aggregated Ethernet interfaces. You cannot enable it on tagged VLAN interfaces.

-
- Related Topics**
- JUNOS MPLS for EX Series Switches Overview on page 3
 - Understanding MPLS and Path Protection on EX Series Switches on page 10
 - Example: Configuring MPLS on EX Series Switches on page 11
 - Configuring MPLS on Provider Edge Switches (CLI Procedure) on page 30
 - Configuring MPLS on Provider Switches (CLI Procedure) on page 35

- Configuring Path Protection in an MPLS Network (CLI Procedure) on page 27
- *JUNOS Software MPLS Applications Configuration Guide* at <http://www.juniper.net/techpubs/software/junos/junos95/index.html>
- *JUNOS Software VPNs Configuration Guide* at <http://www.juniper.net/techpubs/software/junos/junos95/index.html>

Understanding MPLS and Path Protection on EX Series Switches

JUNOS MPLS for Juniper Networks EX Series Ethernet Switches provides path protection to protect your MPLS network from label switched path (LSP) failures.

By default, an LSP routes itself hop-by-hop from the ingress provider edge switch through the provider switches toward the egress provider edge switch. The LSP generally follows 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.

Typically, when an LSP fails, the switch immediately upstream from the failure signals the outage to the ingress provider edge switch. The ingress provider edge switch calculates a new path to the egress provider edge switch, establishes the new LSP, and then directs 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 switch might get lost or the new path might take too long to come up, resulting in significant packet drops.

You can configure path protection by configuring primary and secondary paths on the ingress switch. If the primary path fails, the ingress switch immediately reroutes traffic from the failed path to the standby path, eliminating the need for the ingress switch to calculate a new route and signal a new path. For information about configuring standby LSPs, see “Configuring Path Protection in an MPLS Network (CLI Procedure)” on page 27.

- Related Topics**
- JUNOS MPLS for EX Series Switches Overview on page 3
 - Understanding JUNOS MPLS Components for EX Series Switches on page 6
 - Example: Configuring MPLS on EX Series Switches on page 11
 - Configuring MPLS on Provider Edge Switches (CLI Procedure) on page 30
 - *JUNOS Software MPLS Applications Configuration Guide* at <http://www.juniper.net/techpubs/software/junos/junos95/index.html>

Chapter 2

Example of MPLS Configuration

- Example: Configuring MPLS on EX Series Switches on page 11

Example: Configuring MPLS on EX Series Switches

You can configure MPLS on your switches to increase transport efficiency in your network. MPLS services can be used to connect various sites to a backbone network or to ensure better performance for low-latency applications such as VoIP and other business-critical functions.

This example shows how to configure an MPLS tunnel:

- Requirements on page 11
- Overview and Topology on page 12
- Configuring the Local Provider Edge Switch on page 15
- Configuring the Remote Provider Edge Switch on page 18
- Configuring the Provider Switch on page 21
- Verification on page 23

Requirements

This example uses the following hardware and software components:

- JUNOS Release 9.5 or later for EX Series switches
- Three EX Series switches

Before you begin configuring MPLS, ensure that you have configured either the OSPF or IS-IS routing protocol on the switches. This example includes the configuration of OSPF on the switches.

Overview and Topology

You can configure MPLS on your switches to increase transport efficiency in your network. This example includes an ingress or local provider edge switch, an egress or remote provider edge switch, and one provider (transit) switch. It includes CCCs that tie the customer-edge interface of the local provider edge switch (PE-1) to the customer-edge interface of the remote provider edge switch (PE-2). It also describes how to configure the core interfaces of the provider edge switches and the provider switch to support the transmission of the MPLS packets. In this example, the core interfaces that connect the local provider edge switch and the provider switch are individual interfaces; whereas the core interfaces that connect the remote provider edge switch and the provider switch are aggregated Ethernet interfaces.



NOTE: You do not need to create a LAG for an MPLS connection. This example includes a LAG between the provider switch and the remote provider edge switch, because this type of configuration is another option that you may wish to implement. For information on configuring LAGs, see *Configuring Aggregated Ethernet Interfaces (CLI Procedure)*.

Figure 2 on page 12 shows the topology used in this example.

Figure 2: Configuring MPLS on EX Series Switches

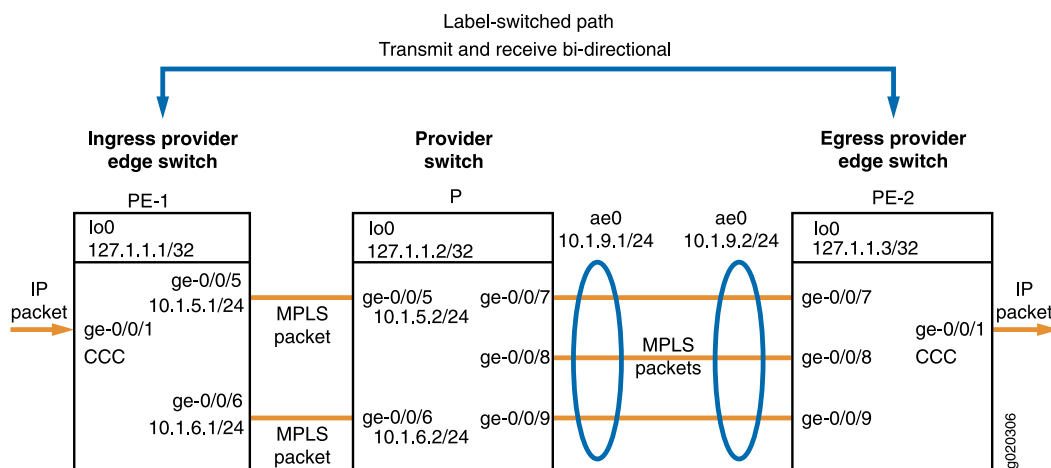


Table 1 on page 13 shows the MPLS configuration components used for the ingress provider edge switch in this example.

Table 1: Components of the Ingress PE Switch in Topology for MPLS with Interface-Based CCC

Property	Settings	Description
Local provider edge switch hardware	EX Series switch	PE-1
Loopback address	lo0 127.1.1.1/32	Identifies PE-1 for interswitch communications.
Routing protocol	ospf traffic-engineering	Indicates that this switch is using OSPF as the routing protocol and that traffic engineering is enabled.
MPLS protocol and definition of label switched path	mpls label-switched-path lsp_to_pe2_ge1 to 127.1.13	Indicates that this provider edge switch is using the MPLS protocol with the specified label switched path (LSP) to reach the other provider edge switch (specified by the loopback address). The statement must also specify the core interfaces to be used for MPLS traffic.
RSVP protocol	rsvp	Indicates that this switch is using the RSVP protocol. The statement must specify the loopback address and the core interfaces that are going to be used for the RSVP session.
Interface family	family inet family mpls family ccc	The logical units of the core interfaces are configured to belong to both family inet and family mpls . The logical unit of the customer edge interface is configured to belong to family ccc .
Customer-edge interface	ge-0/0/1	Interface that connects this network to devices outside the network.
Core interfaces	ge-0/0/5.0 and ge-0/0/6.0 with IP addresses 10.1.5.1/24 and 10.1.6.1/24	Interfaces that connect to other switches within the MPLS network.
CCC definition	connections remote-interface-switch ge-1-to-pe2 interface ge-0/0/1.0 transmit-lsp lsp_to_pe2_ge1 receive-lsp lsp_to_pe1_ge1	Associates the circuit cross-connect (CCC), ge-0/0/1 , with the LSPs that have been defined on the local and remote provider edge switches.

Table 2 on page 14 shows the MPLS configuration components used for the egress provider edge switch in this example.

Table 2: Components of the Egress PE Switch in Topology for MPLS with Interface-Based CCC

Property	Settings	Description
Remote provider edge switch hardware	EX Series switch	PE-2
Loopback address	lo0 127.1.1.3/32	Identifies PE-2 for interswitch communications.
Routing protocol	ospf traffic-engineering	Indicates that this switch is using OSPF as the routing protocol and that traffic engineering is enabled.
MPLS protocol and definition of label-switched path	mpls label-switched-path lsp_to_pe1_ge1 to 127.1.1.1	Indicates that this provider edge switch is using the MPLS protocol with the specified label switched path (LSP) to reach the other PE switch. The statement must also specify the core interfaces to be used for MPLS traffic.
RSVP protocol	rsvp	Indicates that this switch is using the RSVP protocol. The statement must specify the loopback address and the core interfaces that are going to be used for the RSVP session.
Interface family	family inet family mpls family ccc	The logical unit of the core interface is configured to belong to both family inet and family mpls . The logical unit of the customer edge interface is configured to belong to family ccc .
Customer-edge interface	ge-0/0/1	Interface that connects this network to devices outside the network.
Core interface	ae0 with IP address 10.1.9.2/24	Aggregated Ethernet interface on PE-2 that connects to aggregated Ethernet interface ae0 of the provider switch and belongs to family mpls .
CCC definition	connections remote-interface-switch ge-1-to-pe1 interface ge-0/0/1.0 transmit-lsp lsp_to_pe1_ge1; receive-lsp lsp_to_pe2_ge1;	Associates the circuit cross-connect (CCC), ge-0/0/1 , with the LSPs that have been defined on the local and remote provider edge switches.

Table 3 on page 15 shows the MPLS configuration components used for the provider switch in this example.

Table 3: Components of the Provider Switch in Topology for MPLS with Interface-Based CCC

Property	Settings	Description
Provider switch hardware	EX Series switch	Transit switch within the MPLS network configuration.
Loopback address	lo0 127.1.1.2/32	Identifies provider switch for interswitch communications.
Routing protocol	ospf traffic-engineering	Indicates that this switch is using OSPF as the routing protocol and that traffic engineering is enabled.
MPLS protocol	mpls	Indicates that this switch is using the MPLS protocol. The statement must specify the core interfaces that are going to be used for MPLS traffic.
RSVP protocol	rsvp	Indicates that this switch is using the RSVP protocol. The statement must specify the loopback address and the core interfaces that are going to be used for the RSVP session.
Interface family	family inet family mpls	The logical units for the loopback address and core interfaces belong to family inet. The logical units of the core interfaces are also configured to belong to family mpls.
Core interfaces	ge-0/0/5.0 and ge-0/0/6.0 with IP addresses 10.1.5.1/24 and 10.1.6.1/24 and ae0 with IP address 10.1.9.1/24	Interfaces that connect P to PE-1. Aggregated Ethernet interface on P that connects to aggregated Ethernet interface ae0 of PE-2.

Configuring the Local Provider Edge Switch

CLI Quick Configuration To quickly configure the local provider edge switch, copy the following commands and paste them into the switch terminal window of PE-1:

```
[edit]
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface ge-0/0/5.0
set protocols ospf area 0.0.0.0 interface ge-0/0/6.0
set protocols mpls label-switched-path lsp_to_pe2_ge1 to 127.1.1.3
set protocols mpls interface ge-0/0/5.0
```

```

set protocols mpls interface ge-0/0/6.0
set protocols rsvp interface lo0.0
set protocols rsvp interface ge-0/0/5.0
set protocols rsvp interface ge-0/0/6.0
set interfaces lo0 unit 0 family inet address 127.1.1.1/32
set interfaces ge-0/0/5 unit 0 family inet address 10.1.5.1/24
set interfaces ge-0/0/6 unit 0 family inet address 10.1.6.1/24
set interfaces ge-0/0/5 unit 0 family mpls
set interfaces ge-0/0/6 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family ccc
set protocols connections remote-interface-switch ge-1-to-pe2 interface ge-0/0/1.0
set protocols connections remote-interface-switch ge-1-to-pe2 transmit-lsp
lsp_to_pe2_ge1
set protocols connections remote-interface-switch ge-1-to-pe2 receive-lsp
lsp_to_pe1_ge1

```

Step-by-Step Procedure To configure the ingress provider edge switch:

1. Configure OSPF with traffic engineering enabled:

```

[edit protocols]
user@switchPE-1# set ospf traffic-engineering

```

2. Configure OSPF on the loopback address and core interfaces:

```

[edit protocols]
user@switchPE-1# set ospf area 0.0.0.0 interface lo0.0
user@switchPE-1# set ospf area 0.0.0.0 interface ge-0/0/5.0
user@switchPE-1# set ospf area 0.0.0.0 interface ge-0/0/6.0

```

3. Configure MPLS on the switch with a label switched path to the remote provider edge switch:

```

[edit protocols]
user@switchPE-1# set mpls label-switched-path lsp_to_pe2_ge1 to 127.1.1.3

```

4. Configure MPLS on the core interfaces:

```

[edit protocols]
user@switchPE-1# set mpls interface ge-0/0/5.0
user@switchPE-1# set mpls interface ge-0/0/6.0

```

5. Configure RSVP on the loopback address and core interfaces:

```

[edit protocols]
user@switchPE-1# set rsvp interface lo0.0
user@switchPE-1# set rsvp interface ge-0/0/5.0
user@switchPE-1# set rsvp interface ge-0/0/6.0

```

6. Configure IP addresses for the loopback and core interfaces:

```
[edit]
user@switchPE-1# set interfaces lo0 unit 0 family inet address 127.1.1.1/32
user@switchPE-1# set interfaces ge-0/0/5 unit 0 family inet address
10.1.5.1/24
user@switchPE-1# set interfaces ge-0/0/6 unit 0 family inet address
10.1.6.1/24
```

7. Configure family mpls on the logical unit of the core interface addresses:

```
[edit]
user@switchPE-1# set interfaces ge-0/0/5 unit 0 family mpls
user@switchPE-1# set interfaces ge-0/0/6 unit 0 family mpls
```

9. Configure the logical unit of the customer-edge interface as a CCC:

```
[edit interfaces ge-0/0/1 unit 0]
user@PE-1# set family ccc
```

10. Configure the interface-based CCC from PE-1 to PE-2:



NOTE: You can also configure a tagged VLAN interface as a CCC. See “Configuring MPLS on Provider Edge Switches (CLI Procedure)” on page 30.

```
[edit protocols]
user@PE-1# set connections remote-interface-switch ge-1-to-pe2 interface ge-0/0/1.0
user@PE-1# set connections remote-interface-switch ge-1-to-pe2 transmit-lsp
lsp_to_pe2_ge1
user@PE-1# set connections remote-interface-switch ge-1-to-pe2 receive-lsp
lsp_to_pe1_ge1
```

Results Display the results of the configuration:

```
user@switchPE-1> show configuration

interfaces {
  ge-0/0/1 {
    unit 0 {
      family ccc;
    }
  }
  ge-0/0/5 {
    unit 0 {
      family inet {
        address 10.1.5.1/24;
      }
      family mpls;
    }
  }
}
```

```

ge-0/0/6 {
  unit 0 {
    family inet {
      address 10.1.6.1/24;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 127.1.1.1/32;
    }
  }
}
protocols {
  rsvp {
    interface lo0.0;
    interface ge-0/0/5.0;
    interface ge-0/0/6.0;
  }
  mpls {
    label-switched-path lsp_to_pe2_ge1 {
      to 127.1.1.3;
    }
    interface ge-0/0/5.0;
    interface ge-0/0/6.0;
  }
  ospf {
    traffic-engineering;
    area 0.0.0.0 {
      interface lo0.0;
      interface ge-0/0/5.0;
      interface ge-0/0/6.0;
    }
  }
  connections {
    remote-interface-switch ge-1-to-pe2 {
      interface ge-0/0/1.0;
      transmit-lsp lsp_to_pe2_ge1;
      receive-lsp lsp_to_pe1_ge1;
    }
  }
}

```

Configuring the Remote Provider Edge Switch

CLI Quick Configuration To quickly configure the remote provider edge switch, copy the following commands and paste them into the switch terminal window of PE-2:

```

[edit]
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface ae0
set protocols mpls label-switched-path lsp_to_pe1_ge1 to 127.1.1.1
set protocols mpls interface ae0

```

```

set protocols rsvp interface lo0.0
set protocols rsvp interface ae0
set interfaces lo0 unit 0 family inet address 127.1.1.3/32
set interfaces ae0 unit 0 family inet address 10.1.9.2/24
set interfaces ae0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family ccc
set protocols connections remote-interface-switch ge-1-to-pe1 interface ge-0/0/1.0
set protocols connections remote-interface-switch ge-1-to-pe1 transmit-lsp
lsp_to_pe1_ge1
set protocols connections remote-interface-switch ge-1-to-pe1 receive-lsp
lsp_to_pe2_ge1

```

Step-by-Step Procedure To configure PE-2:

1. Configure OSPF with traffic engineering enabled:

```

[edit protocols]
user@switchPE-2# set ospf traffic-engineering

```

2. Configure OSPF on the loopback address and core interface of PE-2:

```

[edit protocols]
user@switchPE-2# set ospf area 0.0.0.0 interface lo0.0
user@switchPE-2# set ospf area 0.0.0.0 interface ae0

```

3. Configure MPLS on the switch with a label switched path to the remote provider edge switch:

```

[edit protocols]
user@switchPE-2# set mpls label-switched-path lsp_to_pe1_ge1 to 127.1.1.1

```

4. Configure MPLS on the core interface:

```

[edit protocols]
user@switchPE-2# set mpls interface ae0

```

5. Configure RSVP on the loopback address and core interface of PE-2:

```

[edit protocols]
ser@switchPE-2# set rsvp interface lo0.0
user@switchPE-2# set rsvp interface ae0

```

6. Configure IP addresses for the loopback and core interfaces:

```

[edit]
user@switchPE-2# set interfaces lo0 unit 0 family inet address 127.1.1.3/32
user@switchPE-2# set interfaces ae0 unit 0 family inet address 10.1.9.2/24

```

7. Configure family mpls on the logical unit of the core interface address of PE-2:

```

[edit]
user@switchPE-2# set interfaces ae0 unit 0 family mpls

```

8. Configure the logical unit of the customer-edge interface as a CCC:

```
[edit interfaces ge-0/0/1 unit 0]
user@PE-2# set family ccc
```

- Configure the interface-based CCC from PE-2 to PE-1:

```
[edit protocols]
user@PE-2# set connections remote-interface-switch ge-1-to-pe2 interface
ge-0/0/1.0
user@PE-2# set connections remote-interface-switch ge-1-to-pe2 transmit-lsp
lsp_to_pe1_ge1
user@PE-2# set connections remote-interface-switch ge-1-to-pe2 receive-lsp
lsp_to_pe2_ge1
```

Results Display the results of the configuration:

```
user@switchPE-2> show configuration

interfaces {
  ge-0/0/1 {
    unit 0 {
      family ccc;
    }
  }
  ae0 {
    unit 0 {
      family inet {
        address 10.1.9.2/24;
      }
      family mpls;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 127.1.1.3/32;
      }
    }
  }
}
protocols {
  rsvp {
    interface lo0.0;
    interface ae0.0;
  }
  mpls {
    label-switched-path lsp_to_pe1_ge1 {
      to 127.1.1.1;
    }
    interface ae0.0;
  }
  ospf {
    traffic-engineering;
    area 0.0.0.0 {
```

```

        interface ae0.0;
    }
}
connections {
    remote-interface-switch ge-1-to-pe1 {
        interface ge-0/0/1.0;
        transmit-lsp lsp_to_pe1_ge1;
        receive-lsp lsp_to_pe2_ge1;
    }
}
}

```

Configuring the Provider Switch

CLI Quick Configuration To quickly configure the provider switch, copy the following commands and paste them into the switch terminal window:

```

[edit]
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface ge-0/0/5.0
set protocols ospf area 0.0.0.0 interface ge-0/0/6.0
set protocols ospf area 0.0.0.0 interface ae0
set protocols mpls interface ge-0/0/5.0
set protocols mpls interface ge-0/0/6.0
set protocols mpls interface ae0
set protocols rsvp interface lo0.0
set protocols rsvp interface ge-0/0/5.0
set protocols rsvp interface ge-0/0/6.0
set protocols rsvp interface ae0
set interfaces lo0 unit 0 family inet address 127.1.1.2/32
set interfaces ge-0/0/5 unit 0 family inet address 10.1.5.1/24
set interfaces ge-0/0/6 unit 0 family inet address 10.1.6.1/24
set interfaces ae0 unit 0 family inet address 10.1.9.1/24
set interfaces ge-0/0/5 unit 0 family mpls
set interfaces ge-0/0/6 unit 0 family mpls
set interfaces ae0 unit 0 family mpls

```

Step-by-Step Procedure To configure the provider switch:

1. Configure OSPF with traffic engineering enabled:

```

[edit protocols]
user@switchP# set ospf traffic-engineering

```

2. Configure OSPF on the loopback and core interfaces:

```

[edit protocols]
user@switchP# set ospf area interface lo0.0
user@switchP# set ospf area interface ae0

```

3. Configure MPLS on the core interfaces on the switch:

```

[edit protocols]
user@switchP# set mpls interface ge-0/0/5

```

```
user@switchP# set mpls interface ge-0/0/6
user@switchP# set mpls interface ae0
```

4. Configure RSVP on the loopback address and core interfaces:

```
[edit protocols]
user@switchP# set rsvp interface lo0.0
user@switchP# set rsvp interface ge-0/0/5
user@switchP# set rsvp interface ge-0/0/6
user@switchP# set rsvp interface ae0
```

5. Configure IP addresses for the loopback and core interfaces:

```
[edit]
user@switchP# set interfaces lo0 unit 0 family inet address 127.1.1.2/32
user@switchP# set interfaces ge-0/0/5 unit 0 family inet address
10.1.5.1/24
user@switchP# set interfaces ge-0/0/6 unit 0 family inet address
10.1.6.1/24
user@switchP# set interfaces ae0 unit 0 family inet address 10.1.9.1/24
```

6. Configure family mpls on the logical unit of the core interface addresses:

```
[edit]
user@switchP# set interfaces ge-0/0/5 unit 0 family mpls
user@switchP# set interfaces ge-0/0/6 unit 0 family mpls
user@switchP# set interfaces ae0 unit 0 family mpls
```

Results Display the results of the configuration:

```
user@switchP> show configuration
```

```
interfaces {
  ge-0/0/5 {
    unit 0 {
      family inet {
        address 10.1.5.1/24;
      }
      family mpls;
    }
  }
  ge-0/0/6 {
    unit 0 {
      family inet {
        address 10.1.6.1/24;
      }
      family mpls;
    }
  }
  ae0 {
    unit 0 {
      family inet {
```

```

        address 10.1.9.1/24;
    }
    family mpls;
}
lo0 {
    unit 0 {
        family inet {
            address 127.1.1.2/32;
        }
    }
}
protocols {
    rsvp {
        interface lo0.0;
        interface ge-0/0/5.0;
        interface ge-0/0/6.0;
        interface ae0.0;
    }
    mpls {
        interface ge-0/0/5.0;
        interface ge-0/0/6.0;
        interface ae0.0;
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface lo0.0;
            interface ge-0/0/5.0;
            interface ge-0/0/6.0;
            interface ae0.0;
        }
    }
}

```

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying the Physical Layer on the Switches on page 23
- Verifying the Routing Protocol on page 24
- Verifying the Core Interfaces Being Used for the MPLS Traffic on page 24
- Verifying RSVP on page 25
- Verifying the Assignment of Interfaces for MPLS Label Operations on page 25
- Verifying the Status of the CCC on page 25

Verifying the Physical Layer on the Switches

Purpose Verify that the interfaces are up. Perform this verification task on each of the switches.

Action user@switchPE-1> **show interfaces ge- terse**

Interface	Admin	Link	Proto	Local	Remote
ge-0/0/0	up	up			
ge-0/0/0.0	up	up			
ge-0/0/1.0	up	up	ccc		
ge-0/0/2.0	up	up	eth-switch		
ge-0/0/3.0	up	up	eth-switch		
ge-0/0/4.0	up	up	eth-switch		
ge-0/0/5.0	up	up	inet	10.1.5.1/24	
		mpls			
ge-0/0/6.0	up	up	inet	10.1.6.1/24	
		mpls			

Meaning The `show interfaces terse` command displays status information about the Gigabit Ethernet interfaces on the switch. This output verifies that the interfaces are `up`. The output for the protocol family (`Proto` column) shows that interface `ge-0/0/0.1` is configured as a circuit cross-connect. The output for the protocol family of the core interfaces (`ge-0/0/0.5` and `ge-0/0/0.6`), shows that these interfaces are configured as both `inet` and `mpls`. The `Local` column for the core interfaces shows the IP address configured for these interfaces.

Verifying the Routing Protocol

Purpose Verify the state of the configured routing protocol. Perform this verification task on each of the switches. The state should be `Full`.

Action `user@switchPE-1> show ospf neighbor`

Address	Interface	State	ID	Pri	Dead
127.1.1.2	ge-0/0/5	Full	10.10.10.10	128	39

Meaning The `show ospf neighbor` command displays the status of the routing protocol. This output shows that the state is `Full`, meaning that the routing protocol is operating correctly—that is, hello packets are being exchanged between directly connected neighbors.

Verifying the Core Interfaces Being Used for the MPLS Traffic

Purpose Verify that the state of the MPLS interface is `Up`. Perform this verification task on each of the switches.

Action `user@switchPE-1> show mpls interface`

Interface	State	Administrative groups
ge-0/0/5	Up	<none>
ge-0/0/6	Up	<none>

Meaning The `show mpls interface` command displays the status of the core interfaces that have been configured to belong to `family mpls`. This output shows that the interface configured to belong to `family mpls` is `Up`.

Verifying RSVP

Purpose Verify the state of the RSVP session. Perform this verification task on each of the switches.

Action user@switchPE-1> **show rsvp session**

```
Ingress RSVP: 1 sessions
To          From          State  Rt Style Labelin Labelout LSPname
127.1.13    127.1.1.1        Up     0  1 FF      -   300064 lsp_to_pe2_ge1
Total 1 displayed, Up 1, Down 0

Egress RSVP: 1 sessions
To          From          State  Rt Style Labelin Labelout LSPname
127.1.1.1    127.1.1.3      Up     0  1 FF  299968      -
lsp_to_pe1_ge1
Total 1 displayed, Up 1, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Meaning This output confirms that the RSVP sessions are Up.

Verifying the Assignment of Interfaces for MPLS Label Operations

Purpose Verify which interface is being used as the beginning of the CCC and which interface is being used to push the MPLS packet to the next hop. Perform this task only on the provider edge switches.

Action user@switchPE-1> **show route forwarding-table family mpls**

```
MPLS:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0      0                  dscd  50   1
0                user  0      0                  recv  49   3
1                user  0      0                  recv  49   3
2                user  0      0                  recv  49   3
299776           user  0      0                  Pop   541  2 ge-0/0/1.0
ge-0/0/1.0 (CCC) user  0  2.0.0.1          Push 299792 540 2 ge-0/0/5.0
```

Meaning This output shows that the CCC has been set up on interface `ge-0/0/1.0`. The switch receives ingress traffic on `ge-0/0/1.0` and pushes label `299792` onto the packet, which goes out through interface `ge-0/0/5.0`. The output also shows when the switch receives an MPLS packet with label `29976`, it pops the label and sends the packet out through interface `ge-0/0/1.0`.

After you have checked the local provider edge switch, run the same command on the remote provider edge switch.

Verifying the Status of the CCC

Purpose Verify the status of the CCC. Perform this task only on the provider edge switches.

Action user@switchPE-1> **show connections**

```

CCC and TCC connections [Link Monitoring On]
Legend for status (St)
UN -- uninitialized
NP -- not present
WE -- wrong encapsulation
DS -- disabled
Dn -- down
-> -- only outbound conn is up
<- -- only inbound conn is up
Up -- operational
RmtDn -- remote CCC down
Restart -- restarting

Legend for connection types
if-sw: interface switching
rmt-if: remote interface switching
lsp-sw: LSP switching
tx-p2mp-sw: transmit P2MP switching
rx-p2mp-sw: receive P2MP switching

Legend for circuit types
intf -- interface
tlsp -- transmit LSP
rlsp -- receive LSP
    
```

Connection/Circuit	Type	St	Time last up	# Up trans
ge1-to-pe2	rmt-if	Up	Feb 17 05:00:09	1
ge-0/0/1.0	intf	Up		
lsp_to_pe1_ge1	tlsp	Up		
lsp_to_pe2_ge1	rlsp	Up		

Meaning The `show connections` command displays the status of the CCC connections. This output verifies that the CCC interface and its associated transmit and receive LSPs are `Up`. After you have checked the local provider edge switch, run the same command on the remote provider edge switch.

- Related Topics**
- Configuring MPLS on Provider Edge Switches (CLI Procedure) on page 30
 - Configuring MPLS on Provider Switches (CLI Procedure) on page 35
 - JUNOS MPLS for EX Series Switches Overview on page 3
 - For information on the interface statement for OSPF, see the *JUNOS Software Routing Protocols Configuration Guide* at <http://www.juniper.net/techpubs/software/junos/junos95/index.html>.

Chapter 3

Configuring MPLS

- Configuring Path Protection in an MPLS Network (CLI Procedure) on page 27
- Configuring MPLS on Provider Edge Switches (CLI Procedure) on page 30
- Configuring MPLS on Provider Switches (CLI Procedure) on page 35

Configuring Path Protection in an MPLS Network (CLI Procedure)

The JUNOS Software implementation of MPLS on EX Series switches provides path protection as a mechanism for protecting against label switched path (LSP) failures. Path protection reduces the time required to recalculate a route in case of a failure within the MPLS tunnel. You configure path protection on the ingress provider edge switch in your MPLS network. You do not configure the egress provider edge switch or the provider switches for path protection. You can explicitly specify which provider switches are used for the primary and secondary paths, or you can let the software calculate the paths automatically.

Before you configure path protection, be sure you have:

- Configured an ingress provider edge switch and an egress provider edge switch. See “Configuring MPLS on Provider Edge Switches (CLI Procedure)” on page 30.
- Configured at least one provider (transit) switch. See “Configuring MPLS on Provider Switches (CLI Procedure)” on page 35.
- Verified the configuration of your MPLS network. See “Verifying That MPLS Is Working Correctly” on page 39.

To configure path protection, complete the following tasks on the ingress provider edge switch:

1. Configuring the Primary Path on page 28
2. Configuring the Secondary Path on page 28
3. Configuring the Revert Timer on page 29

Configuring the Primary Path

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 provider edge switch.

In the tasks described in this topic, the *lsp-name* has already been configured on the ingress provider edge switch as **lsp_to_240** and the loopback interface address on the remote provider edge switch has already been configured as **127.0.0.8**.

When the software switches from the primary to the 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 **revert-timer** statement.

You can configure zero primary paths or one primary path. If you do not configure a primary path, the first secondary path (if a secondary path has been configured) is selected as the 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 for the packets to reach the egress provider edge switch.

To configure a primary path:

1. Create the primary path for the LSP:

```
[edit protocols mpls label-switched-path lsp_to_240 to 127.0.0.8]
user@switch# set primary primary_path_lsp_to_240
```

2. Configure an explicit route for the primary path by specifying the IP address of the loopback interface or the switch IP address or hostname of each switch used in the MPLS tunnel. You can specify the link types as either **strict** or **loose** in each **path** statement. If the link type is **strict**, the LSP must go to the next address specified in the **path** statement without traversing other switches. If the link type is **loose**, the LSP can traverse through other switches before reaching this switch. This configuration uses the default **strict** designation for the paths.



NOTE: You can enable path protection without specifying which provider switches are used. If you do not list the specific provider switches to be used for the MPLS tunnel, the switch calculates the route.



TIP: Do not include the ingress provider edge switch in these statements. List the IP address of the loopback interface or switch address or hostname of all other switch hops in sequence, ending with the egress provider edge switch.

```
[edit protocols mpls label-switched-path lsp_to_240 to 127.0.0.8]
user@switch# set path primary_path_lsp_to_240 127.0.0.2
user@switch# set path primary_path_lsp_to_240 127.0.0.3
user@switch# set path primary_path_lsp_to_240 127.0.0.8
```

Configuring the Secondary 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 first secondary path in the configuration is not available, the next one is tried, as so on. 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 provider edge switch.

To configure the secondary path:

1. Create a secondary path for the LSP:

```
[edit protocols mpls label-switched-path lsp_to_240 to 127.0.0.8]
user@switch# set secondary secondary_path_lsp_to_240 standby
```

2. Configure an explicit route for the secondary path by specifying the IP address of the loopback interface or the switch IP address or hostname of each switch used in the MPLS tunnel. You can specify the link types as either **strict** or **loose** in each **path** statement. This configuration uses the default **strict** designation for the paths.



TIP: Do not include the ingress provider edge switch in these statements. List the IP address of the loopback interface or switch address or hostname of all other switch hops in sequence, ending with the egress provider edge switch.

```
[edit protocols mpls label-switched-path lsp_to_240 to 127.0.0.8]
user@switch# set path secondary_path_lsp_to_240 127.0.0.4
user@switch# set path primary_path_lsp_to_240 127.0.0.8
```

Configuring the Revert Timer

For LSPs configured with both primary and secondary paths, you can optionally configure a revert timer. If the 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 the primary path. If the primary path experiences any connectivity problems or stability problems during this time, the timer is restarted.



TIP: If you do not explicitly configure the revert timer, it is set by default to 60 seconds.

To configure the revert timer for LSPs configured with primary and secondary paths:

- For all LSPs on the switch:

```
[edit protocols mpls]
user@switch# set revert-timer 120
```

- For a specific LSP on the switch:

```
[edit protocols mpls label-switched-path]
user@switch# set lsp_to_240 revert-timer 120
```

- Related Topics** ■ Understanding MPLS and Path Protection on EX Series Switches on page 10

Configuring MPLS on Provider Edge Switches (CLI Procedure)

JUNOS MPLS for EX Series switches supports Layer 2 protocols and Layer 2 virtual private networks (VPNs). You can configure MPLS on your switches to increase transport efficiency in your network. MPLS services can be used to connect various sites to a backbone network or to ensure better performance for low-latency applications such as VoIP and other business-critical functions.

To implement MPLS on EX Series switches, you must configure two provider edge switches: an ingress provider edge switch and an egress provider edge switch.

To configure a provider edge switch, complete the following tasks. When you have completed configuring one provider edge switch, perform the same tasks on the other provider edge switch:

1. Enabling the OSPF or the IS-IS Routing Protocol on the Loopback and Core Interfaces on page 31
2. Enabling Traffic Engineering for the Routing Protocol on page 31
3. Configuring IP Addresses for the Loopback and Core Interfaces on page 31
4. Enabling MPLS, Defining the Label Switched Path, and Applying MPLS to the Core Interfaces on page 32
5. Enabling RSVP and Applying It to the Loopback and Core Interfaces on page 32

6. Enabling Family MPLS on the Core Interfaces on page 33
7. Configuring a Circuit Cross-Connect on a Customer Edge Interface on page 33

Enabling the OSPF or the IS-IS Routing Protocol on the Loopback and Core Interfaces

Enable OSPF or IS-IS on the loopback address and on the core interface addresses.



NOTE: You can use the switch address as an alternative to the loopback address.

1. Configure OSPF on the loopback and core interfaces:

```
[edit protocols]
user@switch# set ospf area 0.0.0.0 interface lo0.0
user@switch# set ospf area 0.0.0.0 interface ge-0/0/5.0
user@switch# set ospf area 0.0.0.0 interface ge-0/0/6.0
user@switch# set ospf area 0.0.0.0 interface ae0
```

Enabling Traffic Engineering for the Routing Protocol

Enable traffic engineering for the routing protocol (either OSPF or IS-IS) on the loopback address and on the core interface addresses.

1. Enable traffic engineering for the routing protocol:

```
[edit protocols]
user@switch# set ospf traffic-engineering
```

Configuring IP Addresses for the Loopback and Core Interfaces

Configure IP addresses for the loopback and core interfaces.

1. Configure an IP address for the loopback interface and for the core interfaces:

```
[edit]
user@switch# set interfaces lo0 unit 0 family inet address 127.1.1.1/32
user@switch# set interfaces ge-0/0/5 unit 0 family inet address 10.1.5.1/24
user@switch# set interfaces ge-0/0/6 unit 0 family inet address 10.1.6.1/24
user@switch# set interfaces ae0 unit 0 family inet address 10.1.9.1/24
```

Enabling MPLS, Defining the Label Switched Path, and Applying MPLS to the Core Interfaces

Enable MPLS within the `protocols` stanza and define a label switched path (LSP) that specifies the loopback interface address or switch address of the destination (egress) provider edge switch. You must also apply MPLS to the core interfaces.

1. Enable MPLS and define the LSP:

```
[edit protocols]
user@switch# set mpls label-switched-path lsp_to_pe2_ge1 to 127.1.1.3
```



TIP: `lsp_to_pe2_ge1` is the LSP name. You will need to use the specified name again when configuring the circuit cross-connect.

2. Configure MPLS on the core interface addresses:

```
[edit protocols]
user@switch# set mpls interface ge-0/0/5.0
user@switch# set mpls interface ge-0/0/6.0
user@switch# set mpls interface ae0
```

Enabling RSVP and Applying It to the Loopback and Core Interfaces

Enable RSVP and apply it to the loopback and the core interfaces.

1. Configure RSVP on the loopback address and the core interface addresses:

```
[edit protocols]
user@switch# set rsvp interface lo0.0
user@switch# set rsvp interface ge-0/0/5.0
user@switch# set rsvp interface ge-0/0/6.0
user@switch# set rsvp interface ae0
```

Enabling Family MPLS on the Core Interfaces

On the logical units of the core interface addresses, configure these interfaces to belong to family `mpls`. This configuration identifies the interfaces used for forwarding MPLS packets.

1. Configure family `mpls` on the logical units of the core interfaces:

```
[edit]
user@switch# set interfaces ge-0/0/5 unit 0 family mpls
user@switch# set interfaces ge-0/0/6 unit 0 family mpls
user@switch# set interfaces ae0 unit 0 family mpls
```



NOTE: You can enable family `mpls` on either individual interfaces or aggregated Ethernet interfaces. You cannot enable it on tagged VLAN interfaces.

Configuring a Circuit Cross-Connect on a Customer Edge Interface

You must configure the customer edge interfaces of the provider edge switches as a circuit cross-connect (CCC). The customer edge interface can be either a simple interface or a tagged VLAN interface. In both cases, you must configure the logical unit of the customer edge interface to belong to **family ccc** and you must configure an association between that interface and two LSPs—one for transmitting MPLS packets to the remote provider edge and the other for receiving MPLS packets from the remote PE.



NOTE: When an interface is configured to belong to **family ccc**, it cannot belong to any other family.

If you are configuring a CCC on a tagged VLAN interface, you must explicitly enable VLAN tagging and specify a VLAN ID.



NOTE: The VLAN tag ID cannot be configured on logical interface unit 0. The logical unit number must be 1 or higher.

This procedure shows how to set up two CCCs:

- If you are configuring a CCC on a simple interface (**ge-0/0/1**), omit step 1 and step 2.
 - If you are configuring a CCC on a tagged VLAN interface (**ge-0/0/2**), include all the steps in this procedure.
1. Enable VLAN tagging on the customer edge interface **ge-0/0/2** of the local provider edge switch:

```
[edit interfaces ge-0/0/2]
user@switch# set vlan-tagging
```

2. Configure the logical unit of the customer edge interface with a VLAN ID:

```
[edit interfaces ge-0/0/2 unit 1]
user@switch# set vlan-id 100
```

3. Configure the logical unit of the customer edge interface to belong to **family ccc**:

```
[edit interfaces ge-0/0/1 unit 0]
user@switch# set family ccc
```

```
[edit interfaces ge-0/0/2 unit 1]
user@switch# set family ccc
```

4. Associate the CCC interface with two LSPs, one for transmitting MPLS packets and the other for receiving MPLS packets:

```
[edit protocols]
user@switch# set connections remote-interface-switch ge-1-to-pe2 interface
ge-0/0/1.0
```

```

user@switch# set connections remote-interface-switch ge-1-to-pe2
transmit-lsp lsp_to_pe2_ge1
user@switch# set connections remote-interface-switch ge-1-to-pe2
receive-lsp lsp_to_pe1_ge1

[edit protocols]
user@switch# set connections remote-interface-switch ge-1-to-pe2 interface
ge-0/0/2.1
user@switch# set connections remote-interface-switch ge-1-to-pe2
transmit-lsp lsp_to_pe2_ge1
user@switch# set connections remote-interface-switch ge-1-to-pe2
receive-lsp lsp_to_pe1_ge1

```



TIP: The `transmit-lsp` option specifies the LSP name that was configured on PE-1 (the local provider edge switch) by the `label-switched-path` statement within the `protocols mpls` stanza.

The `receive-lsp` option specifies the LSP name that was configured on PE-2 (the remote provider edge switch) by the `label-switched-path` statement within the `protocols mpls` stanza.

When you have completed configuring one provider edge switch, follow the same procedures to configure the other provider edge switch.

- Related Topics**
- Example: Configuring MPLS on EX Series Switches on page 11
 - Configuring MPLS on Provider Switches (CLI Procedure) on page 35
 - Configuring an OSPF Network (J-Web Procedure)
 - Verifying That MPLS Is Working Correctly on page 39
 - Understanding JUNOS MPLS Components for EX Series Switches on page 6
 - For information on the interface statement for OSPF, see the *JUNOS Software Routing Protocols Configuration Guide* at <http://www.juniper.net/techpubs/software/junos/junos95/index.html>.

Configuring MPLS on Provider Switches (CLI Procedure)

JUNOS MPLS for EX Series switches supports Layer 2 protocols and Layer 2 virtual private networks (VPNs). You can configure MPLS on your switches to increase transport efficiency in your network. MPLS services can be used to connect various sites to a backbone network or to ensure better performance for low-latency applications such as VoIP and other business-critical functions.

To implement MPLS on EX Series switches, you must configure at least one provider switch as a transit switch for the MPLS packets.

To configure the provider switch, complete the following tasks:

1. Enabling the OSPF or the IS-IS Routing Protocol on the Loopback and Core Interfaces on page 36
2. Enabling Traffic Engineering for the Routing Protocol on page 36
3. Enabling MPLS and Applying MPLS to the Core Interfaces on page 36
4. Enabling RSVP and Applying It to the Loopback and Core Interfaces on page 37
5. Configuring IP Addresses for the Loopback and Core Interfaces on page 37
6. Enabling Family MPLS on the Core Interfaces on page 37

Enabling the OSPF or the IS-IS Routing Protocol on the Loopback and Core Interfaces

Enable OSPF or IS-IS on the loopback address and on the core interface addresses.



NOTE: You can use the switch address as an alternative to the loopback address.

1. Configure OSPF on the loopback and core interfaces:

```
[edit protocols]
user@switch# set ospf area 0.0.0.0 interface lo0.0
user@switch# set ospf area 0.0.0.0 interface ge-0/0/5.0
user@switch# set ospf area 0.0.0.0 interface ge-0/0/6.0
user@switch# set ospf area 0.0.0.0 interface ae0
```

Enabling Traffic Engineering for the Routing Protocol

Enable traffic engineering for the routing protocol (either OSPF or IS-IS) on the loopback address and on the core interface addresses.

1. Enable traffic engineering for the routing protocol:

```
[edit protocols]
user@switch# set ospf traffic-engineering
```

Enabling MPLS and Applying MPLS to the Core Interfaces

Enable MPLS within the protocols stanza and apply it to the core interfaces.

1. Configure MPLS on the core interface addresses:

```
[edit protocols]
user@switch# set mpls interface ge-0/0/5.0
user@switch# set mpls interface ge-0/0/6.0
user@switch# set mpls interface ae0
```

Enabling RSVP and Applying It to the Loopback and Core Interfaces

Enable RSVP and apply it to the loopback and the core interfaces.

1. Configure RSVP on the loopback address and the core interface addresses:

```
[edit protocols]
user@switch# set rsvp interface lo0.0
user@switch# set rsvp interface ge-0/0/5.0
user@switch# set rsvp interface ge-0/0/6.0
user@switch# set rsvp interface ae0
```

Configuring IP Addresses for the Loopback and Core Interfaces

Configure IP addresses for the loopback and core interfaces.

1. Configure an IP address for the loopback interface and for the core interfaces:

```
[edit]
user@switch# set interfaces lo0 unit 0 family inet address 127.1.1.1/32
user@switch# set interfaces ge-0/0/5 unit 0 family inet address 10.1.5.1/24
user@switch# set interfaces ge-0/0/6 unit 0 family inet address 10.1.6.1/24
user@switch# set interfaces ae0 unit 0 family inet address 10.1.9.2/24
```

Enabling Family MPLS on the Core Interfaces

On the logical unit of the core interface addresses, configure these interfaces to belong to family mpls. This configuration identifies the interfaces used for forwarding MPLS packets.

1. Configure family mpls on the logical units of the core interfaces:

```
[edit]
user@switch# set interfaces ge-0/0/5 unit 0 family mpls
user@switch# set interfaces ge-0/0/6 unit 0 family mpls
user@switch# set interfaces ae0 unit 0 family mpls
```



NOTE: You can enable family mpls on either individual interfaces or aggregated Ethernet interfaces. You cannot enable it on tagged VLAN interfaces.

- Related Topics**
- Example: Configuring MPLS on EX Series Switches on page 11
 - Configuring MPLS on Provider Edge Switches (CLI Procedure) on page 30
 - Configuring an OSPF Network (J-Web Procedure)
 - Verifying That MPLS Is Working Correctly on page 39

- Understanding JUNOS MPLS Components for EX Series Switches on page 6
- For information on the interface statement for OSPF, see the *JUNOS Software Routing Protocols Configuration Guide* at <http://www.juniper.net/techpubs/software/junos/junos95/index.html>.

Chapter 4

Verifying MPLS

- Verifying That MPLS Is Working Correctly on page 39

Verifying That MPLS Is Working Correctly

To verify that MPLS is working correctly on EX Series switches, perform the following tasks:

1. Verifying the Physical Layer on the Switches on page 39
2. Verifying the Routing Protocol on page 40
3. Verifying the Core Interfaces Being Used for the MPLS Traffic on page 40
4. Verifying RSVP on page 40
5. Verifying the Assignment of Interfaces for MPLS Label Operations on page 41
6. Verifying the Status of the CCC on page 41

Verifying the Physical Layer on the Switches

Purpose Verify that the interfaces are up. Perform this verification task on each of the switches.

Action user@switch> **show interfaces ge- terse**

Interface	Admin	Link	Proto	Local	Remote
ge-0/0/0	up	up			
ge-0/0/0.0	up	up			
ge-0/0/1.0	up	up	ccc		
ge-0/0/2.0	up	up	ccc		
ge-0/0/3.0	up	up	eth-switch		
ge-0/0/4.0	up	up	eth-switch		
ge-0/0/5.0	up	up	inet	10.1.5.1/24	
mpls					
ge-0/0/6.0	up	up	inet	10.1.6.1/24	
mpls					

Meaning The `show interfaces terse` command displays status information about the Gigabit Ethernet interfaces on the switch. This output verifies that the interfaces are **up**. The output for the protocol family (**Proto** column) shows that interfaces `ge-0/0/1.0` and `ge-0/0/2.0` are configured as circuit cross-connect. The **Local** and **Remote** columns do not display IP addresses, because the **inet** family is not configured for CCC interfaces. The output for the protocol family of the core interfaces (`ge-0/0/0.5` and

ge-0/0/0.6), shows that these interfaces are configured as both `inet` and `mpls`. The `Local` column for the core interfaces shows the IP address configured for these interfaces.

Verifying the Routing Protocol

Purpose Verify the state of the configured routing protocol. You should perform this verification task on each of the switches. The state should be `Full`. If you have configured OSPF as the routing protocol, use the `show ospf neighbor` command to verify that the routing protocol is communicating with the switch neighbors. If you have configured IS-IS as the routing protocol, use the `show isis adjacency` command to verify that the routing protocol is communicating with the switch neighbors.

Action `user@switch> show ospf neighbor`

Address	Interface	State	ID	Pri	Dead
127.1.1.2	ge-0/0/5	Full	10.10.10.10	128	39

Meaning The `show ospf neighbor` command displays the status of the routing protocol that has been configured on this switch. The output shows that the state is `full`, meaning that the routing protocol is operating correctly—that is, hello packets are being exchanged between directly connected neighbors. For additional information on checking and monitoring routing protocols, see the *JUNOS Software Routing Protocols and Policies Command Reference* at <http://www.juniper.net/techpubs/software/junos/junos95/index.html>.

Verifying the Core Interfaces Being Used for the MPLS Traffic

Purpose Verify that the state of the MPLS interface is `Up`. You should perform this verification task on each of the switches.

Action `user@switch> show mpls interface`

Interface	State	Administrative groups
ge-0/05	Up	<none>

Meaning The `show mpls interface` command displays the status of the core interfaces that have been configured to belong to family `mpls`. This output shows that the interface configured to belong to family `mpls` is up.

Verifying RSVP

Purpose Verify the state of the RSVP session. You should perform this verification task on each of the switches.

Action `user@switch> show rsvp session`

```
Ingress RSVP: 1 sessions
To          From          State  Rt Style Labelin Labelout LSPname
127.1.1.3   127.1.1.1       Up     0  1 FF      -   300064
lsp_to_pe2_ge1
Total 1 displayed, Up 1, Down 0
```

```
Egress RSVP: 1 sessions
To          From          State Rt Style Labelin Labelout LSPname
127.1.1.1   127.1.1.3   Up    0  1 FF 299968   -
lsp_to_pe1_ge1
Total 1 displayed, Up 1, Down 0
```

```
Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Meaning This output confirms that the RSVP sessions are Up.

Verifying the Assignment of Interfaces for MPLS Label Operations

Purpose Verify which interface is being used as the beginning of the CCC and which interface is being used to push the MPLS packet to the next hop. You should perform this task only on the provider edge switches.

Action user@switch> **show route forwarding-table family mpls**

```
MPLS:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0
0                user  0                recv  49   3
1                user  0                recv  49   3
2                user  0                recv  49   3
299776           user  0                Pop   541  2 ge-0/0/1.0
ge-0/0/1.0 (CCC) user  0 127.1.2.1         Push 299792 540 2 ge-0/0/5.0
```

Meaning This output shows that CCC has been set up on interface ge-0/0/1.0. The switch receives ingress traffic on ge-0/0/1.0 with label 299776. It pops that label and swaps it to label 299792, which it pushes out on interface ge-0/0/5.0.

Verifying the Status of the CCC

Purpose Verify the status of the CCC. You should perform this task only on the provider edge switches.

Action user@switch> **show connections**

```
CCC and TCC connections [Link Monitoring On]
Legend for status (St)
UN -- uninitialized
NP -- not present
WE -- wrong encapsulation
DS -- disabled
Dn -- down
-> -- only outbound conn is up
<- -- only inbound conn is up
Up -- operational
RmtDn -- remote CCC down
Restart -- restarting

Legend for connection types
if-sw: interface switching
rmt-if: remote interface switching
lsp-sw: LSP switching
tx-p2mp-sw: transmit P2MP switching
rx-p2mp-sw: receive P2MP switching

Legend for circuit types
intf -- interface
tlsp -- transmit LSP
rlsp -- receive LSP
```

```
Connection/Circuit      Type      St      Time last up      # Up trans
ge1-to-pe2              rmt-if    Up      Feb 17 05:00:09   1
  ge-0/0/1.0             intf      Up
  lsp_to_pe1_ge1         tlsp     Up
  lsp_to_pe2_ge1         rlsp     Up
```

Meaning The `show connections` command displays the status of the CCC connections. This output verifies that the CCC interface and its associated transmit and receive LSPs are Up.

- Related Topics**
- [Configuring MPLS on Provider Edge Switches \(CLI Procedure\) on page 30](#)
 - [Configuring MPLS on Provider Switches \(CLI Procedure\) on page 35](#)

Chapter 5

Configuration Statements for MPLS

- [edit protocols] Configuration Statement Hierarchy on page 43

[edit protocols] Configuration Statement Hierarchy

```
protocols {
  connections {
    remote-interface-switch connection-name {
      interface interface-name.unit-number;
      transmit-lsp label-switched-path;
      receive-lsp label-switched-path;
    }
  }
  dot1x {
    authenticator {
      authentication-profile-name profile-name;
      interface (all | [interface-names ]) {
        disable;
        guest-vlan ( vlan-id | vlan-name );
        mac-radius <restrict>;
        maximum-requests number;
        no-reauthentication;
        quiet-period seconds;
        reauthentication {
          interval seconds;
        }
        retries number;
        server-fail (deny | permit | use-cache | vlan-id | vlan-name);
        server-reject-vlan (vlan-id | vlan-name);
        server-timeout seconds;
        supplicant (multiple | single | single-secure);
        supplicant-timeout seconds;
        transmit-period seconds;
      }
      static mac-address {
        interface interface-name;
        vlan-assignment (vlan-id | vlan-name);
      }
    }
  }
  gvrp {
    <enable | disable>;
    interface (all | [interface-name]) {
      disable;
    }
  }
}
```

```

    }
    join-timer milliseconds;
    leave-timer milliseconds;
    leaveall-timer milliseconds;
  }
  igmp-snooping {
    traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable>
        <match regex>;
      flag flag (detail | disable | receive | send);
    }
    vlan (vlan-id | vlan-number) {
      data-forwarding {
        source {
          groups group-prefix;
        }
        receiver {
          source-vlans vlan-list;
          install ;
        }
      }
    }
    disable {
      interface interface-name
    }
    immediate-leave;
    interface interface-name {
      group-limit limit;
      multicast-router-interface;
      static {
        group ip-address;
      }
    }
    proxy ;
    query-interval seconds;
    query-last-member-interval seconds;
    query-response-interval seconds;
    robust-count number;
  }
}
lldp {
  disable;
  advertisement-interval seconds;
  hold-multiplier number;
  interface (all | interface-name) {
    disable;
  }
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>
      <match regex>;
    flag flag (detail | disable | receive | send);
  }
}
lldp-med {
  disable;
  fast-start number;
  interface (all | interface-name) {

```

```

disable;
location {
  elin number;
  civic-based {
    what number;
    country-code code;
    ca-type {
      number {
        ca-value value;
      }
    }
  }
}
}
}
}
}
mpls {
  interface ( all | interface-name );
  label-switched-path lsp-name to remote-provider-edge-switch;
  path destination {
    <address | hostname> <strict | loose>
  }
}
mstp {
  disable;
  bpdu-block-on-edge;
  bridge-priority priority;
  configuration-name name;
  forward-delay seconds;
  hello-time seconds;
  interface (all | interface-name) {
    disable;
    bpdu-timeout-action {
      block;
      alarm;
    }
    cost cost;
    edge;
    mode mode;
    no-root-port;
    priority priority;
  }
  max-age seconds;
  max-hops hops;
  msti msti-id {
    vlan (vlan-id | vlan-name);
    interface interface-name {
      disable;
      cost cost;
      edge;
      mode mode;
      priority priority;
    }
  }
}
revision-level revision-level;
traceoptions {
  file filename <files number > <size size> <no-stamp | world-readable |
  no-world-readable>;
}

```

```

        flag flag;
    }
}
mvrp {
    disable
    interface (all | interface-name) {
        disable;
        join-timer milliseconds;
        leave-timer milliseconds;
        leaveall-timer milliseconds;
        registration (forbidden | normal);
    }
    no-dynamic-vlan;
    traceoptions {
        file filename <files number > <size size> <no-stamp | world-readable |
            no-world-readable>;
        flag flag;
    }
}
oam {
    ethernet{
        link-fault-management {
            action-profile profile-name;
            action {
                syslog;
                link-down;
            }
            event {
                link-adjacency-loss;
                link-event-rate;
                frame-error count;
                frame-period count;
                frame-period-summary count;
                symbol-period count;
            }
            interface interface-name {
                link-discovery (active | passive);
                pdu-interval interval;
                event-thresholds threshold-value;
                remote-loopback;
                event-thresholds {
                    frame-error count;
                    frame-period count;
                    frame-period-summary count;
                    symbol-period count;
                }
            }
            negotiation-options {
                allow-remote-loopback;
                no-allow-link-events;
            }
        }
    }
}
rstp {
    disable;
    bpdu-block-on-edge;
}

```

```

bridge-priority priority;
forward-delay seconds;
hello-time seconds;
interface (all | interface-name) {
  disable;
  bpdu-timeout-action {
    block;
    alarm;
  }
  cost cost;
  edge;
  mode mode;
  no-root-port;
  priority priority;
}
max-age seconds;
}
traceoptions {
  file filename <files number > <size size> <no-stamp | world-readable |
  no-world-readable>;
  flag flag;
}
}
sflow {
  collector {
    ip-address;
    udp-port port-number;
  }
  disable;
  interfaces interface-name {
    disable;
    polling-interval seconds;
    sample-rate number;
  }
  polling-interval seconds;
  sample-rate number;
}
stp {
  disable;
  bridge-priority priority;
  forward-delay seconds;
  hello-time seconds;
  interface (all | interface-name) {
    disable;
    bpdu-timeout-action {
      block;
      alarm;
    }
    cost cost;
    edge;
    mode mode;
    no-root-port;
    priority priority;
  }
  max-age seconds;
}

```

```

    traceoptions {
      file filename <files number > <size size> <no-stamp | world-readable |
        no-world-readable>;
      flag flag;
    }
  vstp {
    bpdu-block-on-edge;
    disable;
    force-version stp;
    vlan (all | vlan-id | vlan-name) {
      bridge-priority priority;
      forward-delay seconds;
      hello-time seconds;
      interface (all | interface-name) {
        bpdu-timeout-action {
          alarm;
          block;
        }
        cost cost;
        disable;
        edge;
        mode mode;
        no-root-port;
        priority priority;
      }
      max-age seconds;
      traceoptions {
        file filename <files number > <size size> <no-stamp | world-readable |
          no-world-readable>;
        flag flag;
      }
    }
  }
}

```

- Related Topics**
- 802.1X for EX Series Switches Overview
 - Example: Configure Automatic VLAN Administration Using GVRP
 - Understanding MAC RADIUS Authentication on EX Series Switches
 - Understanding Server Fail Fallback and 802.1X Authentication on EX Series Switches
 - IGMP Snooping on EX Series Switches Overview
 - Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches
 - Understanding MSTP for EX Series Switches
 - Understanding RSTP for EX Series Switches
 - Understanding STP for EX Series Switches
 - Understanding How to Use sFlow Technology for Network Monitoring on an EX Series Switch
 - Understanding VSTP for EX Series Switches
 - Understanding Ethernet OAM Link Fault Management for an EX Series Switch

connections

Syntax	connections { 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 protocols]
Release Information	Statement introduced in JUNOS Release 9.5 for EX Series switches.
Description	Define the connection between two circuits in a CCC connection. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Topics	<ul style="list-style-type: none"> ■ Example: Configuring MPLS on EX Series Switches on page 11 ■ Configuring MPLS on Provider Edge Switches (CLI Procedure) on page 30 ■ <i>JUNOS Software MPLS Applications Configuration Guide</i> at http://www.juniper.net/techpubs/software/junos/junos95/index.html

family ccc

Syntax	family ccc;
Hierarchy Level	[edit interfaces <i>ge-chassis/slot/port</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced in JUNOS Release 9.5 for EX Series switches.
Description	Configure the logical interface as a circuit cross-connect (CCC).
Default	You must configure a logical interface to be able to use the physical device.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<ul style="list-style-type: none"> ■ Example: Configuring MPLS on EX Series Switches on page 11 ■ Configuring MPLS on Provider Edge Switches (CLI Procedure) on page 30 ■ <i>JUNOS Software MPLS Applications Configuration Guide</i> at http://www.juniper.net/techpubs/software/junos/junos95/index.html

family mpls

Syntax	family mpls;
Hierarchy Level	[edit interfaces <i>ge-chassis/slot/port</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced in JUNOS Release 9.5 for EX Series switches.
Description	Configure MPLS protocol family information for the logical interface.
Default	You must configure a logical interface to be able to use the physical device.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<ul style="list-style-type: none"> ■ Example: Configuring MPLS on EX Series Switches on page 11 ■ Configuring MPLS on Provider Edge Switches (CLI Procedure) on page 30 ■ Configuring MPLS on Provider Switches (CLI Procedure) on page 35 ■ <i>JUNOS Software MPLS Applications Configuration Guide</i> at http://www.juniper.net/techpubs/software/junos/junos95/index.html

interface

Syntax	interface (all <i>interface-name</i>);
Hierarchy Level	[edit protocols mpls]
Release Information	Statement introduced in JUNOS Release 9.5 for EX Series switches.
Description	Enable MPLS on all interfaces on the switch or on the specified interface.
Default	MPLS is disabled.
Options	all—All interfaces on the switch. <i>interface-name</i> —Name of an interface in one of these formats: <ul style="list-style-type: none"> ■ aex —Name of an aggregated Ethernet interface. ■ <i>ge-chassis/slot/port</i> —Name of a Gigabit Ethernet interface.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Topics	<ul style="list-style-type: none"> ■ Example: Configuring MPLS on EX Series Switches on page 11 ■ Configuring MPLS on Provider Edge Switches (CLI Procedure) on page 30 ■ Configuring MPLS on Provider Switches (CLI Procedure) on page 35

label-switched-path

Syntax	label-switched-path <i>lsp-name</i> to <i>remote-provider-edge-switch</i> ;
Hierarchy Level	[edit protocols mpls]
Release Information	Statement introduced in JUNOS Release 9.5 for EX Series switches.
Description	Define a label-switched path (LSP) to the remote provider edge switch to use for MPLS traffic. You must specify this statement on the provider edge switch.
Options	<p><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 on the ingress switch.</p> <p><i>remote-provider-edge-switch</i> —Either the loopback address or the switch address.</p>
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	<ul style="list-style-type: none"> ■ Example: Configuring MPLS on EX Series Switches on page 11 ■ Configuring MPLS on Provider Edge Switches (CLI Procedure) on page 30 ■ <i>JUNOS Software MPLS Applications Configuration Guide</i> at http://www.juniper.net/techpubs/software/junos/junos95/index.html

mpls

Syntax mpls {
 interface (all | *interface-name*);
 label-switched-path *lsp-name* to *remote-provider-edge-switch*;
 path *destination* {
 <*address* | *hostname*> <strict | loose>
 }
 }

Hierarchy Level [edit protocols]

Release Information Statement introduced in JUNOS Release 9.5 for EX Series switches.

Description Enable MPLS on the switch.

The remaining statements are explained separately.

Default MPLS is disabled.

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

- Related Topics**
- Example: Configuring MPLS on EX Series Switches on page 11
 - Configuring MPLS on Provider Edge Switches (CLI Procedure) on page 30
 - Configuring MPLS on Provider Switches (CLI Procedure) on page 35
 - *JUNOS Software MPLS Applications Configuration Guide* at <http://www.juniper.net/techpubs/software/junos/junos95/index.html>

path

Syntax	<pre>path destination { <address hostname> <strict loose> }</pre>
Hierarchy Level	[edit protocols mpls]
Release Information	Statement introduced in JUNOS Release 9.5 for EX Series switches.
Description	Configure path protection on your MPLS network.
Options	<p><i>destination</i> —Name of a label switched path (LSP). In addition to specifying the name of the configured LSP, you can include some other designation such as <i>primary-path</i>.</p> <p><i>address</i> —(Optional) IP address of each transit switch (or the IP address of the loopback interface on the switch) in the LSP. If you want to control exactly which switches are selected for the LSP, specify the address or hostname of each transit switch. Specify the addresses in order, starting with the first provider (transit) switch, and continuing sequentially along the path until reaching the egress provider edge switch. Default: If you do not specify the addresses or hostnames of any switches, the LSP is calculated by the switch.</p> <p><i>hostname</i> —(Optional) See <i>address</i> . Default: If you do not specify the addresses or hostnames of any switches, the LSP is calculated by the switch.</p> <p><i>loose</i>—(Optional) Indicates that the next address in the path statement is a loose link. This means that the LSP can traverse through other switches before reaching this switch. Default: strict</p> <p><i>strict</i>—(Optional) Indicates that the LSP must go to the next address specified in the path statement without traversing other switches. This is the default.</p>
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	<ul style="list-style-type: none"> ■ Configuring Path Protection in an MPLS Network (CLI Procedure) on page 27

primary

Syntax	<code>primary path-name;</code>
Hierarchy Level	[edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced in JUNOS Release 9.5 for EX Series switches.
Description	Specify the primary path to use for a label switched path (LSP). You can configure only one primary path.
Options	<i>path-name</i> —Name of the primary path that you created with the <code>path</code> statement.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Topics	■ Configuring Path Protection in an MPLS Network (CLI Procedure) on page 27

remote-interface-switch

Syntax remote-interface-switch *connection-name* {
 interface *interface-name.unit-number*;
 receive-lsp *label-switched-path*;
 transmit-lsp *label-switched-path*;
 }

Hierarchy Level [edit protocols connections]

Release Information Statement introduced in JUNOS Release 9.5 for EX Series switches.

Description Configure MPLS LSP tunnel cross-connects. This makes an association between a CCC interface and two LSPs, one for transmitting MPLS packets from the local provider edge switch to the remote provider edge switch and the other for receiving MPLS packets on the local provider edge switch from the remote provider edge switch.

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 of the CCC interface.

receive-lsp label-switched-path —Name of the LSP from the connection's source. This LSP name was specified by the *label-switched-path* statement on the remote provider edge switch in the *protocols mpls* stanza.

transmit-lsp label-switched-path —Name of the LSP to the connection's destination. This LSP name was specified by the *label-switched-path* statement on the local provider edge switch in the *protocols mpls* stanza.

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

- Related Topics**
- Example: Configuring MPLS on EX Series Switches on page 11
 - Configuring MPLS on Provider Edge Switches (CLI Procedure) on page 30
 - *JUNOS Software MPLS Applications Configuration Guide* at <http://www.juniper.net/techpubs/software/junos/junos95/index.html>

revert-timer

Syntax	revert-timer <i>seconds</i> ;
Hierarchy Level	[edit protocols mpls], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced in JUNOS Release 9.5 for EX Series switches.
Description	Specify the amount of time that a label switched path (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. If you have configured a value of 0 seconds for the <code>revert-timer</code> 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.
Default	60 seconds
Options	<i>seconds</i> —Value in seconds. Range: 0 through 65,535 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Topics	<ul style="list-style-type: none"> ■ Configuring Path Protection in an MPLS Network (CLI Procedure) on page 27

rsvp

Syntax	rsvp;
Hierarchy Level	[edit protocols]
Release Information	Statement introduced in JUNOS Release 9.5 for EX Series switches.
Description	Enable Resource Reservation Protocol (RSVP) signaling. The primary purpose of RSVP in JUNOS Software for EX Series switches is to support dynamic signaling within label switched paths (LSPs).
Default	RSVP is disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Topics	<ul style="list-style-type: none"> ■ Example: Configuring MPLS on EX Series Switches on page 11 ■ Configuring MPLS on Provider Edge Switches (CLI Procedure) on page 30 ■ Configuring MPLS on Provider Switches (CLI Procedure) on page 35 ■ <i>JUNOS Software MPLS Applications Configuration Guide</i> at http://www.juniper.net/techpubs/software/junos/junos95/index.html

secondary

Syntax	<code>secondary path-name { standby; }</code>
Hierarchy Level	[edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced in JUNOS Release 9.5 for EX Series switches.
Description	Specify one or more secondary paths to use for the label switched path (LSP). You can configure more than one secondary path. All secondary paths are equal, and the first one that is available is chosen.
Options	<i>path-name</i> —Name of a secondary path that you created with the <code>path</code> statement. The remaining statement is explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Topics	■ Configuring Path Protection in an MPLS Network (CLI Procedure) on page 27

standby

Syntax	<code>standby;</code>
Hierarchy Level	[edit protocols mpls label-switched-path <i>lsp-name</i> secondary <i>path-name</i>]
Release Information	Statement introduced in JUNOS Release 9.5 for EX Series switches.
Description	Enable the path to remain up at all times to provide instant switchover if connectivity problems occur.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Topics	■ Configuring Path Protection in an MPLS Network (CLI Procedure) on page 27

traffic-engineering

Syntax	traffic-engineering;
Hierarchy Level	[edit protocols ospf] [edit protocols isis]
Release Information	Statement introduced in JUNOS Release 9.5 for EX Series switches.
Description	Enable the traffic-engineering features of the specified routing protocol.
Default	Traffic-engineering support is disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Topics	<ul style="list-style-type: none">■ Example: Configuring MPLS on EX Series Switches on page 11■ Configuring MPLS on Provider Edge Switches (CLI Procedure) on page 30■ Configuring MPLS on Provider Switches (CLI Procedure) on page 35■ Configuring an OSPF Network (J-Web Procedure)■ <i>JUNOS Software MPLS Applications Configuration Guide</i> at http://www.juniper.net/techpubs/software/junos/junos95/index.html

Chapter 6

Operational Mode Commands for MPLS

show connections

Syntax show connections
 <brief | extensive>
 <all | remote-interface-switch>
 <down | up | up-down>
 <history>
 <labels>
 <name>
 <status>

Release Information Command introduced in JUNOS Release 9.5 for EX Series switches.

Description Display information about the configured circuit cross-connect (CCC) connections.

Options none—Display the standard level of output for all configured CCC connections on all logical systems.

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

all—(Optional) Display all connections.

down | up | up-down—(Optional) Display nonoperational, operational, or both kinds of connections.

history—(Optional) Display information about connection history.

labels—(Optional) Display labels used for transmit and receive LSPs.

name —(Optional) Display information about the specified connection only.

remote-interface-switch—(Optional) Display remote interface switch connections only.

name —(Optional) Display information about the specified connection only.

status—(Optional) Display information about the connection and interface status.

Required Privilege Level view

- Related Topics**
- Example: Configuring MPLS on EX Series Switches on page 11
 - Configuring MPLS on Provider Edge Switches (CLI Procedure) on page 30
 - connections

List of Sample Output

- show connections on page 61
- show connections brief on page 61
- show connections down on page 62
- show connections extensive on page 62
- show connections history on page 62
- show connections labels on page 62
- show connections <name> on page 62

show connections remote-interface-switch on page 62

show connections status on page 63

Output Fields Table 4 on page 61 describes the output fields for the show connections command. Output fields are listed in the approximate order in which they appear.

Table 4: show connections Output Fields

Field Name	Field Description
CCC and TCC connections [Link Monitoring On Off]	Whether link monitoring is enabled: On or Off.
Legend for Status (St)	Connection or circuit status. See the output's legend for an explanation of the status field values.
Legend for connection types	Type of connection: <ul style="list-style-type: none"> ■ if-sw—Layer 2 switching cross-connect. ■ rmt-if—Remote interface switch. While graceful restart is in progress, rmt-if will display a state (St) of Restart.
Legend for circuit types	Type of circuit: <ul style="list-style-type: none"> ■ intf—Interface circuit. ■ t1sp—Transmit LSP circuit. ■ r1sp—Receive LSP circuit.
Connection/Circuit	Name of the configured CCC connection.
Type	Type of connection.
St	State of the connection.
Time last up	Time that the connection or circuit last transitioned to the Up (operational) state.
# Up trans	Number of times that the connection or circuit has transitioned to the Up (operational) state.

show connections user@switch> show connections

```

Connection/Circuit      Type      St      Time last up      # Up trans
ge1-to-pe2              rmt-if    Up      Jun 26 18:37:25
1
ge-0/0/5.0              intf      Up
1sp_pe1_to_ge1_pe2     t1sp     Up
1sp_pe2_to_ge1_pe1     r1sp     Up

```

show connections brief user@switch> show connections brief

```

Connection/Circuit      Type      St      Time last up      # Up trans
ge-1_to_pe2            rmt-if    Up      Jan 29 13:07:56
1

```

show connections down user@switch> **show connections down**
No matching connections found.

show connections extensive user@switch> **show connections extensive**

Connection/Circuit	Type	St	Time last up	# Up trans
ge1_to_pe2	rmt-if	Up	Jan 29 13:07:56	
1				
ge-0/0/5.0	intf	Up		
lsp_pe1_to_ge1_pe2	tlsp	Up		
lsp_pe2_to_ge1_pe1	rlsp	Up		
Incoming labels: 299776				
Outgoing labels: Push 300112				

show connections history user@switch> **show connections history**

Connection/Circuit	Type	St	Time last up	# Up trans
ge1-to-pe2	rmt-if	Up	Jan 29 13:07:56	
1				
				# Paths
Time	Event	Interface/Label		Rcv Xmt
Jan 29 13:07:56	CCC status update			1 1
Jan 29 13:07:55	TLSP up	300112@1:0, 1		1 1
Jan 29 13:07:55	TLSP down	300112@1		1 0
Jan 29 13:07:55	TLSP up	300112@1:0, 4097		1 1
Jan 29 13:07:54	RLSP up	299776		1 0
Jan 29 13:01:08	Remote CCC down			0 0
Jan 29 13:01:08	Interface up	ge-0/0/0.10		0 0
Jan 29 13:01:06	Interface down	ge-0/0/0.10		0 0
Jan 29 13:01:04	Remote CCC down			0 0
Jan 29 13:01:02	Interface down			0 0

show connections labels user@switch> **show connections labels**

Connection/Circuit	Type	St	Time last up	# Up trans
ge1-to-pe2	rmt-if	RmtDn	Jun 26 18:37:25	
1				
Incoming labels: 299776				
Outgoing labels: Push 299792				

show connections <name> user@switch> **show connections ge1-to-pe2**

Connection/Circuit	Type	St	Time last up	# Up trans
ge1_to_pe2	rmt-if	Up	Jan 29 13:07:56	
1				
ge-0/0/5.0	intf	Up		
lsp_pe1_to_ge1_pe2	tlsp	Up		
lsp_pe2_to_ge1_pe1	rlsp	Up		

show connections remote-interface-switch

Connection/Circuit	Type	St	Time last up	# Up trans
xcon10_ge0_to_239	rmt-if	Up	Jan 29 13:07:56	
1				
ge-0/0/0.10	intf	Up		
lsp_to_240_10	tlsp	Up		
lsp_to_239_10	rlsp	Up		
xcon11_ge0_to_239	rmt-if	Up	Jan 29 13:07:57	
1				
ge-0/0/0.11	intf	Up		
lsp_to_240_11	tlsp	Up		
lsp_to_239_11	rlsp	Up		

```

show connections user@switch> show connections status
status Connection/Circuit      Type      St      Time last up      # Up trans
xcon10_ge0_to_239      rmt-if    Up
1
  ge-0/0/0.10          intf     Up
  lsp_to_240_10        t1sp    Up
  lsp_to_239_10        r1sp    Up
xcon11_ge0_to_239      rmt-if    Up      Jan 29 13:07:57
1
  ge-0/0/0.11          intf     Up
  lsp_to_240_11        t1sp    Up
  lsp_to_239_11        r1sp    Up

```

show route forwarding-table

Syntax show route forwarding-table
 <detail | extensive | summary>
 <ccc *ccc-interface-name*>
 <destination>
 <family *family-name*>
 <label *label*>
 <matching *ip_prefix*>
 <multicast>
 <vpn *vpn*>

Release Information Command introduced in JUNOS Release 9.5 for EX Series switches.

Description Display the Routing Engine's forwarding table, including the network-layer prefixes and their next hops. This command is used to help verify that the routing protocol process has relayed the correction information to the forwarding table. The Routing Engine constructs and maintains one or more routing tables. From the routing tables, the Routing Engine derives a table of active routes, called the forwarding table.

Options none—Display the routes in the forwarding table.

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

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

destination —(Optional) Display the destination prefix.

family *family-name* —(Optional) Display routing table entries for the specified family: inet, inet6, iso, mpls.

label *label* —(Optional) Display route entries for the specified label name.

matching *ip_prefix* —(Optional) Display route entries for the specified IP prefix.

multicast—(Optional) Display route entries for multicast routes.

vpn *vpn* —(Optional) Display route entries for the specified VPN.

Required Privilege Level view

- Related Topics**
- Example: Configuring MPLS on EX Series Switches on page 11
 - Configuring MPLS on Provider Edge Switches (CLI Procedure) on page 30
 - Configuring MPLS on Provider Switches (CLI Procedure) on page 35

List of Sample Output

- show route forwarding-table on page 66
- show route forwarding-table summary on page 67
- show route forwarding-table extensive on page 67
- show route forwarding-table ccc on page 68

show route forwarding-table family on page 68
 show route forwarding-table label on page 69
 show route forwarding-table matching on page 69
 show route forwarding-table multicast on page 70

Output Fields Table 5 on page 65 lists the output fields for the `show route forwarding-table` command. Output fields are listed in the approximate order in which they appear. Field names might be abbreviated (as shown in parentheses) when no level of output is specified or when the `detail` keyword is used instead of the `extensive` keyword.

Table 5: show route forwarding-table Output Fields

Field Name	Field Description	Level of Output
Routing table	Name of the routing table (for example, inet, inet6, mpls).	All levels
Address family	Address family (for example, IP, IPv6, ISO, MPLS).	All levels
Destination	Destination of the route.	detail, extensive
Route Type (Type)	How the route was placed into the forwarding table. When the <code>detail</code> keyword is used, the route type might be abbreviated (as shown in parentheses): <ul style="list-style-type: none"> ■ <code>cloned (clon)</code>—(TCP or multicast only) Cloned route. ■ <code>destination (dest)</code>—Remote addresses directly reachable through an interface. ■ <code>destination down (iddn)</code>—Destination route for which the interface is unreachable. ■ <code>interface cloned (ifcl)</code>—Cloned route for which the interface is unreachable. ■ <code>route down (ifdn)</code>—Interface route for which the interface is unreachable. ■ <code>ignore (ignr)</code>—Ignore this route. ■ <code>interface (intf)</code>—Installed as a result of configuring an interface. ■ <code>permanent (perm)</code>—Routes installed by the kernel when the routing table is initialized. ■ <code>user</code>—Routes installed by the routing protocol process or as a result of the configuration. 	All levels
Route reference (RtRef)	Number of routes to reference.	detail, extensive
Flags	Route type flags: <ul style="list-style-type: none"> ■ <code>none</code>—No flags are enabled. ■ <code>accounting</code>—Route has accounting enabled. ■ <code>cached</code>—Cache route. ■ <code>incoming-iface interface-number</code> —Check against incoming interface. ■ <code>prefix load balance</code>—Load balancing is enabled for this prefix. ■ <code>sent to PFE</code>—Route has been sent to the Packet Forwarding Engine. ■ <code>static</code>—Static route. 	extensive
Nexthop	IP address of the next hop to the destination.	detail, extensive

Table 5: show route forwarding-table Output Fields (continued)

Field Name	Field Description	Level of Output
Next hop type (Type)	<p>Next-hop type. When the detail keyword is used, the next-hop type might be abbreviated (as indicated in parentheses):</p> <ul style="list-style-type: none"> ■ broadcast (bcst)—Broadcast. ■ deny—Deny. ■ hold—Next hop is waiting to be resolved into a unicast or multicast type. ■ indexed (idxd)—Indexed next hop. ■ indirect (indr)—Indirect next hop. ■ local (locl)—Local address on an interface. ■ routed multicast (mcr)—Regular multicast next hop ■ multicast (mcr)—Wire multicast next hop (limited to the LAN). ■ multicast discard (mdsc)—Multicast discard. ■ multicast group (mgrp) —Multicast group member. ■ receive (rcv)—Receive. ■ reject (rjct) Discard. An ICMP unreachable message was sent. ■ resolve (rslv)—Resolving the next hop. ■ unicast (ucst)—Unicast. ■ unilist (ulst)—List of unicast next hops. A packet sent to this next hop goes to any next hop in the list. 	detail, extensive
Index	Software index of the next hop that is used to route the traffic for a given prefix.	detail, extensive none
Route interface-index	Logical interface index from which the route is learned. For example, for interface routes, this is the logical interface index of the route itself. For static routes, this field is zero. For routes learned through routing protocols, this is the logical interface index from which the route is learned.	extensive
Reference (NhRef)	Number of routes that refer to this next hop.	none detail, extensive
Next-hop interface (Netif)	Interface used to reach the next hop.	none detail, extensive
Alternate forward nh index	Index number of the alternate next hop interface. Seen with multicast option only.	extensive
Next-hop L3 Interface	The next hop layer 3 interface. This option can be expressed as a VLAN name and is only seen with the multicast option.	extensive
Next-hop L2 Interfaces	The next hop layer 2 interfaces. Seen with multicast option only.	extensive

```

show route forwarding-table user@switch> show route forwarding-table
Routing table: default.inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          user   2 0:12:f2:21:cf:0  ucst  333   5 me0.0
default          perm   0                   rjct   36    2
0.0.0.0/32       perm   0                   dscd   34    1
    
```

```

2.2.2.0/24      intf      0                rslv 1309      1 ae0.0
2.2.2.0/32      dest      0 2.2.2.0        recv 1307      1 ae0.0
2.2.2.1/32      dest      0 0:21:59:cc:89:c0 ucst 1320      1 ae0.0
2.2.2.2/32      intf      0 2.2.2.2        locl 1308       2
2.2.2.2/32      dest      0 2.2.2.2        locl 1308       2
2.2.2.255/32    dest      0 2.2.2.255      bcst 1306      1 ae0.0
3.3.3.0/24      intf      0                rslv 1313      1 ae1.0
3.3.3.0/32      dest      0 3.3.3.0        recv 1311      1 ae1.0
3.3.3.1/32      intf      0 3.3.3.1        locl 1312       2
3.3.3.1/32      dest      0 3.3.3.1        locl 1312       2
3.3.3.2/32      dest      0 0:21:59:cc:89:c1 ucst 1321      24 ae1.0
3.3.3.255/32    dest      0 3.3.3.255      bcst 1310      1 ae1.0
4.4.4.0/24      user      0 3.3.3.2        ucst 1321      24 ae1.0
8.8.8.8/32      user      0 3.3.3.2        ucst 1321      24 ae1.0
9.9.9.9/32      intf      0 9.9.9.9        locl 1280       1
10.10.10.10/32  user      0 3.3.3.2        ucst 1321      24 ae1.0
10.93.8.0/21    intf      0                rslv 323        1 me0.0
10.93.8.0/32    dest      0 10.93.8.0      recv 321        1 me0.0
10.93.13.238/32 intf      0 10.93.13.238   locl 322        2
10.93.13.238/32 dest      0 10.93.13.238   locl 322        2
10.93.15.254/32 dest      0 0:12:f2:21:cf:0 ucst 333         5 me0.0
10.93.15.255/32 dest      0 10.93.15.255   bcst 320         1 me0.0
14.14.14.0/24   ifdn      0                rslv 1319      1 ge-0/0/25.0
14.14.14.0/32   iddn      0 14.14.14.0     recv 1317      1 ge-0/0/25.0
14.14.14.2/32   user      0                rjct 36          2
14.14.14.2/32   intf      0 14.14.14.2     locl 1318       2
14.14.14.2/32   iddn      0 14.14.14.2     locl 1318       2
14.14.14.255/32 iddn      0 14.14.14.255   bcst 1316      1 ge-0/0/25.0
224.0.0.0/4     perm      1                mdsc 35          1
224.0.0.1/32    perm      0 224.0.0.1      mcst 31         3
224.0.0.5/32    user      1 224.0.0.5      mcst 31         3
255.255.255.255/32 perm      0                bcst 32         1

```

```

show route forwarding-table summary user@switch> show route forwarding-table summary
Routing table: default.inet
Internet:

```

```

    user:          6 routes
    perm:          5 routes
    intf:          8 routes
    dest:         12 routes
    ifdn:          1 routes
    iddn:          3 routes

```

```

show route forwarding-table extensive user@switch> show route forwarding-table summary
Routing table: default.inet [Index 0]
Internet:

```

```

Destination: default
Route type: user
Route reference: 2
Flags: sent to PFE, rt nh decoupled
Nexthop: 0:12:f2:21:cf:0
Next-hop type: unicast
Next-hop interface: me0.0
Route interface-index: 0
Index: 333
Reference: 5

```

```

Destination: default
Route type: permanent
Route reference: 0
Flags: none
Route interface-index: 0

```

```

Next-hop type: reject                Index: 36      Reference: 2

Destination: 0.0.0.0/32
Route type: permanent
Route reference: 0                   Route interface-index: 0
Flags: sent to PFE
Next-hop type: discard              Index: 34      Reference: 1

Destination: 2.2.2.0/24
Route type: interface
Route reference: 0                   Route interface-index: 66
Flags: sent to PFE
Next-hop type: resolve              Index: 1309    Reference: 1
Next-hop interface: ae0.0

Destination: 2.2.2.0/32
Route type: destination
Route reference: 0                   Route interface-index: 66
Flags: sent to PFE
Nexthop: 2.2.2.0
Next-hop type: receive              Index: 1307    Reference: 1
Next-hop interface: ae0.0

Destination: 2.2.2.1/32
Route type: destination
Route reference: 0                   Route interface-index: 66
Flags: sent to PFE
Nexthop: 0:21:59:cc:89:c0
Next-hop type: unicast              Index: 1320    Reference: 1
Next-hop interface: ae0.0

Destination: 2.2.2.2/32
Route type: interface
Route reference: 0                   Route interface-index: 0
Flags: sent to PFE
Nexthop: 2.2.2.2
Next-hop type: local                Index: 1308    Reference: 2

Destination: 2.2.2.2/32
Route type: destination
Route reference: 0                   Route interface-index: 66
Flags: none
Nexthop: 2.2.2.2
Next-hop type: local                Index: 1308    Reference: 2

Destination: 2.2.2.255/32
Route type: destination
Route reference: 0                   Route interface-index: 66
Flags: sent to PFE
Nexthop: 2.2.2.255
Next-hop type: broadcast            Index: 1306    Reference: 1
Next-hop interface: ae0.0

```

show route forwarding-table ccc

```

user@switch> show route forwarding-table ccc ge-0/0/0.10
Routing table: default.mpls
MPLS:
Destination      Type RtRef Next hop          Type Index NhRef Netif
ge-0/0/0.10     (CCC) user    0 3.3.3.2          Push 300112 1343 2 ae1.0

```

show route forwarding-table family

```

user@switch> show route forwarding-table family mpls

```

```

Routing table: default.mpls
MPLS:
Destination      Type RtRef Next hop      Type Index NhRef Netif
default          perm  0
0                user  0
1                user  0
2                user  0
299776           user  0      Pop      1334    2 ge-0/0/0.10
299792           user  0      Pop      1339    2 ge-0/0/0.14
299808           user  0      Pop      1341    2 ge-0/0/0.2
299824           user  0      Pop      1344    2 ge-0/0/0.11
299840           user  0      Pop      1345    2 ge-0/0/0.13
299856           user  0      Pop      1346    2 ge-0/0/0.18
299872           user  0      Pop      1347    2 ge-0/0/0.16
299888           user  0      Pop      1348    2 ge-0/0/0.7
299904           user  0      Pop      1349    2 ge-0/0/0.20
299920           user  0      Pop      1350    2 ge-0/0/0.19
299936           user  0      Pop      1351    2 ge-0/0/0.17
299952           user  0      Pop      1352    2 ge-0/0/0.9
299968           user  0      Pop      1353    2 ge-0/0/0.1
299984           user  0      Pop      1354    2 ge-0/0/0.12
300000           user  0      Pop      1355    2 ge-0/0/0.8
300016           user  0      Pop      1356    2 ge-0/0/0.4
300032           user  0      Pop      1357    2 ge-0/0/0.5
300048           user  0      Pop      1358    2 ge-0/0/0.3
300064           user  0      Pop      1359    2 ge-0/0/0.15
ge-0/0/0.1 (CCC) user  0 3.3.3.2      Push 300064 1340    2 ae1.0
ge-0/0/0.2 (CCC) user  0 3.3.3.2      Push 299872 1328    2 ae1.0
ge-0/0/0.3 (CCC) user  0 3.3.3.2      Push 299792 1323    2 ae1.0
ge-0/0/0.4 (CCC) user  0 3.3.3.2      Push 300016 1337    2 ae1.0
ge-0/0/0.5 (CCC) user  0 3.3.3.2      Push 299824 1325    2 ae1.0
ge-0/0/0.7 (CCC) user  0 3.3.3.2      Push 299920 1331    2 ae1.0
ge-0/0/0.8 (CCC) user  0 3.3.3.2      Push 299840 1326    2 ae1.0
ge-0/0/0.9 (CCC) user  0 3.3.3.2      Push 299888 1329    2 ae1.0
ge-0/0/0.10 (CCC) user  0 3.3.3.2      Push 300112 1343    2 ae1.0
ge-0/0/0.11 (CCC) user  0 3.3.3.2      Push 299776 1322    2 ae1.0
ge-0/0/0.12 (CCC) user  0 3.3.3.2      Push 299952 1333    2 ae1.0
ge-0/0/0.13 (CCC) user  0 3.3.3.2      Push 300096 1342    2 ae1.0
ge-0/0/0.14 (CCC) user  0 3.3.3.2      Push 299984 1335    2 ae1.0
ge-0/0/0.15 (CCC) user  0 3.3.3.2      Push 299936 1332    2 ae1.0
ge-0/0/0.16 (CCC) user  0 3.3.3.2      Push 299808 1324    2 ae1.0
ge-0/0/0.17 (CCC) user  0 3.3.3.2      Push 300000 1336    2 ae1.0
ge-0/0/0.18 (CCC) user  0 3.3.3.2      Push 300032 1338    2 ae1.0
ge-0/0/0.19 (CCC) user  0 3.3.3.2      Push 299904 1330    2 ae1.0
ge-0/0/0.20 (CCC) user  0 3.3.3.2      Push 299856 1327    2 ae1.0

```

```

show route forwarding-table label user@switch> show route forwarding-table label 29976

```

```

Routing table: default.mpls
MPLS:
Destination      Type RtRef Next hop      Type Index NhRef Netif
299776           user  0      Pop      1334    2 ge-0/0/0.10

```

```

show route forwarding-table matching user@switch> show route forwarding-table matching 3
Routing table: default.inet
Internet:

```

```

show route forwarding-table multicast
user@switch> show route forwarding-table multicast
Routing table: default.inet
Internet:
Destination      Type RtRef Next hop                Type Index NhRef Netif
224.0.0.0/4      perm  1
224.0.0.1/32    perm  0 224.0.0.1              mcst  31   3
224.0.0.5/32    user  1 224.0.0.5              mcst  31   3

Routing table: __master.anon__.inet
Internet:
Destination      Type RtRef Next hop                Type Index NhRef Netif
224.0.0.0/4      perm  0
224.0.0.1/32    perm  0 224.0.0.1              mcst 1285  1

Routing table: default.inet6
Internet6:
Destination      Type RtRef Next hop                Type Index NhRef Netif
ff00::/8         perm  0
ff02::1/128     perm  0 ff02::1                mcst  39   1
    
```

show mpls interface

- Syntax** show mpls interface
- Release Information** Command introduced in JUNOS Release 9.5 for EX Series switches.
- Description** Display information about MPLS-enabled interfaces. MPLS is enabled on an interface when the interface is configured with both the `set protocols mpls interface interface-name` and `set interfaces interface-name unit 0 family mpls` commands.
- Required Privilege Level** view
- Related Topics**
- Example: Configuring MPLS on EX Series Switches on page 11
 - Configuring MPLS on Provider Edge Switches (CLI Procedure) on page 30
 - Configuring MPLS on Provider Switches (CLI Procedure) on page 35
- List of Sample Output** show mpls interface on page 71
- Output Fields** Table 6 on page 71 describes the output fields for the `show mpls interface` command. Output fields are listed in the approximate order in which they appear.

Table 6: show mpls interface Output Fields

Field Name	Field Description
Interface	Name of the interface.
State	State of the interface: Up or Dn (down).
Administrative groups	Administratively assigned colors of the link.

```

show mpls interface user@switch> show mpls interface
Interface  State      Administrative groups
so-1/0/0.0 Up         Blue Yellow Red

```

show rsvp session

Syntax show rsvp session
 <brief | detail | extensive | terse>
 <bidirectional | unidirectional>
 <down | up>
 <interface *interface-name*>
 <lsp-type>
 <name *session-name*>
 <session-type>
 <statistics>
 <te-link *te-link*>

Release Information Command introduced in JUNOS Release 9.5 for EX Series switches.

Description Display information about Resource Reservation Protocol (RSVP) sessions.

Options none—Display standard information about all RSVP sessions.

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

bidirectional | unidirectional—(Optional) Display information about bidirectional or unidirectional RSVP sessions only, respectively.

down | up—(Optional) Display only LSPs that are inactive or active, respectively.

interface *interface-name* —(Optional) Display RSVP sessions for the specified interface only.

lsp-type —(Optional) Display information about RSVP sessions with regard to LSPs:

- bypass—Sessions used for bypass LSPs.
- lsp—Sessions used to set up LSPs.
- nolsp—Sessions not used to set up LSPs.

name *session-name* —(Optional) Display information about the named session.

session-type —(Optional) Display information about a particular session type:

- egress—Sessions that terminate on this switch.
- ingress—Sessions that originate from this switch.
- transit—Sessions that transit through this switch.

statistics—(Optional) Display packet statistics.

te-link *te-link* —(Optional) Display sessions with reservations on the specified traffic-engineered link name.

Required Privilege Level view

- Related Topics**
- Example: Configuring MPLS on EX Series Switches on page 11
 - Configuring MPLS on Provider Edge Switches (CLI Procedure) on page 30
 - Configuring MPLS on Provider Switches (CLI Procedure) on page 35
 - *JUNOS Software MPLS Applications Configuration Guide* at <http://www.juniper.net/techpubs/software/junos/junos95/index.html>

- List of Sample Output**
- show rsvp session on page 74
 - show rsvp session statistics on page 75
 - show rsvp session detail on page 75
 - show rsvp session extensive on page 75

- Output Fields** Table 7 on page 73 describes the output fields for the `show rsvp session` command. Output fields are listed in the approximate order in which they appear.

Table 7: show rsvp session Output Fields

Field Name	Field Description	Level of Output
Ingress RSVP	Information about ingress RSVP sessions.	detail
Ingress RSVP	Information about ingress RSVP sessions. Each session has one line of output.	All levels
Egress RSVP	Information about egress RSVP sessions.	All levels
Transit RSVP	Information about the transit RSVP sessions.	All levels
To	Destination (egress switch) of the session.	All levels
From	Source (ingress switch) of the session.	All levels
State	State of the path: Up, Down, or AdminDn. AdminDn indicates that the LSP is being taken down gracefully.	All levels
Address	Destination (egress switch) of the LSP.	detail
LSPstate	State of the LSP that is being handled by this RSVP session. It can be either Up, Dn (down), or AdminDn. AdminDn indicates that the LSP is being taken down gracefully.	brief, detail
Rt	Number of active routes (prefixes) that have been installed in the routing table. For ingress RSVP sessions, the routing table is the primary IPv4 table (inet.0). For transit and egress RSVP sessions, the routing table is the primary MPLS table (mpls.0).	brief
ActiveRoute	Number of active routes (prefixes) that have been installed in the forwarding table. For ingress RSVP sessions, the forwarding table is the primary IPv4 table (inet.0). For transit and egress RSVP sessions, the forwarding table is the primary MPLS table (mpls.0).	detail
LSPname	Name of the LSP.	brief, detail
LSPpath	Indicates whether the RSVP session is for the primary or secondary LSP path. LSPpath can be either primary or secondary and can be displayed on the ingress, egress, and transit switches. LSPpath can also indicate when a graceful LSP deletion has been triggered.	detail

Table 7: show rsvp session Output Fields (continued)

Field Name	Field Description	Level of Output
Recovery label received	(When LSP is bidirectional) Label the upstream node suggests for use in the Resv message that is sent.	detail
Recovery label sent	(When LSP is bidirectional) Label the downstream node suggests for use in its Resv messages that is returned.	detail
Suggested label received	(When LSP is bidirectional) Label the upstream node suggests for use in the Resv message that is sent.	detail
Suggested label sent	(When LSP is bidirectional) Label the downstream node suggests for use in its Resv message that is returned.	detail
Resv style or Style	RSVP reservation style. This field consists of two parts. The first is the number of active reservations. The second is the reservation style, which can be FF (fixed filter), SE (shared explicit), or WF (wildcard filter).	brief detail
Label in	Incoming label for this LSP.	brief, detail
Label out	Outgoing label for this LSP.	brief, detail
Time left	Number of seconds remaining in the lifetime of the reservation.	brief, detail
Since	Date and time when the RSVP session was initiated.	detail
Tspec	Sender's traffic specification, which describes the sender's traffic parameters.	detail
Port number	Protocol ID and sender/receiver port used in this RSVP session.	detail
Creating backup LSP, link down	A link down event occurred, and traffic is being switched over to the bypass LSP.	extensive
Deleting backup LSP, protected LSP restored	Link has come back up and the LSP has been restored. Because the backup LSP is no longer needed, it is deleted.	extensive
PATH rcvfrom	Address of the previous-hop (upstream) switch or client, interface the neighbor used to reach this switch, and number of packets received from the upstream neighbor.	detail

```

show rsvp session user@switch> show rsvp session
Ingress RSVP: 1 sessions
To          From          State  Rt Style Labelin Labelout LSPname
10.255.245.214 10.255.245.212 AdminDn 0 1 FF      -      22293 LSP Bidir
Total 1 displayed, Up 1, Down 0

Egress RSVP: 2 sessions
To          From          State  Rt Style Labelin Labelout LSPname
10.255.245.194 10.255.245.195 Up     0 1 FF   39811  - Gpro3-ba Bidir
10.255.245.194 10.255.245.195 Up     0 1 FF     3      - pro3-ba
Total 2 displayed, Up 2, Down 0

Transit RSVP: 1 sessions
To          From          State  Rt Style Labelin Labelout LSPname

```

```
10.255.245.198 10.255.245.197 Up    0 1 SE 100000      3 pro3-de
Total 1 displayed, Up 1, Down 0
```

show rsvp session statistics

```
user@switch> show rsvp session statistics
Ingress RSVP: 2 sessions
To          From          State  Packets    Bytes    LSPname
10.255.245.24 10.255.245.22 Up      0          0      pro3-bd
10.255.245.24 10.255.245.22 Up    44868    2333136  pro3-bd-2
Total 2 displayed, Up 2, Down 0
Egress RSVP: 2 sessions
To          From          State  Packets    Bytes    LSPname
10.255.245.22 10.255.245.24 Up      0          0      pro3-db
10.255.245.22 10.255.245.24 Up      0          0      pro3-db-2
Total 2 displayed, Up 2, Down 0
Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

show rsvp session detail

```
user@switch> show rsvp session detail
Ingress RSVP: 1 sessions
1.1.1.1
  From: 2.2.2.2, LSPstate: Up, ActiveRoute: 0
  LSPname: to-a, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: -, Label out: 3
  Time left: -, Since: Fri Mar 26 18:42:42 2004
  Tspec: rate 300kbps size 300kbps peak Infbps m 20 M 1500
  DiffServ info: diffServ-TE LSP, bandwidth: <ct1 300kbps>
  Port number: sender 1 receiver 15876 protocol 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  PATH sentto: 192.168.37.16 (t1-0/2/1.0) 1 pkt
```

show rsvp session extensive

```
user@switch> show rsvp session extensive
8.8.8.8
  From: 9.9.9.9, LSPstate: Up, ActiveRoute: 0
  LSPname: lsp_to_240, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 322832
  Resv style: 1 FF, Label in: -, Label out: 322832
  Time left: -, Since: Thu Feb 26 16:25:39 2009
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 2 receiver 44542 protocol 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 3.3.3.2 (xe-0/1/0.0) 238 pkts
  RESV rcvfrom: 3.3.3.2 (xe-0/1/0.0) 234 pkts
  Explct route: 3.3.3.2 4.4.4.2
```

