JUNIPER NETWORKS

Junos® OS

MPLS Fast Reroute Network Operations Guide

Published: 2011-01-19

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

Junos® OS MPLS Fast Reroute Network Operations Guide

Revision History
12 January 2007—Revision 1
July 2010—Revision 2

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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 présentés confirment leur volonté que cette convention de même que tous les documents y compris tout avis qui s'y rattaché, soient redigés en langue anglaise. (Translation: The parties confirm that this Agreement and all related documentation is and will be in the English language)).

Abbreviated Table of Contents

Table of Contents

About This Guide

This preface provides the following guidelines for using the *Junos® operating system (Junos OS) MPLS Fast Reroute Network Operations Guide*:

Objectives

This guide describes operational information helpful in monitoring and troubleshooting an MPLS network configured with fast reroute (FRR) and load balancing. This guide is not directly related to any particular release of the Junos operating system (Junos OS).

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

- *Junos OS MPLS Applications Configuration Guide*—Provides an overview of traffic engineering concepts and describes how to configure traffic engineering protocols..
- *Junos OS Feature Guide*—Provides a detailed explanation and configuration examples for several of the most complex features in the Junos OS.

For information about related tasks performed by Network Operations Center (NOC) personnel, see the following network operations guides:

- *Junos OS MPLS Fast Reroute Network Operations Guide*
- *Junos OS MPLS Log Reference Network Operations Guide*
- *Junos OS Baseline Network Operations Guide*
- *Junos OS Interfaces Network Operations Guide*

NOTE: To obtain the most current version of this manual, see the product documentation page on the Juniper Networks Web site, located at http://www.juniper.net/.

Audience

This guide is designed for Network Operations Center (NOC) personnel who monitor a Juniper Networks M Series or T Series routing platform.

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

- Border Gateway Protocol (BGP)
- Routing Information Protocol (RIP)
- Intermediate System-to-Intermediate System (IS-IS)
- Open Shortest Path First (OSPF)
- Internet Control Message Protocol (ICMP) router discovery
- Internet Group Management Protocol (IGMP)
- Distance Vector Multicast Routing Protocol (DVMRP)
- Protocol-Independent Multicast (PIM)
- Multiprotocol Label Switching (MPLS)
- Resource Reservation Protocol (RSVP)
- Simple Network Management Protocol (SNMP)

Supported Routing Platforms

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

- M Series
- T Series

Using the Index

This guide contains a complete index. For a list and description of glossary terms, see the *Junos OS Comprehensive Index and Glossary*.

Using the Examples in This Manual

If you want to use the examples in this manual, you can use the **load merge** or the **load merge relative** command. These commands cause the software to merge the incoming

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

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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

   ```
   commit {
   ```

file ex-script-snippet.xsl; }

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

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

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

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

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

Document Conventions

Table 1 on page xvi defines notice icons used in this guide.

Table 1: Notice Icons

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

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

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the **configure** command: user@host> **configure**
`Fixed-width text like this`	Represents output that appears on the terminal screen.	`user@host>` **show chassis alarms** `No alarms currently active`

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

Convention	Description	Examples
Italic text like this	• Introduces important new terms. • Identifies book names. • Identifies RFC and Internet draft titles.	• A policy *term* is a named structure that defines match conditions and actions. • *Junos OS System Basics Configuration Guide* • RFC 1997, *BGP Communities Attribute*
Italic text like this	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** *domain-name*
Text like this	Represents names of configuration statements, commands, files, and directories; IP addresses; configuration hierarchy levels; or labels on routing platform components.	• To configure a stub area, include the **stub** statement at the **[edit protocols ospf area area-id]** hierarchy level. • The console port is labeled **CONSOLE**.
< > (angle brackets)	Enclose optional keywords or variables.	**stub** <**default-metric** *metric*>**;**
| (pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	**broadcast | multicast** (***string1*** **|** ***string2*** **|** ***string3***)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	**rsvp { # Required for dynamic MPLS only**
[ ] (square brackets)	Enclose a variable for which you can substitute one or more values.	**community name members [** ***community-ids*** **]**

Indention and braces ({ })	Identify a level in the configuration hierarchy.	[edit]
		routing-options {
		  static {
		  route default {
		  nexthop *address*;
		  retain;
		  }
		  }
		}

; (semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
J-Web GUI Conventions		
Bold text like this	Represents J-Web graphical user interface (GUI) items you click or select.	• 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**.

List of Technical Publications

Table 3 on page xviii lists the software and hardware guides and release notes for Juniper Networks M-series, MX-series, and T-series routing platforms and describes the contents of each document. Table 4 on page xix lists the books included in the *Network Operations Guide* series. Table 5 on page xx lists the manuals and release notes supporting Junos OS for J-series and SRX-series platforms. All documents are available at http://www.juniper.net/techpubs/.

Table 6 on page xxi lists additional books on Juniper Networks solutions that you can order through your bookstore. A complete list of such books is available at http://www.juniper.net/books.

Table 3: Technical Documentation for Supported Routing Platforms

Book	Description
Hardware Documentation	
Hardware Guide	Describes how to install, maintain, and troubleshoot routing platforms and components. Each platform has its own hardware guide.
PIC Guide	Describes the routing platform's Physical Interface Cards (PICs). Each platform has its own PIC guide.
DPC Guide	Describes the Dense Port Concentrators (DPCs) for all MX-series routers.
Junos Scope Documentation	
Junos Scope Software User Guide	Describes the Junos Scope software graphical user interface (GUI), how to install and administer the software, and how to use the software to manage routing platform configuration files and monitor routing platform operations.
Advanced Insight Solutions (AIS) Documentation	
Advanced Insight Solutions Guide	Describes the Advanced Insight Manager (AIM) application, which provides a gateway between Junos devices and Juniper Support Systems (JSS) for case management and intelligence updates. Explains how to run AI-Scripts on Juniper Networks devices.
Release Notes	
Junos OS Release Notes	Summarize new features and known problems for a particular software release, provide corrections and updates to published Junos, Junos XML protocol, and NETCONF manuals, provide information that might have been omitted from the manuals, and describe upgrade and downgrade procedures.
Hardware Release Notes	Describe the available documentation for the routing platform and summarize known problems with the hardware and accompanying software. Each platform has its own release notes.

Table 3: Technical Documentation for Supported Routing Platforms *(continued)*

Book	Description
Junos Scope Release Notes	Contain corrections and updates to the published Junos Scope manual, provide information that might have been omitted from the manual, and describe upgrade and downgrade procedures.
AIS Release Notes	Summarize AIS new features and guidelines, identify known and resolved problems, provide information that might have been omitted from the manuals, and provide initial setup, upgrade, and downgrade procedures.
AIS AI-Scripts Release Notes	Summarize AI-Scripts new features, identify known and resolved problems, provide information that might have been omitted from the manuals, and provide instructions for automatic and manual installation, including deleting and rolling back.

Table 4: Junos OS Network Operations Guides

Book	Description
Baseline	Describes the most basic tasks for running a network using Juniper Networks products. Tasks include upgrading and reinstalling Junos OS, gathering basic system management information, verifying your network topology, and searching log messages.
Interfaces	Describes tasks for monitoring interfaces. Tasks include using loopback testing and locating alarms.
MPLS	Describes tasks for configuring, monitoring, and troubleshooting an example MPLS network. Tasks include verifying the correct configuration of the MPLS and RSVP protocols, displaying the status and statistics of MPLS running on all routing platforms in the network, and using the layered MPLS troubleshooting model to investigate problems with an MPLS network.
MPLS Log Reference	Describes MPLS status and error messages that appear in the output of the **show mpls lsp extensive** command. The guide also describes how and when to configure Constrained Shortest Path First (CSPF) and RSVP trace options, and how to examine a CSPF or RSVP failure in a sample network.
MPLS Fast Reroute	Describes operational information helpful in monitoring and troubleshooting an MPLS network configured with fast reroute (FRR) and load balancing.
Hardware	Describes tasks for monitoring M-series and T-series routing platforms.

To configure and operate a J-series Services Router or an SRX-series Services Gateway running Junos OS, you must also use the configuration statements and operational mode commands documented in Junos configuration guides and command references. To

configure and operate a WX Integrated Services Module, you must also use WX documentation.

Table 5: Junos OS for J-series Services Routers and SRX-series Services Gateways Documentation

Book	Description
J-series and SRX-series Platforms	
Junos OS Interfaces and Routing Configuration Guide	Explains how to configure SRX-series and J-series interfaces for basic IP routing with standard routing protocols, ISDN service, firewall filters (access control lists), and class-of-service (CoS) traffic classification.
Junos OS Security Configuration Guide	Explains how to configure and manage SRX-series and J-series security services such as stateful firewall policies, IPsec VPNs, firewall screens, Network Address Translation (NAT), Public Key Cryptography, chassis clusters, Application Layer Gateways (ALGs), and Intrusion Detection and Prevention (IDP).
Junos OS Administration Guide for Security Devices	Shows how to monitor SRX-series and J-series devices and routing operations, firewall and security services, system alarms and events, and network performance. This guide also shows how to administer user authentication and access, upgrade software, and diagnose common problems.
Junos OS CLI Reference	Provides the complete configuration hierarchy available on SRX-series and J-series devices. This guide also describes the configuration statements and operational mode commands unique to these devices.
Network and Security Manager: Configuring J Series Services Routers and SRX Series Services Gateways Guide	Explains how to configure, manage, and monitor J-series Services Routers and SRX-series services gateways through NSM.
Junos OS Release Notes	Summarize new features and known problems for a particular release of Junos OS, including Junos OS for J-series and SRX-series devices.The release notes also contain corrections and updates to the manuals and software upgrade and downgrade instructions for Junos OS.
J-series Only	
Junos OS Design and Implementation Guide	Provides guidelines and examples for designing and implementing IPsec VPNs, firewalls, and routing on J-series Services Routers running Junos OS.
J Series Services Routers Quick Start	Explains how to quickly set up a J-series Services Router. This document contains router declarations of conformity.
JUNOS Software with Enhanced Services J-series Services Router Hardware Guide	Provides an overview, basic instructions, and specifications for J-series Services Routers. This guide explains how to prepare a site, unpack and install the router, replace router hardware, and establish basic router connectivity. This guide contains hardware descriptions and specifications.

Table 5: Junos OS for J-series Services Routers and SRX-series Services Gateways Documentation *(continued)*

Book	Description
Junos OS Migration Guide	Provides instructions for migrating an SSG device running ScreenOS software to Junos OS or upgrading a J-series device to a later version of the Junos OS.
WXC Integrated Services Module Installation and Configuration Guide	Explains how to install and initially configure a WXC Integrated Services Module in a J-series Services Router for application acceleration.

Table 6: Additional Books Available Through http://www.juniper.net/books

Book	Description
Interdomain Multicast Routing	Provides background and in-depth analysis of multicast routing using Protocol Independent Multicast sparse mode (PIM SM) and Multicast Source Discovery Protocol (MSDP); details any-source and source-specific multicast delivery models; explores multiprotocol BGP (MBGP) and multicast IS-IS; explains Internet Gateway Management Protocol (IGMP) versions 1, 2, and 3; lists packet formats for IGMP, PIM, and MSDP; and provides a complete glossary of multicast terms.
Junos Cookbook	Provides detailed examples of common Junos OS configuration tasks, such as basic router configuration and file management, security and access control, logging, routing policy, firewalls, routing protocols, MPLS, and VPNs.
MPLS-Enabled Applications	Provides an overview of Multiprotocol Label Switching (MPLS) applications (such as Layer 3 virtual private networks [VPNs], Layer 2 VPNs, virtual private LAN service [VPLS], and pseudowires), explains how to apply MPLS, examines the scaling requirements of equipment at different points in the network, and covers the following topics: point-to-multipoint label switched paths (LSPs), DiffServ-aware traffic engineering, class of service, interdomain traffic engineering, path computation, route target filtering, multicast support for Layer 3 VPNs, and management and troubleshooting of MPLS networks.
OSPF and IS-IS: Choosing an IGP for Large-Scale Networks	Explores the full range of characteristics and capabilities for the two major link-state routing protocols: Open Shortest Path First (OSPF) and IS-IS. Explains architecture, packet types, and addressing; demonstrates how to improve scalability; shows how to design large-scale networks for maximum security and reliability; details protocol extensions for MPLS-based traffic engineering, IPv6, and multitopology routing; and covers troubleshooting for OSPF and IS-IS networks.
Routing Policy and Protocols for Multivendor IP Networks	Provides a brief history of the Internet, explains IP addressing and routing (Routing Information Protocol [RIP], OSPF, IS-IS, and Border Gateway Protocol [BGP]), explores ISP peering and routing policies, and displays configurations for both Juniper Networks and other vendors' routers.
The Complete IS-IS Protocol	Provides the insight and practical solutions necessary to understand the IS-IS protocol and how it works by using a multivendor, real-world approach.

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can send your comments to

techpubs-comments@juniper.net, or fill out the documentation feedback form at https://www.juniper.net/cgi-bin/docbugreport/. If you are using e-mail, be sure to include the following information with your comments:

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

Requesting Technical Support

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

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

Self-Help Online Tools and Resources

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

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

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

Opening a Case with JTAC

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

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

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

PART 1

Investigating Fast Reroute in an MPLS Network

CHAPTER 1

MPLS FRR Protection Introduction

MPLS FRR Protection Overview

Multiprotocol Label Switching (MPLS) fast reroute (FRR) refers to local protection methods such as one-to-one and many-to-one (facility) backup. In the general networking community, the term FRR has become a shorthand way of describing the entire spectrum of MPLS traffic protection mechanisms. This should not be confused with the Junos OS fast reroute feature. In this book, the acronym FRR is used to describe general MPLS traffic protection, while the distinct Junos OS feature is described as fast reroute.

In the Junos OS, general MPLS traffic protection for Resource Reservation Protocol (RSVP)-signaled label-switched path (LSP) failures is provided by several complementary mechanisms. These protection mechanisms include local protection (fast reroute, link protection, and node-link protection), and path protection (primary and secondary paths). Local protection in conjunction with path protection can provide minimum packet loss for an LSP, and control the way the LSP is rerouted after a failure.

Traditionally, both types of protection rely on fast detection of connectivity failure at the physical level. However, for transmission media without fast physical level detection, the Junos OS supports the configuration of bidirectional forwarding detection (BFD) and MPLS ping for fast-failure detection. It is beyond the scope of this document to cover BFD or MPLS ping. For more information on BFD and MPLS ping, see the *Junos MPLS Applications Configuration Guide*.

The terms *node* and *router* are used interchangeably throughout the topics related to this subject.

Related Documentation

- Checklist for Path Protection on page 9
- Local Protection Checklist on page 23
- Checklist for RSVP Reservation Styles on page 51
- Checklist for Load Balancing in an MPLS Network on page 63

MPLS Protection Background

During network failure, MPLS FRR protects against link or node failure in the path of an RSVP-signaled LSP with “Local Protection” on page 4 at the level of the link or node, and “Path Protection” on page 5at the level of the entire LSP. For a list of terms and acronyms, see “Terms and Acronyms” on page 6

Local Protection

Local protection includes two methods:

- One-to-one (fast reroute) backup is one dedicated detour that protects one LSP.
- Many-to-one (facility) backup is one bypass path that protects many LSPs.

In the Juniper Networks implementation, one-to-one backup corresponds to the **fast-reroute** statement, while many-to-one (facility) backup corresponds to the **link-protection** and **node-link-protection** statements. This implementation is based on RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*. Local protection is included at the MPLS and RSVP hierarchy levels, as illustrated in the sample output below. It is not recommended that you configure both types of local protection (fast reroute and facility backup) together. They are included together for illustration purposes only.

The following sample output shows the configuration of the **fast-reroute** statement:

```
[edit]
protocols {
    mpls {
        label-switched-path  lsp-path-name  {
            fast-reroute;
        }
    }
}
```

The following sample output shows the configuration of link protection (many-to-one or facility backup):

```
[edit]
protocols {
    rsvp {
        interface  type-fpc/pic/port {
            link-protection;
        }
    }
    mpls {
        label-switched-path  lsp-path-name  {
            link-protection;
        }
    }
}
```

The following sample output shows the configuration of node-link protection (many-to-one or facility backup):

```
[edit]
protocols {
    rsvp {
        interface  type-fpc/pic/port {
```

```
            link-protection;
        }
    }
    mpls {
        label-switched-path  lsp-path-name  {
            node-link-protection;
        }
    }
}
```

Local protection in the Junos OS is described as follows:

- One-to-one (fast reroute) backup—A router upstream from a failure quickly builds a detour LSP around the failure to the router downstream from the failure, providing protection against link or node failure. The upstream router then signals the outage to the ingress router, thereby maintaining connectivity before a new LSP is established. You can configure one-to-one backup by including the **fast-reroute** statement at the **[edit protocols mpls label-switched-path *path-name*]** hierarchy level.
- Link protection (many-to-one or facility backup)—Each router establishes a bypass LSP to its neighbor, avoiding the link connecting them, and ensuring traffic flow for the LSP when a link connecting two nodes fails. You can configure many-to-one backup by including the **link-protection** statement at the **[edit protocols mpls label-switched-path *path-name*]** hierarchy level.
- Node-link protection (many-to-one or facility backup)—Each router dynamically signals a bypass LSP and determines if the protected LSP needs a node bypass or a link bypass, thereby ensuring traffic flow when a node or link in the LSP fails. You can configure node-link protection by including the **node-link-protection** statement at the **[edit protocols mpls label-switched-path *path-name*]** hierarchy level. To enable node-link protection, you must also include the **link-protection statement** at the **[edit protocols rsvp interface *interface-name*]** hierarchy level.

The important difference between using the **fast-reroute** statement and either of the link-protection statements is that the **fast-reroute** statement, regardless of whether a link or node fails, always protects one LSP with one detour path. The **link-protection** and **node-link-protection** statements always protect any LSPs crossing the node with one bypass path.

There are a couple of things to consider when deciding to configure fast reroute or link protection. The first is interoperability with equipment from other vendors, for example, Cisco Systems supports FRR, but does not support one-to-one backup. The second is that protection paths consume forwarding resources. In this regard, facility backup has better scaling because the protection paths are shared.

Path Protection

Complementary to local protection methods, Junos OS supports the configuration of path protection with primary and secondary paths. By configuring path protection together with local protection, you can obtain minimum packet loss for an LSP while at the same time maintaining control over the path after the failure.

In the Junos OS, path protection is included at the MPLS hierarchy level, as illustrated in the sample output below. The sample output shows the primary, secondary, and path statements you must include to an MPLS LSP configuration.

```
[edit]
protocols {
    mpls {
        label-switched-path  lsp-path-name  {
            primary  path-name ;
            secondary  path-name  {
                standby;
            }
        }
        path  path-name  {
        }
        path  path-name  {
     }
}
```

Path protection in the Junos OS is described as follows:

- Primary paths—Dictate the physical path for the LSP and are used in normal operations. When not configured and when Constrained Shortest Path First (CSPF) is used, the label-switched router (LSR) determines the path to reach the egress router based on user constraints, such as LSP bandwidth, link color, or other constraints. You can configure primary paths by issuing the **primary *path-name*** statement at the **[edit protocols mpls label-switched-path *path-name*]** hierarchy level. For an example and more information about configuring and verifying primary paths, see "Configuring and Verifying a Primary Path" on page 11.
- Secondary paths—Become operational when the primary path fails. There are two types of secondary paths: standby and non-standby. A standby secondary path is precomputed and pre-signaled while a non-standby secondary path is precomputed but is not pre-signaled. You can configure secondary paths by issuing the **secondary *path-name*** statement at the **[edit protocols mpls label-switched-path *path-name*]** hierarchy level. To configure a standby secondary path, include the **standby** statement at the **[edit protocols mpls label-switched-path lsp-path-name secondary]** hierarchy level. For an example and more information about configuring and verifying secondary paths, see "Configuring and Verifying a Secondary Path" on page 16.

Terms and Acronyms

- *Bypass tunnel*—A label-switched path (LSP) that is used to protect multiple LSPs in many-to-one (facility) backup.
- *CSPF*—Constrained Shortest Path First. An MPLS algorithm that has been modified to take into account specific restrictions when calculating the shortest path across the network.
- *Detour LSP*—The LSP that is used to reroute traffic around a failure in one-to-one backup.
- *DMP*—Detour Merge Point. In the case of one-to-one backup, this is an LSR where multiple detours converge. Only one detour is signaled beyond that LSR.
- *Facility backup*—A local repair method in which a bypass tunnel is used to protect one or more protected LSPs that traverse the point of local repair, the resource being protected, and the merge point, in that order.
- *Local repair*–Techniques used to repair LSP tunnels quickly when a node or link along the LSP fails.

- *LSP*—An MPLS label-switched path (LSP). In this document, an LSP is always explicitly routed.
- *LSR*—Label-switching router. A router on which MPLS is enabled and that can process label-switched packets.
- *Merge point*—The LSR where one or more backup tunnels rejoin the path of the protected LSP downstream of the potential failure. The same LSR may simultaneously be a merge point and a point of local repair.
- *Next-hop bypass tunnel*—A backup tunnel that bypasses a single link for different LSPs.
- *Next-next-hop bypass tunnel*—A backup tunnel that bypasses a single node of the protected LSP.
- *One-to-one backup*—A local repair method in which a detour LSP is separately created for each protected LSP at a point of local repair.
- *Point of local repair*—The ingress (head-end) LSR of a backup tunnel or a detour LSP.
- *Protected LSP*—An LSP is protected at a given hop if it has one or multiple detours or bypass paths.

Related Documentation

For additional information about MPLS fast reroute and MPLS protection methods, see the following:

- Local Protection Overview on page 25
- Path Protection Overview on page 10
- Configuring and Verifying One-to-One Backup on page 27
- Configuring and Verifying Link Protection on page 35
- Configuring and Verifying Node-Link Protection on page 42

CHAPTER 2

Path Protection in an MPLS Network

The Junos OS implementation of Mutiprotocol Label Switching (MPLS) provides several complementary mechanisms for protecting against Resource Reservation Protocol (RSVP)-signaled LSP failures, including path protection (primary and secondary paths), and local protection (the **fast reroute** statement, link protection, and node-link protection). This chapter describes path protection supported by the Junos OS.

Checklist for Path Protection

This checklist provides the steps and commands for configuring and verifying path protection supported by the Junos OS. The checklist provides links to an overview of path protection and more detailed information about the commands used to configure and verify path protection in different scenarios.

Table 7 on page 9 provides commands for checking for path protection.

Table 7: Checklist for Path Protection

Tasks	Command or Action
"Path Protection Overview" on page 10	
"Configuring and Verifying a Primary Path" on page 11	

Table 7: Checklist for Path Protection *(continued)*

Tasks	Command or Action
1. Configure a Primary Path on page 13	**[edit]** **edit protocols mpls** **[edit protocols mpls]** **set path** *path-name address* < strict | loose > **set label-switched-path** *lsp-path-name* **to** *destination* **[edit protocols mpls label-switched-path** *lsp-path-name***]** **set primary** *primary-name* **set primary** *primary-name* **bandwidth** *bandwidth* **set primary** *primary-name* **priority** *reservation-priority setup-priority* **show** **commit**
2. Verify That the Primary Path Is Operational on page 15	**show mpls lsp extensive ingress** **show rsvp interface**
"Configuring and Verifying a Secondary Path" on page 16	
1. Configure a Standby Secondary Path on page 17	**[edit]** **edit protocols mpls** **[edit protocols mpls]** **set path** *path-name destination* **loose** **set label-switched-path** *lsp-path-name* **secondary** *secondary-name***standby** **show** **commit**
2. Verify That the Secondary Path Is Established on page 18	Deactivate a link or node critical to the primary path. **show mpls lsp extensive**
"Ensuring That Secondary Paths Establish When Resources Are Diminished" on page 20	Configure different bandwidth values for the primary and secondary paths. For example: **[edit protocols mpls]** **edit label-switched-path** *lsp-path-name* **set primary** *primary-name* **bandwidth** *bandwidth* **show** **commit** In this example, no bandwidth is configured for the secondary path.
"Preventing Use of a Path That Previously Failed" on page 21	Configure only multiple secondary paths.

Path Protection Overview

The main advantages of path protection are control over where the traffic goes after a failure and minimum packet loss when combined with fast reroute (one-to-one backup

or link protection). Path protection is the configuration, within a label-switched path (LSP), of two types of paths: a primary path, used in normal operations, and a secondary path used when the primary fails, as shown in Figure 1 on page 11.

In Figure 1 on page 11, an MPLS network consisting of eight routers has a primary path between **R1** and **R5** which is protected by the secondary path between **R1** and **R5**. When a failure is detected, such as an interface down event, an Resource Reservation Protocol (RSVP) error message is sent to the ingress router which switches traffic to the secondary path, maintaining traffic flow.

Figure 1: Path Protection

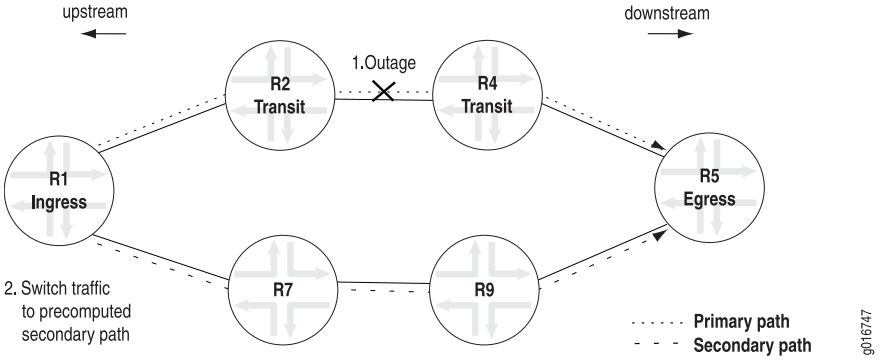

If the secondary path is pre-signaled or on standby, recovery time from a failure is faster than if the secondary path is not pre-signaled. When the secondary path is not pre-signaled a call-setup delay occurs during which the new physical path for the LSP is established, extending the recovery time. If the failure in the primary path is corrected, and after a few minutes of hold time, the ingress router switches traffic back from the secondary path to the primary path.

Because path protection is provided by the ingress router for the entire path, there can be some disadvantages, for example, double-booking of resources and unnecessary protection of links. By protecting a single resource at a time, local protection can remedy these disadvantages.

Related Documentation

-

Configuring and Verifying a Primary Path

Purpose Primary paths are optional and when configured, limit the RSVP calculation of the complete path to the routers specified in the primary Explicit Route Object (ERO) list, which determines the physical path for the LSP. When primary paths are not configured, the ingress router determines the path to the egress router. Only one primary path is permitted per LSP, as shown in Figure 2 on page 12.

Figure 2: Primary Path

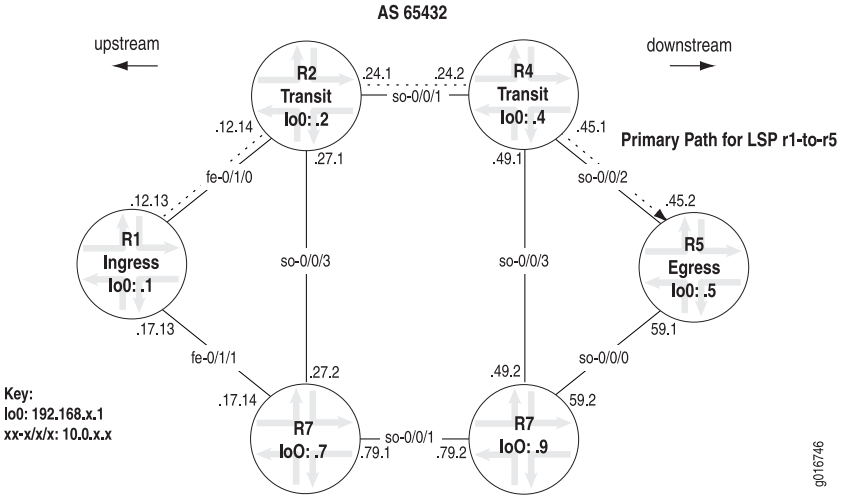

Within the configuration of the primary physical path, you can specify strict or loose ERO values and parameters that affect only the primary physical path, such as bandwidth or priority. The ERO list for the primary path includes an address for each transit router. Specifying the ingress and/or egress routers is optional. For each router address, you can specify the type, which can be one of the following:

- Strict—The route taken from the previous router to this router is a direct path and cannot include any other routers. This is the default. If the address is an interface address, this router also ensures that the incoming interface is the one specified. Specifying the incoming interface is important when there are parallel links between the previous router and this router, and because it ensures that routing can be enforced on a per-link basis.

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

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

If you are listing more than one address, specify the addresses in order, starting with the ingress router (optional) or the first transit router, and continuing sequentially along the path up to the egress router (optional) or the router immediately before the egress router. You need to specify only one address per router hop. If you specify more than one address for the same router, only the first address is used; the additional addresses are ignored and truncated.

When configuring a primary path, you can specify the bandwidth and priority values associated with that primary path.

The bandwidth value is included in the sender's Tspec field in RSVP path setup messages. You specify the bandwidth value in bits per second, with a higher value implying a greater user traffic volume. The default bandwidth is 0 bits per second. A nonzero bandwidth

requires transit routers to reserve capacity along the outbound links for the path. The RSVP reservation scheme is used to reserve this capacity. Any failure in bandwidth reservation (such as failures at RSVP policy control or admission control) might cause the LSP setup to fail.

The priority value is composed of two distinct values: a setup and a hold priority. The setup priority value is used to determine if there is enough bandwidth available at that priority level to establish the primary path. The priority level is between 0 (best) and 7 (worst).The hold priority value is used by an established primary path to retain its bandwidth reservations in the network. If insufficient link bandwidth is available during session establishment, the setup priority is compared to the hold priorities of other established sessions to determine whether some of them should be preempted to accommodate the new session. Sessions with worse hold priorities are preempted.

To configure and verify a primary path, follow these steps:

1. Configure a Primary Path on page 13
2. Verify That the Primary Path Is Operational on page 15

Configure a Primary Path

Action To configure a primary path with an ERO list, bandwidth, and priority, follow these steps:

1. In configuration mode, go to the following hierarchy level:

   ```
   [edit]
   user@host# edit protocols mpls
   ```

2. Configure the primary ERO list:

   ```
   [edit protocols mpls]
   user@host# set path path-name address strict
   ```

   For example:

   ```
   [edit protocols mpls]
   user@R1# set path via-r2 10.0.12.14 strict
   user@R1# set path via-r2 10.0.24.2 strict
   ```

3. Configure the LSP:

   ```
   [edit protocols mpls]
   user@host# set label-switched-path lsp-path-name to destination;
   ```

   For example:

   ```
   [edit protocols mpls]
   user@R1# set label-switched-path r1-to-r5 to 192.168.5.1;
   ```

4. Configure the primary path:

   ```
   [edit protocols mpls label-switched-path lsp-path-name]
   user@host# set primary primary-name
   ```

   For example:

   ```
   [edit protocols mpls label-switched-path r1-to-r5]
   user@R1# set primary via-r2
   ```


5. Configure the bandwidth:

   [edit protocols mpls label-switched-path *lsp-path-name*]
   user@host# set primary *primary-name* bandwidth *bandwidth*

   For example:

   [edit protocols mpls label-switched-path r1-to-r5]
   user@R1# set primary via-r2 bandwidth 35m

6. Configure the priority value:

   [edit protocols mpls label-switched-path *lsp-path-name*]
   user@host# set primary *primary-name* priority *reservation-priority setup-priority*

   For example:

   [edit protocols mpls label-switched-path r1-to-r5]
   user@R1# set primary via-r2 priority 6 6

7. Verify and commit the configuration:

   [edit protocols mpls label-switched-path *lsp-path-name*]
   user@host# show
   user@host# commit

Sample Output The sample output below illustrates the configuration of the primary path on ingress router **R1** in the network shown in Figure 2 on page 12.

```
[edit protocols mpls]
user@R1# show
label-switched-path r1-to-r5 {
    to 192.168.5.1;
    primary via-r2 {   #  Bandwidth and priority configured at the primary path

        bandwidth 35m; #  level of the hierarchy
        priority 6 6; #  Priority setup and hold values
    }
}
path via-r2 { #  Primary ERO list
    10.0.12.14 strict;
    10.0.24.2 strict;
[...Output truncated...]

[edit protocols mpls]
user@R1# commit
commit complete
```

Meaning The sample output shows a label-switched path (LSP) with bandwidth and priority applied to only one primary path. The same parameters specified one level up in the hierarchy, at the **[edit protocols mpls label-switched-path *lsp-path-name*]** hierarchy level, affect all paths.

The path, **via-r2**, specifies the complete strict path from the ingress to the egress routers through **10.0.12.14**, **10.0.24.2**, in that order. There cannot be any intermediate routers except the ones specified. However, there can be intermediate routers between **10.0.24.2** and the egress router because the egress router is not specifically listed in the path statement. To prevent intermediate routers before egress, configure the egress router as the last router, with a strict type.

For more information on configuring a primary path, see the *Junos MPLS Applications Configuration Guide*.

Verify That the Primary Path Is Operational

Purpose Primary paths must always be used in the network if they are available, therefore an LSP always moves back to the primary path after a failure, unless the configuration is adjusted. For more information on adjusting the configuration to prevent a failed primary path from reestablishing, see "Preventing Use of a Path That Previously Failed" on page 21.

Action To verify that the primary path is operational, enter the following Junos OS command-line interface (CLI) operational mode commands:

```
user@host> show mpls lsp extensive ingress
user@host> show rsvp interface
```

Sample Output 1

```
user@R1> show mpls lsp extensive ingress
Ingress LSP: 1 sessions

192.168.5.1
  From: 192.168.1.1, State: Up, ActiveRoute: 0, LSPname: r1-to-r5
   ActivePath: via-r2 (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary  via-r2              State: Up
     Priorities: 6 6
     Bandwidth: 35Mbps
    SmartOptimizeTimer: 180
     Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 11)
 10.0.12.14 S 10.0.24.2 S
     Received RRO  (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node
10=SoftPreempt):
           10.0.12.14 10.0.24.2
    5 Apr 29 14:40:43 Selected as active path
    4 Apr 29 14:40:43 Record Route:  10.0.12.14 10.0.24.2
    3 Apr 29 14:40:43 Up
    2 Apr 29 14:40:43 Originate Call
    1 Apr 29 14:40:43 CSPF: computation result accepted
  Standby   via-r7            State: Dn
    SmartOptimizeTimer: 180
        No computed ERO.
  Created: Sat Apr 29 14:40:43 2006
Total 1 displayed, Up 1, Down 0
```

Sample Output 2

```
user@R1> show rsvp interface
RSVP interface: 3 active
                    Active Subscr- Static      Available   Reserved    Highwater
Interface   State resv   iption  BW          BW          BW          mark
fe-0/1/0.0  Up         2   100%   100Mbps      100Mbps     0bps        0bps

fe-0/1/1.0  Up         1   100%   100Mbps      100Mbps     0bps        0bps

so-0/0/3.0  Up         1   100%   155.52Mbps   155.52Mbps  0bps        0bps
```

Meaning Sample output 1 shows that the LSP is operational and is using the primary path (**via-r2**) with **R2** (**10.0.12.14**) and **R4** (**10.0.24.2**) as transit routers. The priority values are the same for setup and hold, **6 6**. Priority 0 is the highest (best) priority and 7 is the lowest (worst)

priority. The Junos OS default for setup and hold priority is 7:0. Unless some LSPs are more important than others, preserving the default is a good practice. Configuring a setup priority that is better than the hold priority is not allowed, resulting in a failed commit in order to avoid preemption loops.

Configuring and Verifying a Secondary Path

Secondary paths (also known as secondary LSPs) are optional and protect against link and transit node failures. If the primary path can no longer reach the egress router, the alternative, secondary path is used, as shown in Figure 3 on page 16.

Figure 3: Standby Secondary Paths

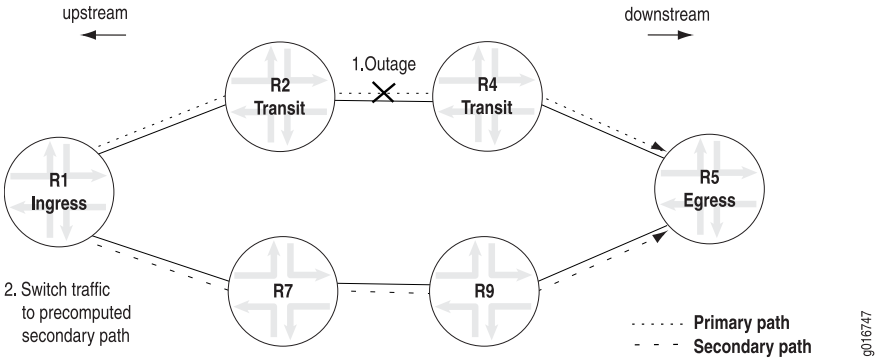

In Figure 3 on page 16, a secondary path **R1-R7-R9-R5** is activated when the primary path **R1-R2-R4-R5** fails. **R2** notifies **R1** of the outage and **R1** switches traffic to the precomputed secondary path.

Two types of secondary paths, standby and non-standby, can become active when a primary path fails, depending on which is configured. A standby secondary path, configured with the **standby** statement, is precomputed and pre-signaled. A non-standby secondary path, configured without the **standby** statement, is precomputed but is not pre-signaled.

Secondary paths configured with the **standby** statement consume more resources because the router must maintain state when the secondary path is not active. However, standby secondary paths do reduce recovery time by eliminating the call-setup delay that is required to establish a new physical path for the LSP.

If the problem with the primary path is corrected, after a few minutes of hold-down to ensure that the primary path remains stable, the ingress router switches traffic from the secondary path back to the primary path. It may not be always prudent for the router to switch back to the primary path. For information on how to keep the router from switching back to the primary path, see “Preventing Use of a Path That Previously Failed” on page 21.

To configure and verify a secondary path, follow these steps:

1. Configure a Standby Secondary Path on page 17
2. Verify That the Secondary Path Is Established on page 18

Configure a Standby Secondary Path

Configuring a standby secondary path is a two-part process. In the first part, you define the path, and in the second part, you specify a secondary path for the LSP that refers to the defined path.

NOTE: To configure a non-standby secondary path, simply omit the standby statement from the secondary path configuration.

To configure a standby secondary path, follow these steps:

Action

1. In configuration mode, go to the following hierarchy level:

   ```
   [edit]
   user@host# edit protocols mpls
   ```

2. Configure the secondary ERO list:

   ```
   [edit protocols mpls]
   user@host# set path path-name destination loose
   ```

   For example:

   ```
   [edit protocols mpls]
   user@R1# set path via-r7 10.0.17.14 loose
   ```

3. Configure the LSP and the secondary path:

   ```
   [edit protocols mpls]
   user@host# set label-switched-path lsp-path-name secondary secondary-name
     standby
   ```

   For example:

   ```
   [edit protocols mpls]
   user@R1# set label-switched-path r1-to-r4 secondary via-r7 standby
   ```

4. Verify and commit the configuration:

   ```
   [edit protocols mpls]
   user@host# show
   user@host# commit
   ```

The sample output below illustrates the configuration of the standby secondary path on ingress router **R1** in the network shown in Figure 2 on page 12.

Sample Output

```
[edit protocols mpls]
user@R1# show
label-switched-path r1-to-r4 {
    to 192.168.4.1;
    ldp-tunneling;
    fast-reroute;
    primary via-r2;
     secondary via-r7 {
        standby; #  Omit the standby statement to configure a non-standby secondary
 path
    }
}
```

```
path via-r2 {
    10.0.12.14 loose;
}
path via-r7 {
    10.0.17.14 loose;
}
[...Output truncated...]
```

Meaning The sample output shows one standby secondary path **via-r7**, which includes the **standby** statement at the **[edit protocols mpls label-switched-path *lsp-path-name* secondary *secondary-name*]** hierarchy level. The standby secondary path is defined in the **path** statement **path via-r7** and specifies a loose hop, indicating that the route taken from the previous router to this router need not be a direct path, can include other routers, and can be received on any interface.

If you have many secondary paths configured for an LSP, and you want them all to be standby, include the **standby** statement one level up in the hierarchy, at the **[edit protocols mpls label-switched-path *lsp-path-name*]** hierarchy level, as shown in the sample output below.

```
[edit protocols mpls]
user@R1# show
label-switched-path r1-to-r4 {
    to 192.168.4.1;
     standby; #  Standby configured at the label-switched-path level of the
hierarchy
    primary via-r2;
    }
    secondary via-r7;
    }
}
[...Output truncated...]
```

For more information on configuring a secondary path, see the *Junos MPLS Applications Configuration Guide*.

Verify That the Secondary Path Is Established

Purpose When the secondary path is configured with the **standby** statement, the secondary path should be *up* but *not active*; it will become active if the primary path fails. A secondary path configured without the **standby** statement will not come up unless the primary path fails. To test that the secondary path is correctly configured and would come up if the primary path were to fail, you must deactivate a link or node critical to the primary path, then issue the **show mpls lsp *lsp-path-name* extensive** command.

Action To verify that the secondary path is established, enter the following Junos OS CLI operational mode command:

```
user@R1> show mpls lsp extensive
```

Sample Output The following sample output shows a correctly configured secondary path before and after it comes up. In the example, interface **fe-0/1/0** on **R2** is deactivated, which brings down the primary path **via-r2**. The ingress router **R1** switches traffic to the secondary path **via-r7**.

```
user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

192.168.5.1
  From: 192.168.1.1, State: Up, ActiveRoute: 0, LSPname: r1-to-r5
  ActivePath: via-r2 (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary  via-r2              State: Up
    Priorities: 6 6
    Bandwidth: 35Mbps
    SmartOptimizeTimer: 180
     Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
 10.0.12.14 S 10.0.24.2 S 10.0.45.2 S
   Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

         10.0.12.14 10.0.24.2 10.0.45.2
    5 Apr 29 14:40:43 Selected as active path
    4 Apr 29 14:40:43 Record Route:  10.0.12.14 10.0.24.2
    3 Apr 29 14:40:43 Up
    2 Apr 29 14:40:43 Originate Call
    1 Apr 29 14:40:43 CSPF: computation result accepted
   Secondary via-r7              State: Dn
    SmartOptimizeTimer: 180
        No computed ERO.
  Created: Sat Apr 29 14:40:43 2006
Total 1 displayed, Up 1, Down 0

[edit interfaces]
user@R2# deactivate fe-0/1/0

[edit interfaces]
user@R2# show
inactive: fe-0/1/0 {
    unit 0 {
        family inet {
            address 10.0.12.14/30;
        }
        family iso;
        family mpls;
    }
}

user@R1> show mpls lsp name r1-to-r4 extensive
Ingress LSP: 1 sessions

192.168.4.1
  From: 192.168.1.1, State: Up, ActiveRoute: 0, LSPname: r1-to-r4
   ActivePath: via-r7 (secondary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
   Primary  via-r2              State: Dn
    Priorities: 6 6
    Bandwidth: 35Mbps
    SmartOptimizeTimer: 180
    Will be enqueued for recomputation in 14 second(s).
   10 Apr 29 14:52:33 CSPF failed: no route toward 10.0.12.1 4[21 times]
    9 Apr 29 14:42:48 Clear Call
    8 Apr 29 14:42:48 Deselected as active
    7 Apr 29 14:42:48 Session preempted
    6 Apr 29 14:42:48 Down
```

```
    5 Apr 29 14:40:43 Selected as active path
    4 Apr 29 14:40:43 Record Route:  10.0.12.14 10.0.24.2
    3 Apr 29 14:40:43 Up
    2 Apr 29 14:40:43 Originate Call
    1 Apr 29 14:40:43 CSPF: computation result accepted
  *Standby  via-r7            State: Up
    SmartOptimizeTimer: 180
     Computed ERO   (S [L] denotes strict [loose] hops): (CSPF metric: 11)
 10.0.17.14 S 10.0.47.1 S
   Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

         10.0.17.14 10.0.47.1
    5 Apr 29 14:42:48 Selected as active path
    4 Apr 29 14:41:12 Record Route:  10.0.17.14 10.0.47.1
    3 Apr 29 14:41:12 Up
    2 Apr 29 14:41:12 Originate Call
    1 Apr 29 14:41:12 CSPF: computation result accepted
  Created: Sat Apr 29 14:40:43 2006
Total 1 displayed, Up 1, Down 0
```

Meaning The sample output from egress router **R1** shows a correctly configured standby secondary path in a down state because the primary path is still up. Upon deactivation of an interface (**interface fe-0/1/0** on **R2**) critical to the primary path, the primary path **via-r2** goes down and the standby secondary path **via-r7** comes up, allowing **R1** to switch traffic to the standby secondary path.

Ensuring That Secondary Paths Establish When Resources Are Diminished

The Junos OS does not require that a primary and secondary path share the same parameters. You may decide to configure your primary paths with strict resource requirements, and configure your secondary paths with less strict requirements, allowing your secondary paths to establish more readily during periods of diminished resources.

Action To ensure that secondary paths establish when resources are diminished, follow these steps:

1. In configuration mode, go to the following hierarchy level:

   ```
   [edit protocols mpls]
   user@host# edit label-switched-path lsp-path-name
   ```

   For example:

   ```
   [edit protocols mpls]
   user@R1# edit label-switched-path r1-to-r4
   ```

2. Configure the bandwidth for the primary path, and do not configure any bandwidth for the secondary path:

   ```
   [edit protocols mpls label-switched-path lsp-path-name]
   user@host# set primary primary-name bandwidth bandwidth
   ```

   For example:

   ```
   [edit protocols mpls label-switched-path r1-to-r4]
   user@R1# set primary via-r2 bandwidth 35m
   ```

3. Verify and commit the configuration:

[edit protocols mpls label-switched-path *lsp-path-name*]
user@host# **show**
user@host# **commit**

Sample Output The sample output below illustrates a bandwidth configuration on ingress router R1 in the network shown in Figure 2 on page 12.

```
[edit protocols mpls]
user@R1# show
label-switched-path r1-to-r4 {
    to 192.168.4.1;
     primary via-r2 {
         bandwidth 35m;
    }
    secondary via-r7 { #  In this example, bandwidth is not configured for the
secondary path.
        standby;       #  However you could configure a bandwidth value different
 from
    }                   #  that on the primary path.
}
[...Output truncated...]
```

Meaning The sample output shows the primary path **via-r2** requires 35 Mbps of bandwidth, while secondary path **via-r7** has no constraints. The primary path is configured with strict resource requirements, while the secondary path is configured with no bandwidth requirements, allowing the secondary path to establish more readily during periods of diminished resources. One thing to keep in mind when configuring a secondary path without bandwidth requirements is that it can be subject to traffic loss due to congestion.

Preventing Use of a Path That Previously Failed

If you configure an alternate path through the network in case the active path fails, you may not want traffic to revert back to the failed path, even if it is no longer failing. When you configure a primary path, the traffic switches over to the secondary path during a failure, and reverts back to the primary path when it returns.

At times, switching traffic back to a primary path that has previously failed may not be a particularly sound idea. In this case, only configure secondary paths, resulting in the next configured secondary path establishing when the first secondary path fails. Later, if the first secondary path becomes operational, the Junos OS will not revert to it, but will continue using the second secondary path.

CHAPTER 3

Local Protection in an MPLS Network

The Junos OS implementation of Multiprotocol Label Switching (MPLS) provides several complementary mechanisms for protecting against Resource Reservation Protocol (RSVP)-signaled LSP failures, including path protection (primary and secondary paths), and local protection (fast reroute, link protection, and node-link protection). This chapter describes local protection supported by the Junos OS.

The terms *node* and *router* are used interchangeably throughout this book.

Local Protection Checklist

This checklist provides the steps and commands for configuring and verifying local protection supported by the Junos OS. Also, the checklist provides links to an overview of local protection and more details about the commands used to configure and verify one-to-one backup, many-to-one (facility) protection, and node-link protection. See Table 8 on page 23

Table 8:

Tasks	Command or Action
"Local Protection Overview" on page 25	
"One-to-One Backup Overview" on page 26	
"Configuring and Verifying One-to-One Backup" on page 27	
1. Configure One-to-One Backup on page 27	**[edit]** **edit protocols mpls** **[edit protocols mpls]** **set label-switched-path *lsp-path-name* to *address*** **set label-switched-path *lsp-path-name* fast-reroute** **(optional) set label-switched-path *lsp-path-name* primary *primary-name*** **(optional) set path *path-name address* loose** **show** **commit**
2. Verify One-to-One Backup on page 28	**show mpls lsp ingress extensive** **show rsvp session**

Table 8: *(continued)*

Tasks	Command or Action
"Many-to-One Link Protection (Facility Backup) Overview" on page 34	
"Configuring and Verifying Link Protection" on page 35	
1. Configure Link Protection on page 35	**[edit]** **edit protocols rsvp interface** *type-fpc/pic/port* **[edit protocols rsvp interface** *type-fpc/pic/port***]** **set link-protection** **show** **top** **edit protocols mpls label-switched-path** *lsp-path-name* **[edit protocols mpls label-switched-path** *lsp-path-name***]** **set link-protection** **show** **commit**
2. Verify That Link Protection Is Up on page 36	**show mpls lsp extensive** **show rsvp session detail** **show rsvp interface**
"Configuring and Verifying Node-Link Protection" on page 42	
1. Configure Node-Link Protection on page 42	**[edit]** **edit protocols mpls label-switched-path** *lsp-path-name* **[edit protocols mpls label-switched-path** *lsp-path-name***]** **set node-link-protection** **show** **edit protocols rsvp interface** *type-fpc/pic/port* **[edit protocols rsvp interface** *type-fpc/pic/port***]** **set link-protection** **show** **commit** Include the **node-link-protection** statement on any other ingress routers that have LSPs requiring use of the bypass path. Include the **link-protection** statement on routers with outgoing interfaces in the LSP.
2. Verify That Node-Link Protection Is Up on page 43	**show mpls lsp** **show mpls lsp extensive** **show rsvp interface** **show rsvp interface extensive** **show rsvp session detail**

Local Protection Overview

Local protection attempts to address the disadvantages of path protection by focusing on a single resource at a time (link or node), in contrast to path protection which attempts to provide protection for the entire path from the ingress router to the egress router. Double-booking of resources, unnecessary protection and nondeterministic switchover times are the main disadvantages of path protection, arising from protection at the ingress router for the entire path. By providing focused protection from the ingress of a single resource at a time, local protection addresses the disadvantages of path protection, minimizing the amount of time during which traffic is lost, while utilizing resources efficiently.

In Figure 4 on page 25, if the LSP from **R1** to **R5** fails on the link between **R2** and **R4**, a detour or bypass path is pre-established quickly, and traffic is redirected around the failure, until the ingress router moves the LSP to a new path that does not use the failed link.

Figure 4: Local Protection

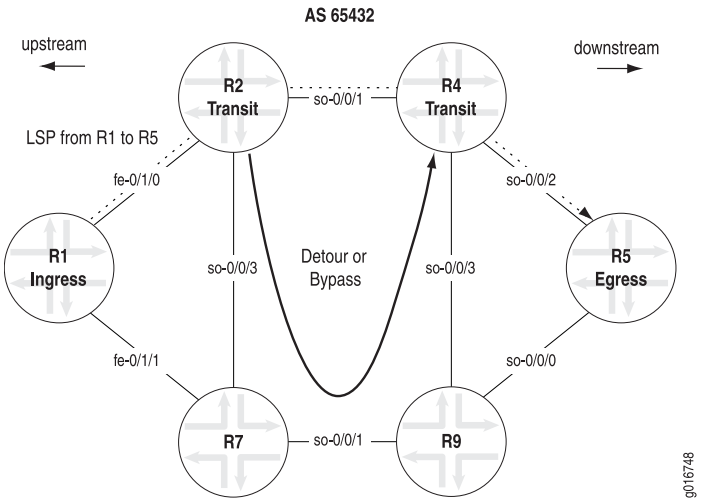

In the Juniper Networks implementation, local protection methods are defined by the number of LSPs protected by the backup path. When one LSP is protected by one backup path, the backup path is referred to as a detour and the protection method is called fast reroute (one-to-one backup). When many LSPs are protected by one backup path, the backup path is referred to as a bypass and the protection method is called facility backup. The purpose of facility backup is to protect a link or node (facility). Facility backup can be used for protecting either a link or a node (and its associated links), also referred to as node-link protection.

The following local protection methods are discussed in this section:

- One-to-One Backup Overview on page 26
- Many-to-One Link Protection (Facility Backup) Overview on page 34
- Node-Link Protection Overview on page 40

One-to-One Backup Overview

Fast reroute or one-to-one backup is a short-term solution to reduce packet loss associated with a particular LSP. One-to-one backup is appropriate under the following circumstances:

- Protection of a small number of LSPs relative to the total number of LSPs.
- Path selection criteria, such as bandwidth, priority, and link coloring for detour paths is critical.
- Control of individual LSPs is important.

In one-to-one backup, the ingress router adds the fast reroute object to the RSVP Path message requesting that downsteam routers establish detours. Downstream routers generate Path messages and establish detours to avoid the downstream link or node. Detours are always calculated to avoid the immediate downstream link and node, providing against both link and node failure, as shown in Figure 5 on page 26.

Figure 5: One-to-One Backup Detours

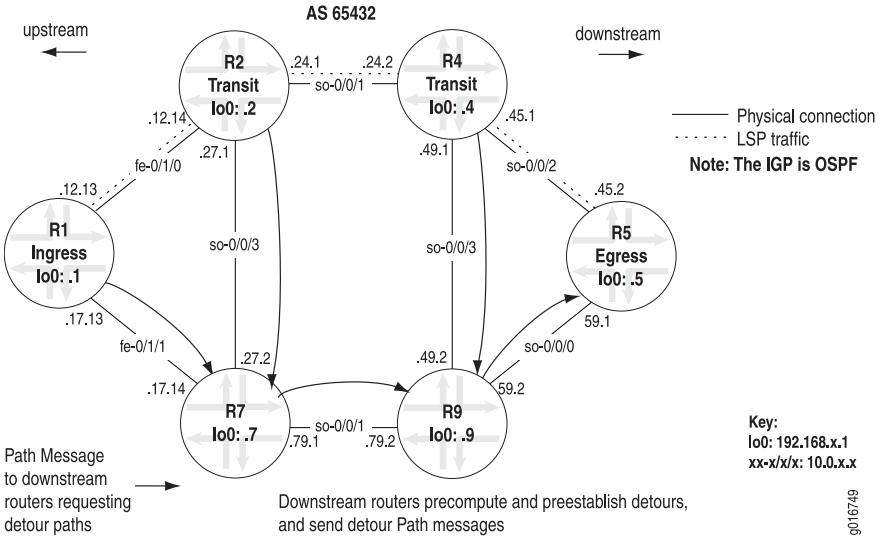

Figure 5 on page 26 shows a network with one LSP configured from the ingress router **R1** to the egress router **R5**, transiting **R2** and **R4**. The following detours are established:

- **R1** creates a detour to **R5** via **R7** and **R9**
- **R2** creates a detour to **R5** via **R7** and **R9**
- **R4** creates a detour to **R5** via **R9**

Each detour is dedicated to a particular LSP traversing the router (one detour to one LSP). If the network topology has insufficient links and nodes, it may be impossible to establish a detour. Also, detour paths are not meant for long-term use because they may provide inadequate bandwidth and can result in congestion on the links. As soon as the

ingress router calculates a new path avoiding the failure, traffic is redirected along the new path, detours are torn down, and new detours established.

Configuring and Verifying One-to-One Backup

The following sections describe the steps you must take to configure and verify one-to-one backup.

Configure One-to-One Backup

The following steps show the commands you must issue to configure a LSP with fast reroute and a primary path. The **show** command output includes bandwidth and hop limit for your information only. Bandwidth and hop limit are not configured on **R1**. You can configure bandwidth and hop limit using the **bandwidth** and **hop-limit** statements at the **[edit protocols mpls lsp *lsp-path-name*]** hierarchy level.

NOTE: It is not necessary to issue the fast-reroute statement on the transit or egress routers.

Action To configure one-to-one backup on the ingress router, follow these steps:

1. In configuration mode, go to the following hierarchy level:

   ```
   [edit]
   user@R1# edit protocols mpls
   ```

2. Configure the LSP:

   ```
   [edit protocols mpls]
   user@R1# set label-switched-path lsp-path-name to address
   ```

   For example:

   ```
   [edit protocols mpls]
   user@R1# set label-switched-path r1-to-r5 to 192.168.5.1
   ```

3. Configure one-to-one backup (fast reroute):

   ```
   [edit protocols mpls]
   user@R1# set label-switched-path lsp-path-name fast-reroute
   ```

   For example:

   ```
   [edit protocols mpls]
   user@R1# set label-switched-path r1-to-r5 fast-reroute
   ```

4. (Optional) Configure a primary path:

   ```
   [edit protocols mpls]
   user@R1# set label-switched-path lsp-path-name  primary primary-name
   ```

   For example:

   ```
   [edit protocols mpls]
   ```

```
user@R1# set label-switched-path r1-to-r5 primary via-r2
```

5. (Optional) Configure the primary ERO list:

```
[edit protocols mpls]
user@R1# set path path-name address loose
```

For example:

```
[edit protocols mpls]
user@R1# set path via-r2 10.0.12.14 loose
```

6. Verify and commit the configuration:

```
[edit protocols mpls]
user@R1# show
user@R1# commit
```

Sample Output

```
[edit protocols mpls]
user@R1# show
label-switched-path r1-to-r5 {
    to 192.168.5.1;
    fast-reroute;
    primary via-r2;
    bandwidth  bps; # Bandwidth for the LSP
    hop-limit  number; # Maximum number of routers the LSP can traverse
}
path via-r2 {
    10.0.12.14 loose;
}
[...Output truncated...]

[edit protocols mpls]
user@R1# commit
commit complete
```

Meaning When the **fast-reroute** statement is configured, the ingress router signals all downstream routers to compute and preestablish a detour path for the LSP, using the Constrained Shortest Path First (CSPF) algorithm on the information in the local router's traffic engineering database (TED). By default, when the detour path is calculated by CSPF, the detour path inherits the same administrative group constraints (link coloring or resource classes) as the main LSP.

Verify One-to-One Backup

Purpose You can verify that one-to-one backup is established by examining the ingress router and the other routers in the network.

Action To verify one-to-one backup, enter the following Junos OS CLI operational mode commands:

```
user@host> show mpls lsp ingress extensive
user@host> show rsvp session
```

Sample Output The following sample output is from the ingress router **R1** in the network shown in Figure 5 on page 26:

```
user@R1> show mpls lsp ingress extensive
Ingress LSP: 1 sessions

192.168.5.1
  From: 192.168.1.1, State: Up, ActiveRoute: 0, LSPname: r1-to-r5
  ActivePath: via-r2 (primary)
   FastReroute desired
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   via-r2           State: Up
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
 10.0.12.14 S 10.0.24.2 S 10.0.45.2 S
     Received RRO  (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node
10=SoftPreempt):
           10.0.12.14(flag=9) 10.0.24.2(flag=1) 10.0.45.2
     8 May 11 14:51:46 Fast-reroute Detour Up
     7 May 11 14:50:55 Record Route: 10.0.12.14(flag=9) 10.0.24.2(flag=1) 10.0.45.2
    6 May 11 14:50:55 Record Route:  10.0.12.14(flag=9) 10.0.24.2 10.0.45.2
    5 May 11 14:50:52 Selected as active path
    4 May 11 14:50:52 Record Route:  10.0.12.14 10.0.24.2 10.0.45.2
    3 May 11 14:50:52 Up
    2 May 11 14:50:52 Originate Call
    1 May 11 14:50:52 CSPF: computation result accepted
  Created: Thu May 11 14:50:52 2006
Total 1 displayed, Up 1, Down 0
```

Meaning The sample output from **R1** shows that the **FastReroute desired** object was included in the Path messages for the LSP, allowing **R1** to select the active path for the LSP and establish a detour path to avoid **R2**.

In line 8, **Fast-reroute Detour Up** shows that the detour is operational. Lines 6 and 7 indicate that transit routers **R2** and **R4** have established their detour paths.

R2, **10.0.12.14**, includes **(flag=9)**, indicating that node protection is available for the downstream node and link. **R4**, **10.0.24.2**, includes **(flag=1)**, indicating that link protection is available for the next downstream link. In this case, **R4** can protect only the downstream link because the node is the egress router **R5**, which cannot be protected. For more information about flags, see the *Junos Feature Guide*.

The output for the **show mpls lsp extensive** command does not show the actual path of the detour. To see the actual links used by the detour paths, you must use the **show rsvp session ingress detail** command.

Sample Output The following sample output is from the ingress router **R1** in the network shown in Figure 5 on page 26.

```
user@R1> show rsvp session ingress detail
Ingress RSVP: 1 sessions

192.168.5.1
  From: 192.168.1.1,  LSPstate: Up,  ActiveRoute: 0
  LSPname: r1-to-r5, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 100848
  Resv style: 1 FF, Label in: -, Label out: 100848
  Time left:    -, Since: Thu May 11 14:17:15 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
```

```
  Port number: sender 1 receiver 9228 protocol 0
   FastReroute desired
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.12.14 (fe-0/1/0.0) 35 pkts
  RESV rcvfrom: 10.0.12.14 (fe-0/1/0.0) 25 pkts
  Explct route: 10.0.12.14 10.0.24.2 10.0.45.2
  Record route: <self> 10.0.12.14 10.0.24.2 10.0.45.2
     Detour is Up
    Detour Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Detour adspec: sent MTU 1500
    Path MTU: received 1500
    Detour PATH sentto: 10.0.17.14 (fe-0/1/1.0) 23 pkts
    Detour RESV rcvfrom: 10.0.17.14 (fe-0/1/1.0) 20 pkts
     Detour Explct route: 10.0.17.14 10.0.79.2 10.0.59.1
    Detour Record route: <self> 10.0.17.14 10.0.79.2 10.0.59.1
    Detour Label out: 100848
Total 1 displayed, Up 1, Down 0
```

Meaning The sample output from **R1** shows the RSVP session of the main LSP. The detour path is established, **Detour is Up**. The physical path of the detour is displayed in **Detour Explct route**. The detour path uses **R7** and **R9** as transit routers to reach **R5**, the egress router.

Sample Output The following sample output is from the first transit router R2 in the network shown in Figure 5 on page 26:

```
user@R2> show rsvp session transit detail
Transit RSVP: 1 sessions

192.168.5.1
  From: 192.168.1.1, LSPstate: Up, ActiveRoute: 1
  LSPname: r1-to-r5, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 100448
  Resv style: 1 FF, Label in: 100720, Label out: 100448
  Time left:  126, Since: Wed May 10 16:12:21 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 5 receiver 9216 protocol 0
   FastReroute desired
  PATH rcvfrom: 10.0.12.13 (fe-0/1/0.0) 173 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 10.0.24.2 (so-0/0/1.0) 171 pkts
  RESV rcvfrom: 10.0.24.2 (so-0/0/1.0) 169 pkts
  Explct route: 10.0.24.2 10.0.45.2
  Record route: 10.0.12.13 <self> 10.0.24.2 10.0.45.2
     Detour is Up
    Detour Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Detour adspec: received MTU 1500 sent MTU 1500
    Path MTU: received 1500
    Detour PATH sentto: 10.0.27.2 (so-0/0/3.0) 169 pkts
    Detour RESV rcvfrom: 10.0.27.2 (so-0/0/3.0) 167 pkts
     Detour Explct route: 10.0.27.2 10.0.79.2 10.0.59.1
    Detour Record route: 10.0.12.13 <self> 10.0.27.2 10.0.79.2 10.0.59.1
    Detour Label out: 100736
Total 1 displayed, Up 1, Down 0
```


Meaning The sample output from **R2** shows the detour is established (**Detour is Up**) and avoids **R4**, and the link connecting **R4** and **R5** (**10.0.45.2**). The detour path is through **R7** (**10.0.27.2**) and **R9** (**10.0.79.2**) to **R5** (**10.0.59.1**), which is different from the explicit route for the detour from **R1**. **R1** has the detour passing through the **10.0.17.14** link on **R7**, while **R1** is using the **10.0.27.2** link. Both detours merge at **R9** through the **10.0.79.2** link to **R5** (**10.0.59.1**).

Sample Output The following sample output is from the second transit router R4 in the network shown in Figure 5 on page 26:

```
user@R4> show rsvp session transit detail
Transit RSVP: 1 sessions

192.168.5.1
  From: 192.168.1.1,  LSPstate: Up, ActiveRoute: 1
   LSPname: r1-to-r5, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: 100448, Label out: 3
  Time left:  155, Since: Wed May 10 16:15:38 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 5 receiver 9216 protocol 0
   FastReroute desired
  PATH rcvfrom: 10.0.24.1 (so-0/0/1.0) 178 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 10.0.45.2 (so-0/0/2.0) 178 pkts
  RESV rcvfrom: 10.0.45.2 (so-0/0/2.0) 175 pkts
  Explct route: 10.0.45.2
  Record route: 10.0.12.13 10.0.24.1 <self> 10.0.45.2
      Detour is Up
     Detour Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
     Detour adspec: received MTU 1500 sent MTU 1500
     Path MTU: received 1500
     Detour PATH sentto: 10.0.49.2 (so-0/0/3.0) 176 pkts
     Detour RESV rcvfrom: 10.0.49.2 (so-0/0/3.0) 175 pkts
      Detour Explct route: 10.0.49.2 10.0.59.1
      Detour Record route: 10.0.12.13 10.0.24.1 <self> 10.0.49.2 10.0.59.1
     Detour Label out: 100352
Total 1 displayed, Up 1, Down 0
```

Meaning The sample output from **R4** shows the detour is established (**Detour is Up**) and avoids the link connecting **R4** and **R5** (**10.0.45.2**). The detour path is through R9 (**10.0.49.2**) to **R5** (**10.0.59.1**). Some of the information is similar to that found in the output for **R1** and **R2**. However, the explicit route for the detour is different, going through the link connecting **R4** and **R9** (**so-0/0/3** or **10.0.49.2.**

Sample Output The following sample output is from **R7**, which is used in the detour path in the network shown in Figure 5 on page 26:

```
user@R7> show rsvp session transit detail
Transit RSVP: 1 sessions, 1 detours

192.168.5.1
  From: 192.168.1.1, LSPstate: Up, ActiveRoute: 1
   LSPname: r1-to-r5, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
```

```
  Recovery label received: -, Recovery label sent: 100368
  Resv style: 1 FF, Label in: 100736, Label out: 100368
  Time left:  135, Since: Wed May 10 16:14:42 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 5 receiver 9216 protocol 0
   Detour branch from 10.0.27.1, to skip 192.168.4.1, Up
    Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Adspec: received MTU 1500
    Path MTU: received 0
    PATH rcvfrom: 10.0.27.1 (so-0/0/3.0) 179 pkts
    Adspec: received MTU 1500 sent MTU 1500
    PATH sentto: 10.0.79.2 (so-0/0/1.0) 177 pkts
    RESV rcvfrom: 10.0.79.2 (so-0/0/1.0) 179 pkts
    Explct route: 10.0.79.2 10.0.59.1
     Record route: 10.0.12.13 10.0.27.1 <self> 10.0.79.2 10.0.59.1
     Label in: 100736, Label out: 100368
   Detour branch from 10.0.17.13, to skip 192.168.2.1, Up
    Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Adspec: received MTU 1500
    Path MTU: received 0
    PATH rcvfrom: 10.0.17.13 (fe-0/1/1.0) 179 pkts
    Adspec: received MTU 1500
    PATH sentto: 10.0.79.2 (so-0/0/1.0) 0 pkts
    RESV rcvfrom: 10.0.79.2 (so-0/0/1.0) 0 pkts
    Explct route: 10.0.79.2 10.0.59.1
     Record route: 10.0.17.13 <self> 10.0.79.2 10.0.59.1
     Label in: 100752, Label out: 100368
Total 1 displayed, Up 1, Down 0
```

Meaning The sample output from **R7** shows the same information as for a regular transit router used in the primary path of the LSP: the ingress address (**192.168.1.1**), the egress address (**192.168.5.1**), and the name of the LSP (**r1-to-r5**). Two detour paths are displayed; the first to avoid **R4** (**192.168.4.1**) and the second to avoid **R2** (**192.168.2.1**). Because **R7** is used as a transit router by **R2** and **R4**, **R7** can merge the detour paths together as indicated by the identical **Label out** value (**100368**) for both detour paths. Whether **R7** receives traffic from **R4** with a label value of **100736** or from **R2** with a label value of **100752**, **R7** forwards the packet to **R5** with a label value of **100368**.

Sample Output The following sample output is from R9, which is a router used in the detour path in the network shown in Figure 5 on page 26:

```
user@R9> show rsvp session transit detail
Transit RSVP: 1 sessions, 1 detours

192.168.5.1
  From: 192.168.1.1, LSPstate: Up, ActiveRoute: 1
  LSPname: r1-to-r5, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: 100352, Label out: 3
  Time left:  141, Since: Wed May 10 16:16:40 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 5 receiver 9216 protocol 0
   Detour branch from 10.0.49.1, to skip 192.168.5.1, Up
    Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Adspec: received MTU 1500
    Path MTU: received 0
    PATH rcvfrom: 10.0.49.1 (so-0/0/3.0) 183 pkts
```



```
    Adspec: received MTU 1500 sent MTU 1500
    PATH sentto: 10.0.59.1 (so-0/0/0.0) 182 pkts
    RESV rcvfrom: 10.0.59.1 (so-0/0/0.0) 183 pkts
    Explct route: 10.0.59.1
     Record route: 10.0.12.13 10.0.24.1 10.0.49.1 <self> 10.0.59.1
    Label in: 100352, Label out: 3
   Detour branch from 10.0.27.1, to skip 192.168.4.1, Up
    Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Adspec: received MTU 1500
    Path MTU: received 0
   Detour branch from 10.0.17.13, to skip 192.168.2.1, Up
    Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Adspec: received MTU 1500
    Path MTU: received 0
    PATH rcvfrom: 10.0.79.1 (so-0/0/1.0) 181 pkts
    Adspec: received MTU 1500
    PATH sentto: 10.0.59.1 (so-0/0/0.0) 0 pkts
    RESV rcvfrom: 10.0.59.1 (so-0/0/0.0) 0 pkts
     Explct route: 10.0.59.1
    Record route: 10.0.12.13 10.0.27.1 10.0.79.1 <self> 10.0.59.1
     Label in: 100368, Label out: 3
Total 1 displayed, Up 1, Down 0
```

Meaning The sample output from **R9** shows that **R9** is the penultimate router for the detour path, the explicit route includes only the egress link address (**10.0.59.1**), and the **Label out** value (**3**) indicates that **R9** has performed penultimate-hop label popping. Also, the detour branch from **10.0.27.1** does not include path information because **R7** has merged the detour paths from **R2** and **R4**. Notice that the **Label out** value in the detour branch from **10.0.17.13** is **100368**, the same value as the **Label out** value on **R7**.

Sample Output The following sample output is from the egress router R5 in the network shown in Figure 5 on page 26:

```
user@R5> show rsvp session egress detail
Egress RSVP: 1 sessions, 1 detours

192.168.5.1
  From: 192.168.1.1, LSPstate: Up, ActiveRoute: 0
  LSPname: r1-to-r5, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left:  119, Since: Thu May 11 14:44:31 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 9230 protocol 0
   FastReroute desired
  PATH rcvfrom: 10.0.45.1 (so-0/0/2.0) 258 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
   Record route: 10.0.12.13 10.0.24.1 10.0.45.1 <self>
   Detour branch from 10.0.49.1, to skip 192.168.5.1, Up
    Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Adspec: received MTU 1500
    Path MTU: received 0
   Detour branch from 10.0.27.1, to skip 192.168.4.1, Up
    Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Adspec: received MTU 1500
```

```
    Path MTU: received 0
   Detour branch from 10.0.17.13, to skip 192.168.2.1, Up
    Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Adspec: received MTU 1500
    Path MTU: received 0
    PATH rcvfrom: 10.0.59.2 (so-0/0/0.0) 254 pkts
    Adspec: received MTU 1500
    PATH sentto: localclient
    RESV rcvfrom: localclient
    Record route: 10.0.12.13 10.0.24.1 10.0.49.1 10.0.59.2 <self>
     Label in: 3, Label out: -
Total 1 displayed, Up 1, Down 0
```

Meaning The sample output from **R5** shows the main LSP in the **Record route** field and the detours through the network.

Many-to-One Link Protection (Facility Backup) Overview

Many-to-one (facility backup) is based on interface rather than on LSP. While fast reroute protects interfaces or nodes along the entire path of a LSP, many-to-one protection can be applied on interfaces as needed, as shown in Figure 6 on page 34. In Figure 6 on page 34, a bypass path is set up around the link to be protected (**10.0.12.14**) using an alternate interface to forward traffic. The bypass path is shared by all protected LSPs traversing the failed link (many LSPs protected by one bypass path).

Figure 6: Many-to-One or Link Protection

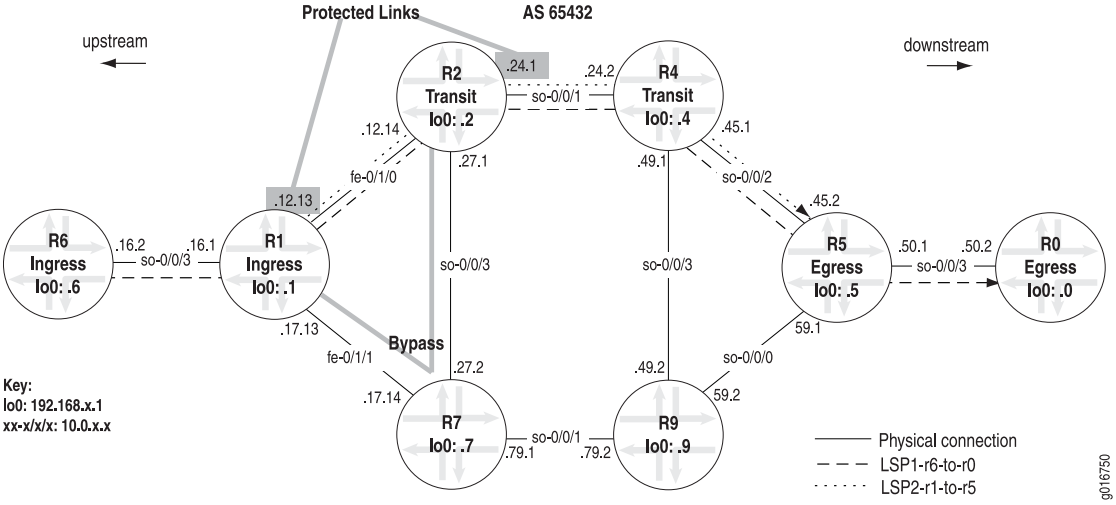

In Figure 6 on page 34, two LSPs (**lsp1-r6-to-r0** and **lsp2-r1-to-r5**) are protected by one preestablished bypass path from **R1** to **R2** through **R7**. Both LSPs have strict paths configured that go through interface **fe-0/1/0**. On **R1**, the interface **10.0.12.13** has link protection configured that protects the next hop **10.0.12.14**.

Link protection (many-to-one or facility backup) allows a router immediately upstream from a link failure to use an alternate interface to forward traffic to its downstream neighbor. This is accomplished by preestablishing a bypass path that is shared by all protected LSPs traversing the failed link. A single bypass path can safeguard a set of

protected LSPs. When an outage occurs, the router immediately upstream from the link outage switches protected traffic to the bypass link, then signals the link failure to the ingress router.

Like fast reroute, link protection provides local repair and restores connectivity faster than the ingress router switching traffic to a standby secondary path. However, unlike fast reroute, link protection does not provide protection against the failure of the downstream neighbor.

Link protection is appropriate in the following situations:

- The number of LSPs to be protected is large.
- Satisfying path selection criteria (priority, bandwidth, and link coloring) for bypass paths is less critical.
- Control at the granularity of individual LSPs is not required.

Configuring and Verifying Link Protection

The following sections describe the steps you must take to configure and verify link protection (many-to-one backup):

1. Configure Link Protection on page 35
2. Verify That Link Protection Is Up on page 36

Configure Link Protection

Purpose Configuring link protection is a two-part process. The first part involves configuring link protection on the RSVP interface, and the second part sets link protection for any LSPs traversing the protected link that require use of the bypass path.

Action To configure link protection, follow these steps:

1. In configuration mode, go to the following hierarchy level:

   ```
   [edit]
   user@R1# edit protocols rsvp interface type-fpc/pic/port
   ```

   For example:

   ```
   [edit]
   user@R1# edit protocols rsvp interface fe-0/1/0
   ```

2. Configure link protection for the interface:

   ```
   [edit protocols rsvp interface type-fpc/pic/port]
   user@R1# set link-protection
   ```

3. Verify the link protection configuration for the interface:

   ```
   [edit protocols rsvp interface type-fpc/pic/port]
   user@R1# show
   ```

4. Configure link protection for LSPs requiring use of the bypass path:

   ```
   [edit protocols rsvp interface fe-0/1/0.0]
   ```

```
user@R1# top
[edit]
user@R1# edit protocols mpls label-switched-path lsp-path-name
```

For example:

```
[edit]
user@R1# edit protocols mpls label-switched-path lsp2-r1-to-r5
```

5. Configure link protection for the LSP:

```
[edit protocols mpls label-switched-path lsp-path-name]
user@R1# set link-protection
```

6. Verify and the link protection configuration for the LSP:

```
[edit protocols mpls label-switched-path lsp-path-name]
user@R1# show
user@R1# commit
```

Sample Output The following sample output illustrates the configuration of the link protection on ingress router R1 in the network shown in Figure 6 on page 34:

```
[edit protocols rsvp]
user@R1# show
interface fe-0/1/0.0 {
     link-protection; #Protection for the RSVP interface
}

[edit protocols mpls label-switched-path lsp2-r1-to-r5]
user@R1# up

[edit protocols mpls]
user@R1# show
label-switched-path lsp2-r1-to-r5  { #Path level of the hierarchy
    #to 192.168.5.1;
    link-protection;

[edit protocols mpls]
user@R1# commit
commit complete
```

Meaning The sample output shows link protection for a specific interface. After link protection is configured, a bypass path is signaled to avoid that link in case of a failure. Having a bypass path available does not in itself provide protection for LSPs that traverse the protected link. You must configure link protection on the ingress router for each LSP that will benefit from the bypass path.

Verify That Link Protection Is Up

Purpose When you verify link protection, you must check that the bypass LSP is up. You can also check the number of LSPs protected by the bypass. In the network shown in Figure 6 on page 34, a bypass path should be up to protect the link between **R1** and **R2**, or next-hop **10.0.12.14**, and the two LSPs traversing the link, **lsp2-r1-to-r5** and **lsp1-r6-to-r0**.

Action To verify link protection (many-to-one backup), enter the following Junos OS CLI operational mode commands on the ingress router:

```
user@host> show mpls lsp extensive
user@host> show rsvp session detail
user@host> show rsvp interface
```

Sample Output

```
user@R1> show mpls lsp extensive | no-more
Ingress LSP: 1 sessions

192.168.5.1
  From: 192.168.1.1, State: Up, ActiveRoute: 0, LSPname: lsp2-r1-to-r5
  ActivePath:  via-r2 (primary)
 Link protection desired
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   via-r2           State: Up
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
 10.0.12.14 S 10.0.24.2 S 10.0.45.2 S
     Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):
          10.0.12.14(Label=101264) 10.0.24.2(Label=100736) 10.0.45.2(Label=3)
    6 Jun 16 14:06:33  Link-protection Up
    5 Jun 16 14:05:39 Selected as active path
    4 Jun 16 14:05:39 Record Route:  10.0.12.14(Label=101264)
10.0.24.2(Label=100736) 10.0.45.2(Label=3)
    3 Jun 16 14:05:39 Up
    2 Jun 16 14:05:39 Originate Call
    1 Jun 16 14:05:39 CSPF: computation result accepted
  Created: Fri Jun 16 14:05:38 2006
Total 1 displayed, Up 1, Down 0

[...Output truncated...]

Transit LSP: 2 sessions

192.168.0.1
  From: 192.168.6.1, LSPstate: Up, ActiveRoute: 0
   LSPname: lsp1-r6-to-r0, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 101296
  Resv style: 1 SE, Label in: 100192, Label out: 101296
  Time left:  116, Since: Mon Jun 19 10:26:32 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 58739 protocol 0
  Link protection desired
  Type: Link protected LSP, using Bypass->10.0.12.14
       1 Jun 19 10:26:32 Link protection up, using Bypass->10.0.12.14
  PATH rcvfrom: 10.0.16.2 (so-0/0/3.0) 579 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 10.0.12.14 (fe-0/1/0.0) 474 pkts
  RESV rcvfrom: 10.0.12.14 (fe-0/1/0.0) 501 pkts
  Explct route: 10.0.12.14 10.0.24.2 10.0.45.2 10.0.50.2
  Record route: 10.0.16.2 <self> 10.0.12.14 10.0.24.2 10.0.45.2 10.0.50.2
[...Output truncated...]
```

Meaning The sample output from ingress router **R1** shows that **lsp2-r1-to-r5** and **lsp1-r6-to-r0** have link protection up, and both LSPs are using the bypass path, **10.0.12.14**. However, the **show mpls lsp** command does not list the bypass path. For information about the bypass path, use the **show rsvp session** command.

Sample Output

```
user@R1> show rsvp session detail
Ingress RSVP: 2 sessions
```

```
192.168.2.1
  From: 192.168.1.1, LSPstate: Up, ActiveRoute: 0
   LSPname: Bypass->10.0.12.14
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 101456
  Resv style: 1 SE, Label in: -, Label out: 101456
  Time left:    -, Since: Fri May 26 18:38:09 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 18709 protocol 0
   Type: Bypass LSP
     Number of data route tunnel through: 2
    Number of RSVP session tunnel through: 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.17.14 (fe-0/1/1.0) 51939 pkts
  RESV rcvfrom: 10.0.17.14 (fe-0/1/1.0) 55095 pkts
   Explct route: 10.0.17.14 10.0.27.1
  Record route: <self> 10.0.17.14 10.0.27.1

192.168.5.1
  From: 192.168.1.1, LSPstate: Up, ActiveRoute: 0
  LSPname: lsp2-r1-to-r5, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 101264
  Resv style: 1 SE, Label in: -, Label out: 101264
  Time left:    -, Since: Fri Jun 16 14:05:39 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 18724 protocol 0
  Link protection desired
  Type: Link protected LSP
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.12.14 (fe-0/1/0.0) 8477 pkts
  RESV rcvfrom: 10.0.12.14 (fe-0/1/0.0) 8992 pkts
  Explct route: 10.0.12.14 10.0.24.2 10.0.45.2
  Record route: <self> 10.0.12.14 10.0.24.2 10.0.45.2
Total 2 displayed, Up 2, Down 0

Egress RSVP: 1 sessions

192.168.1.1
  From: 192.168.5.1, LSPstate: Up, ActiveRoute: 0
  LSPname: r5-to-r1, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left:  159, Since: Mon May 22 22:08:16 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 64449 protocol 0
  PATH rcvfrom: 10.0.17.14 (fe-0/1/1.0) 63145 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.0.59.1 10.0.79.2 10.0.17.14 <self>
Total 1 displayed, Up 1, Down 0
Transit RSVP: 2 sessions
192.168.0.1
  From: 192.168.6.1, LSPstate: Up, ActiveRoute: 0
  LSPname: lsp1-r6-to-r0, LSPpath: Primary
```



```
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 101296
  Resv style: 1 SE, Label in: 100192, Label out: 101296
  Time left:  129, Since: Mon Jun 19 10:26:32 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 58739 protocol 0
  Link protection desired
  Type: Link protected LSP
  PATH rcvfrom: 10.0.16.2 (so-0/0/3.0) 3128 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 10.0.12.14 (fe-0/1/0.0) 2533 pkts
  RESV rcvfrom: 10.0.12.14 (fe-0/1/0.0) 2685 pkts
  Explct route: 10.0.12.14 10.0.24.2 10.0.45.2 10.0.50.2
  Record route: 10.0.16.2 <self> 10.0.12.14 10.0.24.2 10.0.45.2 10.0.50.2

192.168.6.1
  From: 192.168.0.1, LSPstate: Up, ActiveRoute: 0
  LSPname: r0-to-r6, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: 100128, Label out: 3
  Time left:  143, Since: Thu May 25 12:30:26 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 4111 protocol 0
  PATH rcvfrom: 10.0.17.14 (fe-0/1/1.0) 57716 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 10.0.16.2 (so-0/0/3.0) 54524 pkts
  RESV rcvfrom: 10.0.16.2 (so-0/0/3.0) 50534 pkts
  Explct route: 10.0.16.2
  Record route: 10.0.50.2 10.0.59.1 10.0.79.2 10.0.17.14 <self> 10.0.16.2
Total 2 displayed, Up 2, Down 0
```

Meaning The sample output from ingress router **R1** shows the ingress, egress, and transit LSPs for **R1**. Some information is similar to that found in the **show mpls lsp** command. However, because link protection is an RSVP feature, information about bypass paths is provided. The bypass path appears as a separate RSVP ingress session for the protected interface, as indicated by the **Type** field.

The bypass path name is automatically generated. By default, the name appears as **Bypass >** ***interface-address***, where the interface address is the next downstream router's interface (**10.0.12.14**). The explicit route **10.0.17.14 10.0.27.1** for the session shows **R7** as the transit node and **R2** as the egress node.

Within the ingress RSVP section of the output, the LSP originating at **R1** (**lsp2-r1-to-r5)** is shown requesting link protection. Since a bypass path is in place to protect the downstream link, **lsp2-r1-to-r5** is associated with the bypass, as indicated by the **Link protected LSP** field.

The egress section of the output shows the return LSP **r5-to-r1**, which is not protected.

The transit section of the output shows link protection requested by **lsp1-r6-to-r0.** Since a bypass path is in place to protect the downstream link, **lsp1-r6-to-r0** is associated with the bypass, as indicated by the **Link protected LSP** field. Also included in the transit section of the output is the return LSP **r0-to-r6**, which is not protected.

Sample Output

```
user@R1> show rsvp interface
RSVP interface: 4 active
                    Active Subscr- Static      Available   Reserved    Highwater
Interface   State resv   iption  BW          BW          BW          mark
fe-0/1/0.0  Up           2    100%  100Mbps     100Mbps     0bps        35Mbps
fe-0/1/1.0  Up           1    100%  100Mbps     100Mbps     0bps        0bps
fe-0/1/2.0  Up           0    100%  100Mbps     100Mbps     0bps        0bps
so-0/0/3.0  Up           1    100%  155.52Mbps  155.52Mbps  0bps        0bps
```

Meaning The sample output from ingress router **R1** shows the number of LSPs going through the interfaces configured on **R1**. The **Active resv** field shows the number of LSPs for each interface. For example, interface **fe-0/1/0.0** between **R1** and **R2** has two active reservations, **lsp1-r6-to-r0** and **lsp2-r1-to-r5**; interface **fe-0/1/1.0** between **R1** and **R7** has one, the bypass (**10.0.12.14**); interface **fe-0/1/2.0** between **R6** and **R1** has no LSP reservations; and interface **so-0/0/3.0** between **R6** and **R1** has one LSP reservation, **lsp1-r6-to-r0**.

Node-Link Protection Overview

Node-link protection (many-to-one or facility backup) extends the capabilities of link protection and provides slightly different protection from fast reroute. While link protection is useful for selecting an alternate path to the same router when a specific link fails, and fast reroute protects interfaces or nodes along the entire path of an LSP, node-link protection establishes a bypass path that avoids a particular node in the LSP path.

When you enable node-link protection for an LSP, you must also enable link protection on all RSVP interfaces in the path. Once enabled, the following types of bypass paths are established:

- Next-hop bypass LSP—Provides an alternate route for an LSP to reach a neighboring router. This type of bypass path is established when you enable either node-link protection or link protection.
- Next-next-hop bypass LSP—Provides an alternate route for an LSP through a neighboring router en route to the destination router. This type of bypass path is established exclusively when node-link protection is configured.

Figure 7 on page 41 illustrates the example MPLS network topology used in this topic. The example network uses OSPF as the interior gateway protocol (IGP) and a policy to create traffic.

Figure 7: Node-Link Protection

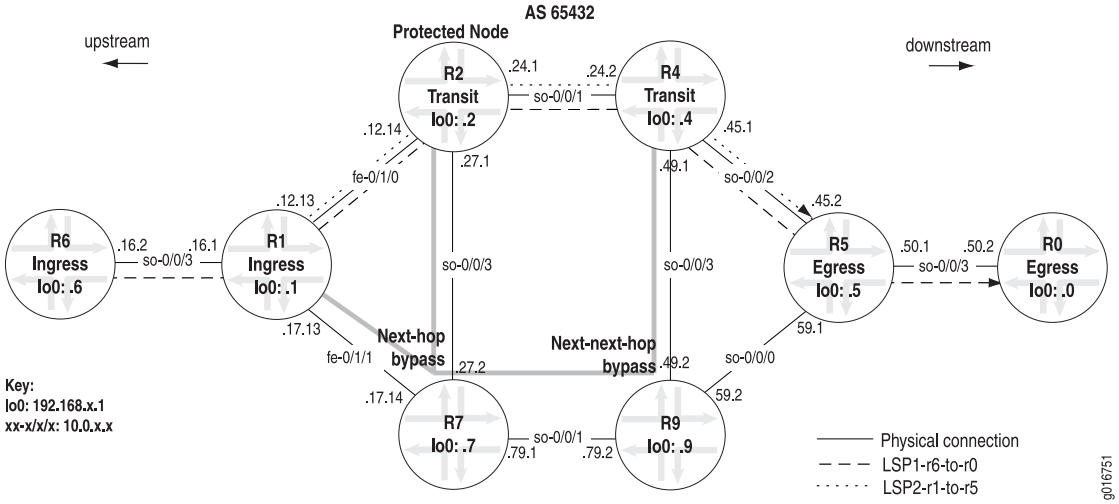

The MPLS network in Figure 7 on page 41 illustrates a router-only network that consists of unidirectional LSPs between **R1** and **R5**, (**lsp2-r1-to-r5**) and between **R6** and **R0** (**lsp1-r6-to-r0**). Both LSPs have strict paths configured that go through interface **fe-0/1/0**.

In the network shown in Figure 7 on page 41, both types of bypass paths are preestablished around the protected node (**R2).** A next-hop bypass path avoids interface **fe-0/1/0** by going through **R7**, and a next-next-hop bypass path avoids **R2** altogether by going through **R7** and **R9** to **R4**. Both bypass paths are shared by all protected LSPs traversing the failed link or node (many LSPs protected by one bypass path).

Node-link protection (many-to-one or facility backup) allows a router immediately upstream from a node failure to use an alternate node to forward traffic to its downstream neighbor. This is accomplished by preestablishing a bypass path that is shared by all protected LSPs traversing the failed link.

When an outage occurs, the router immediately upstream from the outage switches protected traffic to the bypass node, and then signals the failure to the ingress router. Like fast reroute, node-link protection provides local repair, restoring connectivity faster than the ingress router can establish a standby secondary path or signal a new primary LSP.

Node-link protection is appropriate in the following situations:

- Protection of the downstream link and node is required.
- The number of LSPs to be protected is large.
- Satisfying path selection criteria (priority, bandwidth, and link coloring) for bypass paths is less critical.
- Control at the granularity of individual LSPs is not required.

Configuring and Verifying Node-Link Protection

The following section describes the steps you must take to configure and verify many-to-one backup.

1. Configure Node-Link Protection on page 42
2. Verify That Node-Link Protection Is Up on page 43

Configure Node-Link Protection

Configuring node-link protection is a two-part process. The first part involves configuring node-link protection for any LSPs traversing the protected node that require use of the bypass path, and the second part sets link protection on the outgoing RSVP interface on routers in the LSP.

Action To configure node-link protection, follow these steps:

1. In configuration mode, go to the following hierarchy level:

   ```
   [edit]
   user@R1# edit protocols mpls label-switched-path lsp-path-name
   ```

   For example:

   ```
   [edit]
   user@R1# edit protocols mpls label-switched-path lsp2-r1-to-r5
   ```

2. Configure node-link protection for the LSP:

   ```
   [edit protocols mpls label-switched-path lsp-path-name]
   user@R1# set node-link-protection
   ```

3. Verify the node-link protection configuration for the LSP:

   ```
   [edit protocols mpls label-switched-path lsp-path-name]
   user@R1# show
   ```

4. Configure link protection for the interface:

   ```
   [edit protocols]
   user@R1# edit protocols rsvp interface interface-name
   ```

   For example:

   ```
   [edit protocols]
   user@R1# edit protocols rsvp interface fe-0/1/0
   ```

5. Configure link protection:

   ```
   [edit protocols rsvp interface interface-name]
   user@R1# set link-protection
   ```

6. Verify the link protection configuration for the interface, and commit both configurations:

   ```
   [edit protocols rsvp interface interface-name]
   user@R1# show
   user@R1# commit
   ```

7. Repeat Step 1 through Step 3 on any other ingress routers that have LSPs requiring use of the bypass path.
8. Repeat Step 4 and Step 5 on routers with outgoing interfaces in the LSP.

Sample Output The following sample output shows the configuration of node-link protection on ingress router R1 in the network shown in Figure 6 on page 34:

```
[edit protocols mpls label-switched-path lsp2-r1-to-r5]
user@R1# up

[edit protocols mpls]
user@R1# show
label-switched-path lsp2-r1-to-r5  { #Label-switched-path level of the hierarchy

     to 192.168.5.1;
     node-link-protection; #LSP node-link protection

[edit protocols rsvp]
user@R1# show
interface fe-0/1/0.0 {
     link-protection; #Link protection for the RSVP interface
}

[edit protocols rsvp]
user@R1# commit
commit complete
```

Meaning The sample output shows the configuration of node-link protection for an LSP. After node-link protection is configured, bypass paths are signaled to avoid the protected link or node in case of failure. Having bypass paths available does not in itself provide protection for LSPs that traverse the protected node. You must include the **node-link-protection** statement on the ingress router for each LSP that will benefit from the bypass path.

Verify That Node-Link Protection Is Up

Purpose After you configure node-link protection, you must check that bypass paths are up. You can also check the number of LSPs protected by the bypass paths. In the network shown in Figure 7 on page 41, two bypass paths should be up: one next-hop bypass path protecting the link between **R1** and **R2** (or next-hop **10.0.12.14**), and a next-next-hop bypass path avoiding **R2**.

Action To verify node-link protection (many-to-one backup), enter the following Junos OS CLI operational mode commands on the ingress router. You can also issue the commands on transit routers and other routers used in the bypass path for slightly different information.

show mpls lsp (See **Sample Output** on page 44)
show mpls lsp extensive (See Sample Output on page 44)
show rsvp interface (See Sample Output on page 45)
show rsvp interface extensive (See Sample Output on page 46)
show rsvp session detail (See Sample Output on page 46)

Sample Output

```
user@R1> show mpls lsp
Ingress LSP: 1 sessions
To              From            State Rt ActivePath       P     LSPname
192.168.5.1     192.168.1.1     Up     0 via-r2           *      lsp2-r1-to-r5
Total 1 displayed,  Up 1 , Down 0

Egress LSP: 1 sessions
To              From            State   Rt Style Labelin Labelout LSPname
192.168.1.1     192.168.5.1     Up       0  1 FF       3        - r5-to-r1
Total 1 displayed,  Up 1 , Down 0

Transit LSP: 2 sessions
To              From            State   Rt Style Labelin Labelout LSPname
192.168.0.1     192.168.6.1      Up        0  1 FF 100464   101952  lsp1-r6-to-r0
192.168.6.1     192.168.0.1      Up        0  1 FF  100448        3  r0-to-t6
Total 2 displayed, Up 2, Down 0
```

Meaning Sample output from **R1** for the **show mpls lsp** command shows a brief description of the state of configured and active LSPs for which **R1** is the ingress, transit, and egress router. All LSPs are up. **R1** is the ingress router for **lsp2-r1-to-r5**, and the egress router for return LSP **r5-to-r1**. Two LSPs transit **R1**, **lsp1-r6-to-r0** and the return LSP **r0-to-t6**. For more detailed information about the LSP, include the **extensive** option when you issue the **show mpls lsp** command.

Sample Output

```
user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

192.168.5.1
  From: 192.168.1.1,  State: Up , ActiveRoute: 0,  LSPname: lsp2-r1-to-r5
  ActivePath: via-r2 (primary)
   Node/Link protection desired
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   via-r2           State: Up
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
 10.0.12.14 S 10.0.24.2 S 10.0.45.2 S
   Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

         10.0.12.14(Label=101872) 10.0.24.2(Label=101360) 10.0.45.2(Label=3)
    11 Jul 11 14:30:58 Link-protection Up
   10 Jul 11 14:28:28 Selected as active path
    [...Output truncated...]
  Created: Tue Jul 11 14:22:58 2006
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

192.168.1.1
  From: 192.168.5.1,  LSPstate: Up, ActiveRoute: 0
   LSPname: r5-to-r1, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left:  146, Since: Tue Jul 11 14:28:36 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 29228 protocol 0
  PATH rcvfrom: 10.0.12.14 (fe-0/1/0.0) 362 pkts
  Adspec: received MTU 1500
```

```
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.0.45.2 10.0.24.2 10.0.12.14 <self>
Total 1 displayed, Up 1, Down 0

Transit LSP: 2 sessions

192.168.0.1
  From: 192.168.6.1,  LSPstate: Up, ActiveRoute: 0
   LSPname: lsp1-r6-to-r0, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 101952
  Resv style: 1 SE, Label in: 100464, Label out: 101952
  Time left:  157, Since: Tue Jul 11 14:31:38 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 11131 protocol 0
   Node/Link protection desired
   Type: Node/Link protected LSP, using Bypass->10.0.12.14->10.0.24.2
       1 Jul 11 14:31:38 Node protection up, using Bypass->10.0.12.14->10.0.24.2
  PATH rcvfrom: 10.0.16.2 (so-0/0/3.0) 509 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 10.0.12.14 (fe-0/1/0.0) 356 pkts
  RESV rcvfrom: 10.0.12.14 (fe-0/1/0.0) 358 pkts
  Explct route: 10.0.12.14 10.0.24.2 10.0.45.2 10.0.50.2
  Record route: 10.0.16.2 <self> 10.0.12.14 10.0.24.2 10.0.45.2 10.0.50.2

192.168.6.1
  From: 192.168.0.1,  LSPstate: Up, ActiveRoute: 0
  LSPname: r0-to-t6, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: 100448, Label out: 3
  Time left:  147, Since: Tue Jul 11 14:31:36 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 23481 protocol 0
  PATH rcvfrom: 10.0.12.14 (fe-0/1/0.0) 358 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 10.0.16.2 (so-0/0/3.0) 350 pkts
  RESV rcvfrom: 10.0.16.2 (so-0/0/3.0) 323 pkts
  Explct route: 10.0.16.2
  Record route: 10.0.50.2 10.0.45.2 10.0.24.2 10.0.12.14 <self> 10.0.16.2
Total 2 displayed, Up 2, Down 0
```

Meaning Sample output from **R1** for the **show mpls lsp extensive** command shows detailed information about all LSPs for which **R1** is the ingress, egress, or transit router, including all past state history and the reason why an LSP failed. All LSPs are up. The main two LSPs **lsp2-r1-to-r5** and **lsp1-r6-to-r0** have node-link protection as indicated by the **Node/Link protection desired** field in the ingress and transit sections of the output. In the ingress section of the output, the **Link-protection Up** field shows that **lsp2-r1-to-r5** has link protection up. In the transit section of the output, the **Type: Node/Link protected LSP** field shows that **lsp1-r6-to-r0** has node-link protection up, and in case of failure will use the bypass **LSP Bypass->10.0.12.14->10.0.24.2.**

Sample Output

```
user@R1> show rsvp interface
RSVP interface: 4 active
                  Active Subscr- Static      Available   Reserved    Highwater
Interface   State resv   iption  BW          BW          BW          mark
fe-0/1/0.0 Up    2    100%  100Mbps     100Mbps     0bps        0bps
```

```
fe-0/1/1.0  Up          1   100%  100Mbps     100Mbps     0bps        0bps
fe-0/1/2.0  Up          0   100%  100Mbps     100Mbps     0bps        0bps
so-0/0/3.0  Up          1   100%  155.52Mbps  155.52Mbps  0bps        0bps
```

Meaning Sample output from **R1** for the **show rsvp interface** command shows four interfaces enabled with RSVP (**Up**). Interface **fe-0/1/0.0** has two active RSVP reservations (**Active resv**) that might indicate sessions for the two main LSPs, **lsp1-r6-to-r0** and **lsp2-r1-to-r5.** Interface **fe-0/1/0.0** is the connecting interface between R1 and R2, and both LSPs are configured with a strict path through **fe-0/1/0.0**. For more detailed information about what is happening on interface **fe-0/1/0.0**, issue the **show rsvp interface extensive** command.

Sample Output

```
user@R1> show rsvp interface extensive
RSVP interface: 3 active
fe-0/1/0.0 Index 67, State Ena/Up
  NoAuthentication, NoAggregate, NoReliable, LinkProtection
  HelloInterval 9(second)
  Address 10.0.12.13
  ActiveResv 2, PreemptionCnt 0, Update threshold 10%
  Subscription 100%,
  bc0 = ct0, StaticBW 100Mbps
  ct0: StaticBW 100Mbps, AvailableBW 100Mbps
    MaxAvailableBW 100Mbps = (bc0*subscription)
    ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
   Protection: On, Bypass: 2, LSP: 2, Protected LSP: 2, Unprotected LSP: 0
      2 Jul 14 14:49:40 New bypass Bypass->10.0.12.14
      1 Jul 14 14:49:34 New bypass Bypass->10.0.12.14->10.0.24.2
     Bypass: Bypass->10.0.12.14, State: Up, Type: LP, LSP: 0, Backup: 0
      3 Jul 14 14:49:42 Record Route:  10.0.17.14 10.0.27.1
      2 Jul 14 14:49:42 Up
      1 Jul 14 14:49:42 CSPF: computation result accepted
     Bypass: Bypass->10.0.12.14->10.0.24.2, State: Up, Type: NP, LSP: 2, Backup:0
      4 Jul 14 14:50:04 Record Route:  10.0.17.14 10.0.79.2 10.0.59.1 10.0.45.1
      3 Jul 14 14:50:04 Up
      2 Jul 14 14:50:04 CSPF: computation result accepted
      1 Jul 14 14:49:34 CSPF failed: no route toward 10.0.24.2
[...Output truncated...]
```

Meaning Sample output from **R1** for the **show rsvp interface extensive** command shows more detailed information about the activity on all RSVP interfaces (**3**). However, only output for **fe-0/1/0.0** is shown. Protection is enabled (**Protection: On**), with two bypass paths (**Bypass: 2**) protecting two LSPs (**Protected LSP: 2**). All LSPs are protected, as indicated by the **Unprotected LSP: 0** field. The first bypass **Bypass->10.0.12.14**is a link protection bypass path (**Type: LP**), protecting the link between **R1** and **R2 fe-0/1/0.0.** The second bypass path **10.0.12.14->10.0.24.2** is a node-link protected LSP, avoiding **R2** in case of node failure.

Sample Output

```
user@R1> show rsvp session detail
Ingress RSVP: 2 sessions

192.168.4.1
  From: 192.168.1.1,  LSPstate: Up, ActiveRoute: 0
   LSPname: Bypass->10.0.12.14->10.0.24.2
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 102000
```



```
  Resv style: 1 SE, Label in: -, Label out: 102000
  Time left:    -, Since: Tue Jul 11 14:30:53 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 60120 protocol 0
   Type: Bypass LSP
    Number of data route tunnel through: 2
    Number of RSVP session tunnel through: 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.17.14 (fe-0/1/1.0) 336 pkts
  RESV rcvfrom: 10.0.17.14 (fe-0/1/1.0) 310 pkts
   Explct route: 10.0.17.14 10.0.79.2 10.0.59.1 10.0.45.1
  Record route: <self> 10.0.17.14 10.0.79.2 10.0.59.1 10.0.45.1

192.168.5.1
  From: 192.168.1.1,  LSPstate: Up, ActiveRoute: 0
   LSPname: lsp2-r1-to-r5, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 101872
  Resv style: 1 SE, Label in: -, Label out: 101872
  Time left:    -, Since: Tue Jul 11 14:28:28 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 2 receiver 60118 protocol 0
   Node/Link protection desired
  Type: Node/Link protected LSP
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.12.14 (fe-0/1/0.0) 344 pkts
  RESV rcvfrom: 10.0.12.14 (fe-0/1/0.0) 349 pkts
   Explct route: 10.0.12.14 10.0.24.2 10.0.45.2
  Record route: <self> 10.0.12.14 10.0.24.2 10.0.45.2
Total 2 displayed, Up 2, Down 0

Egress RSVP: 1 sessions

192.168.1.1
  From: 192.168.5.1,  LSPstate: Up, ActiveRoute: 0
  LSPname: r5-to-r1, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left:  147, Since: Tue Jul 11 14:28:36 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 29228 protocol 0
  PATH rcvfrom: 10.0.12.14 (fe-0/1/0.0) 348 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.0.45.2 10.0.24.2 10.0.12.14 <self>
Total 1 displayed, Up 1, Down 0

Transit RSVP: 2 sessions

192.168.0.1
  From: 192.168.6.1,  LSPstate: Up, ActiveRoute: 0
   LSPname: lsp1-r6-to-r0, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 101952
  Resv style: 1 SE, Label in: 100464, Label out: 101952
```

```
  Time left:  134, Since: Tue Jul 11 14:31:38 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 11131 protocol 0
   Node/Link protection desired
  Type: Node/Link protected LSP
  PATH rcvfrom: 10.0.16.2 (so-0/0/3.0) 488 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 10.0.12.14 (fe-0/1/0.0) 339 pkts
  RESV rcvfrom: 10.0.12.14 (fe-0/1/0.0) 343 pkts
   Explct route: 10.0.12.14 10.0.24.2 10.0.45.2 10.0.50.2
  Record route: 10.0.16.2 <self> 10.0.12.14 10.0.24.2 10.0.45.2 10.0.50.2

192.168.6.1
  From: 192.168.0.1,  LSPstate: Up, ActiveRoute: 0
  LSPname: r0-to-t6, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: 100448, Label out: 3
  Time left:  158, Since: Tue Jul 11 14:31:36 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 23481 protocol 0
  PATH rcvfrom: 10.0.12.14 (fe-0/1/0.0) 344 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 10.0.16.2 (so-0/0/3.0) 337 pkts
  RESV rcvfrom: 10.0.16.2 (so-0/0/3.0) 310 pkts
  Explct route: 10.0.16.2
  Record route: 10.0.50.2 10.0.45.2 10.0.24.2 10.0.12.14 <self> 10.0.16.2
Total 2 displayed, Up 2, Down 0
```

Meaning Sample output from **R1** shows detailed information about the RSVP sessions active on **R1**. All sessions are up, with two ingress sessions, one egress session, and two transit sessions.

Within the ingress section, the first session is a bypass path, as indicated by the **Type: Bypass LSP** field; and the second session is a protected LSP (**lsp2-r1-to-r5)** originating on **R1**, as indicated by the **Type: Node/Link protected LSP** field.

Conclusion Multiprotocol Label Switching (MPLS) label-switched path (LSP) link protection and node-link protection are facility-based methods used to reduce the amount of time needed to reroute LSP traffic. These protection methods are often compared to fast reroute—the other Junos OS LSP protection method.

While fast reroute protects LSPs on a one-to-one basis, link protection and node-link protection protect multiple LSPs by using a single, logical bypass LSP. Link protection provides robust backup support for a link, node-link protection bypasses a node or a link, and both types of protection are designed to interoperate with other vendor equipment. Such functionality makes link protection and node-link protection excellent choices for scalability, redundancy, and performance in MPLS-enabled networks.

Related Information For additional information about MPLS fast reroute and MPLS protection methods, see the following:

- *Junos Feature Guide*
- *Junos MPLS Applications Configuration Guide*

- Semeria, Chuck. *RSVP Signaling Extensions for MPLS Traffic Engineering.* White paper. 2002
- Semeria, Chuck. *IP Dependability: Network Link and Node Protection*. White paper. 2002
- RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*

CHAPTER 4

RSVP Reservation Styles in an MPLS Network

A Resource Reservation Protocol (RSVP) reservation style is a request for a bandwidth reservation that includes a set of options. The options are described in the three reservation styles: fixed filter (FF), shared explicit (SE), and wildcard filter (WF) that determine how senders, receivers, and sessions are treated. The reservation style also determines how RSVP signaling reroutes an existing label-switched path (LSP). The egress router must establish each LSP with one of the three reservation styles.

This chapter describes explicitly routed LSPs that are established using the FF or SE styles. The RSVP WF reservation style is not used for explicitly routed LSPs because of its lack of applicability for traffic engineering.

The terms *node* and *router* are used interchangeably throughout this book.

Checklist for RSVP Reservation Styles

This checklist provides the steps and commands for working with RSVP reservation styles, specifically the fixed filter and shared explicit styles. In addition, the checklist provides links to overview information about RSVP reservation styles and detailed information about the commands used for configuring and verifying an adaptive label-switched path (LSP). (See Table 9 on page 51)

Table 9: Checklist for RSVP Reservation Styles

Tasks	Command or Action
"RSVP Reservation Styles Overview" on page 52	
"Fixed Filter Style Overview" on page 53	**show configuration protocols mpls** **show rsvp session detail** **show mpls lsp extensive**
"Shared Explicit Style Overview" on page 55	

Table 9: Checklist for RSVP Reservation Styles *(continued)*

Tasks	Command or Action
"Configuring and Verifying an Adaptive LSP" on page 56	**[edit]** **edit protocol mpls** **[edit protocols mpls]** **set label-switched-path *lsp-path-name* adaptive** **show** **commit**
"Rerouting the LSP Tunnel for the SE Reservation Style" on page 60	
"Establish the Initial LSP Tunnel" on page 60	Not applicable.
"Reroute an LSP Tunnel" on page 61	Not applicable.

RSVP Reservation Styles Overview

RSVP was originally a protocol for resource reservation. In the context of resource reservation, different reservation styles were developed to determine the degree to which resources are shared; the FF and SE reservation styles.

The FF reservation style dedicates a particular reservation to an individual sender (ingress router). This reservation style is useful for concurrent and independent traffic from different senders. When used with Multiprotocol Label Switching (MPLS), the FF reservation style allows the establishment of multiple parallel unicast point-to-point LSPs to support load balancing. It can also be used with primary and secondary paths to achieve minimal disruption to traffic. Examples of applications that use FF-style reservations are video applications and unicast applications, which both require flows that have a separate reservation for each sender.

The SE reservation style allows an explicit list of senders to share the largest bandwidth request across shared links. In an MPLS environment, this style is important for rerouting LSPs with no disruption to the flow of subscriber traffic. An example application for shared explicit reservations is an audio application in which each sender transmits a distinct data stream. Typically, only a few senders are transmitting at any one time. Such a flow does not require a separate reservation for each sender; a single reservation is sufficient.

In RSVP with traffic engineering, the ingress router can request the SE style by setting the appropriate bit in the Session Attribute object. If the Session Attribute object is present but the particular bit is *not* present, the egress router can use either style (FF or SE). All values in the Session Attribute object are advisory, so an egress router can ignore the bits when it selects a style; however, to date, this behavior has not been implemented. Selection of a style can be determined by non-support of a particular style, an explicit policy, or available resources.

In the context of traffic engineering, FF is the default reservation style. The SE style allows an LSP to share reservations which is useful when the ingress router is trying to set up

an alternate path before tearing down the existing path. Clearly traffic is sent on the active path only, but from the point of view of reservations, sharing resources avoids double counting the resources.

Fixed Filter Style Overview

The FF reservation style specifies an explicit list of senders and a distinct bandwidth reservation for each sender. The distinct bandwidth reservation is not shared with other senders, and is identified by an IP address and a local identification number (LSP_ID). Because each sender has its own particular reservation, a unique label and a separate LSP are constructed for each sender-receiver pair.

In RSVP with traffic engineering, each sender and receiver represent a different sender or receiver on a router, not necessarily different end systems. (See Figure 8 on page 53).

Figure 8: Fixed Filter Reservation Style

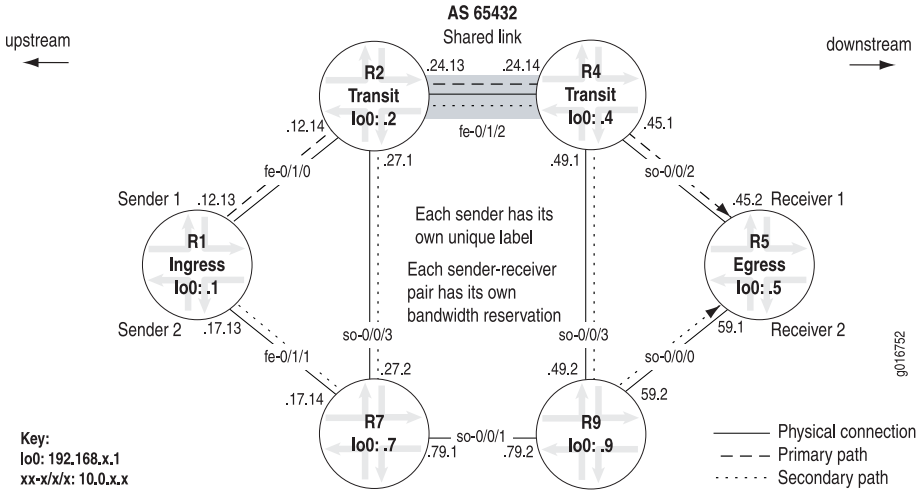

Figure 8 on page 53 shows a primary and secondary path that share the Fast Ethernet link **fe-0/1/2** between **R2** and **R4**. Each path has a separate RSVP session in the FF reservation style. When sessions share a link, the total amount of reserved bandwidth on the shared link is the sum of the reservations for each individual session. If the sum of reservations is larger than the available bandwidth, the LSP cannot be established, as illustrated in the example network in Figure 8 on page 53.

In the example network in Figure 8 on page 53, **R1** requests a 75-Mbps bandwidth reservation for all configured primary and secondary paths. Therefore, to establish a primary and standby secondary path, a 150-Mbps bandwidth reservation is required. Because the Fast Ethernet link has a total of 100 Mbps of bandwidth available, 75 Mbps of which is reserved for the primary path, leaving 25 Mbps for the standby secondary path, the standby secondary path cannot be established.

Action For an illustration of this situation, see the output for the following commands:

show configuration protocols mpls (See Sample Output on page 54)
show rsvp session detail (See Sample Output on page 54)
show mpls lsp extensive (See Sample Output on page 55)

Sample Output

```
user@R1>show configuration protocols mplsbandwidth 75m;
label-switched-path lsp1 {
    to 192.168.5.1;
     primary via-r2;
    secondary via-r7 {
         standby;
    }
}
path via-r7 {
    10.0.17.14 strict;
    10.0.27.1 strict;
     10.0.24.14 strict;
    10.0.49.2 strict;
}
path via-r2 {
    10.0.12.14 strict;
     10.0.24.14 strict;
}
interface fe-0/1/0.0;
interface fe-0/1/1.0;
interface so-0/0/3.0;
```

Meaning

Sample output from **R1** for the **show configuration protocols mpls** command shows the MPLS configuration that includes a bandwidth of 75 Mbps for all paths, LSP **lsp1**, a primary path, and a standby secondary path. Both named paths, **path via-r7** and **path via-r2**, specify all transit routers up to the egress. The egress router is not specified. Both paths are strict, indicating that the route taken from one router to the next router is a direct path and cannot include any other routers. All specified addresses are interface addresses, ensuring that the incoming interface is the one specified and enforcing routing on a per-link basis.

From the network topology shown in Figure 8 on page 53, the link shared by both paths is from **R2** to **R4, fe-0/1/2,** or address **10.0.24.14**.

Sample Output

```
user@R1>  show rsvp session detail
Ingress RSVP: 1 sessions

192.168.5.1
  From: 192.168.1.1, LSPstate: Up, ActiveRoute: 0
   LSPname: lsp1, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 102720
   Resv style: 1 FF, Label in: -, Label out: 102720
  Time left:    -, Since: Fri Jul 21 11:08:12 2006
  Tspec: rate 75Mbps size 75Mbps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 60165 protocol 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.12.14 (fe-0/1/0.0) 6 pkts
  RESV rcvfrom: 10.0.12.14 (fe-0/1/0.0) 6 pkts
  Explct route: 10.0.12.14 10.0.24.14 10.0.45.2
  Record route: <self> 10.0.12.14 10.0.24.14 10.0.45.2
Total 1 displayed, Up 1, Down 0
[...Output truncated...]
```

Meaning The sample output from **R1** for the **show rsvp session detail** command shows that **R1** has one ingress RSVP session established in the FF style and associated with the primary path, indicating that the standby secondary path is not established. If the secondary standby path was established, we would expect to see two ingress sessions, one for the primary path and another for the secondary standby path.

Sample Output

```
user@R1> show mpls lsp extensive

Ingress LSP: 1 sessions

192.168.5.1
  From: 192.168.1.1, State: Up, ActiveRoute: 0,  LSPname: lsp1
   ActivePath: via-r2 (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary  via-r2      State: Up
     Bandwidth: 75Mbps
     SmartOptimizeTimer: 180
     Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
 10.0.12.14 S 10.0.24.14 S 10.0.45.2 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

          10.0.12.14 10.0.24.14 10.0.45.2
     5 Jul 21 11:08:12 Selected as active path
     4 Jul 21 11:08:12 Record Route:  10.0.12.14 10.0.24.14 10.0.45.2
     3 Jul 21 11:08:12 Up
     2 Jul 21 11:08:12 Originate Call
     1 Jul 21 11:08:12 CSPF: computation result accepted
    Standby  via-r7      State: Dn      Bandwidth: 75Mbps
     SmartOptimizeTimer: 180
         No computed ERO.
  Created: Fri Jul 21 11:08:11 2006
Total 1 displayed, Up 1, Down 0
[...Output truncated...]
```

Meaning Sample output from **R1** for the **show mpls lsp extensive** command shows that 75 Mbps of bandwidth is allocated for each path. The secondary standby path is down (**State: Dn**) because there is not enough available bandwidth.

Shared Explicit Style Overview

The SE RSVP reservation style creates shared reservations among explicit senders. For a single RSVP bandwidth reservation, the egress router (receiver) lists the senders sharing the reservation in a Resv message, resulting in the following:

- A multipoint-to-point LSP if the Path message does not contain an ERO or if the ERO is identical across senders. A common label is assigned.
- A separate LSP for each sender if the Path message contains a different ERO for each sender. A different label is assigned to different senders.
- Each LSP shares the bandwidth of the largest request across the shared link.

While any LSP can be established with an SE style reservation, the SE reservation is most useful during LSP reroute, for example, when a standby secondary path or link protection

is configured. In general, the secondary LSP inherits the reservation style of the primary LSP, which is FF by default, or SE if link protection is used, unless the secondary LSP is configured with the **adaptive** statement at the secondary path level. See Figure 9 on page 56.

Figure 9: Shared Explicit Style

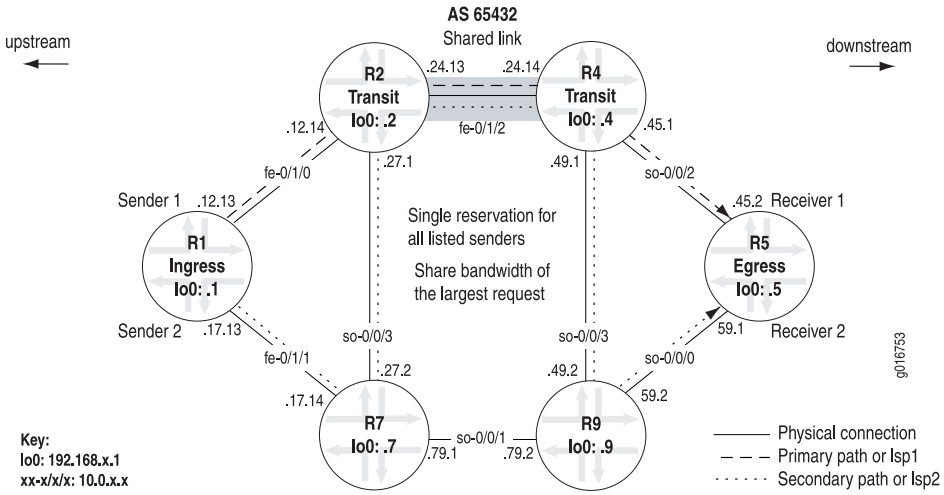

The network shown in Figure 9 on page 56 shows two paths, **R1-R2-R4-R5** and **R1-R7-R9-R5**. The paths are configured as either a strict primary or strict standby secondary path for an adaptive LSP, or as **lsp1** and **lsp2**. Both configurations originate from **R1** and share the link between **R2** and **R4**. For more information about adaptive LSPs, see “Configuring and Verifying an Adaptive LSP” on page 56.

If a network problem results in an LSP reroute, the SE reservation style allows a smooth transition from either a primary path to a standby secondary path, or from on old LSP to a new LSP with the make-before-break operation. This style also permits the old and new LSPs to share a single reservation over links they have in common, preventing double counting of resources.

Configuring and Verifying an Adaptive LSP

When you include the **adaptive** statement in the configuration, the LSP becomes adaptive and is established with the SE reservation style. The **adaptive** statement can be configured at two hierarchy levels:

- The **[edit protocols mpls label-switched-path *lsp-path-name*]** hierarchy level, which keeps the RSVP session information the same for all primary and secondary paths.
- The **[edit protocols mpls label-switched-path *lsp-path-name* secondary *secondary-name*]** hierarchy level, resulting in different Tunnel ID values for each path and causes the paths to be viewed as separate RSVP sessions, that may not share the same bandwidth reservation and possibly double-count resources.

Using an adaptive LSP at the **[edit protocols mpls label-switched-path *lsp-path-name*]** hierarchy level provides two advantages. The first advantage is the prevention of double-counting of bandwidth for links that share old and new paths. Double-counting

occurs when an intermediate router does not recognize that the new and old paths belong to the same LSP and counts them as two separate LSPs, requiring separate bandwidth allocations. If some links are close to saturation, double-counting might cause the setup of the new path to fail. When the **adaptive** statement is included at the **[edit protocols mpls label-switched-path *lsp-path-name*]** hierarchy level, a standby secondary path is established, sharing physical links in common with the LSP's primary path.

The second advantage is the prevention of disruption to subscriber traffic by performing a make-before-break operation. When an established path attempts to reroute onto a new path, the ingress router maintains existing paths and allocated bandwidths, ensuring that the existing path is not prematurely torn down and allowing the current traffic to continue flowing while the new path is set up.

The following steps describe the process of configuring an adaptive LSP that keeps the RSVP session information the same for all primary and secondary paths. Before you can configure an adaptive LSP, you must have an LSP already configured with the primary and secondary paths you want to use, and any other options. For information on configuring a LSP with a primary path and secondary path, see "Checklist for Path Protection" on page 9.

Action To configure an adaptive LSP, follow these steps:

1. In configuration mode, go to the following hierarchy level:

   ```
   [edit]
   user@R1# edit protocol mpls
   ```

2. Configure adaptive mode for the LSP:

   ```
   [edit protocols mpls]
   user@R1# set label-switched-path lsp-path-name adaptive
   ```

   For example:

   ```
   [edit protocols mpls]
   user@R1# set label-switched-path lsp1 adaptive
   ```

3. Verify and commit the configuration:

   ```
   [edit protocols mpls]
   user@R1# show
   user@R1# commit
   ```

Sample Output

```
[edit protocols mpls]
user@R1# show
bandwidth 75m;
label-switched-path  lsp1 {
    to 192.168.5.1;
    adaptive;
    primary via-r2;
    secondary via-r7 {
        standby;
    }
}
path via-r7 {
    10.0.17.14 strict;
    10.0.27.1 strict;
```



```
    10.0.24.14 strict;
    10.0.49.2 strict;
}
path via-r2 {
    10.0.12.14 strict;
    10.0.24.14 strict;
}
interface fe-0/1/0.0;
interface fe-0/1/1.0;
interface so-0/0/3.0;

[edit protocols mpls]
user@R1# commit
commit complete
```

Meaning Sample output from **R1** for the **show** command shows bandwidth of 75 Mbps, the **adaptive** statement, and strict primary and secondary paths. 75 Mbps of bandwidth for each path is more combined bandwidth than the Fast Ethernet link **fe-0/1/2** can accommodate. Because **lsp1** is adaptive, both paths are up, indicating that the bandwidth is not double-counted, as shown in the following output for the **show mpls lsp extensive** command.

Sample Output

```
[edit protocols mpls]
user@R1# run show mpls lsp extensive
Ingress LSP: 1 sessions
192.168.5.1
  From: 192.168.1.1, State: Up, ActiveRoute: 0,  LSPname: lsp1
  ActivePath: via-r2 (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary  via-r2      State: Up
     Bandwidth: 75Mbps
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
 10.0.12.14 S 10.0.24.14 S 10.0.45.2 S
   Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

          10.0.12.14 10.0.24.14 10.0.45.2
    5 Jul 21 14:34:16 Selected as active path
    4 Jul 21 14:34:16 Record Route:  10.0.12.14 10.0.24.14 10.0.45.2
    3 Jul 21 14:34:16 Up
    2 Jul 21 14:34:16 Originate Call
    1 Jul 21 14:34:16 CSPF: computation result accepted
   Standby  via-r7      State: Up
     Bandwidth: 75Mbps
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 5)
 10.0.17.14 S 10.0.27.1 S 10.0.24.14 S 10.0.49.2 S 10.0.59.1 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node ...
          10.0.17.14 10.0.27.1 10.0.24.14 10.0.49.2 10.0.59.1
    4 Jul 21 14:34:45 Record Route:  10.0.17.14 10.0.27.1 10.0.24.14 10.0.49.2
10.0.59.1
    3 Jul 21 14:34:45 Up
    2 Jul 21 14:34:45 Originate Call
    1 Jul 21 14:34:45 CSPF: computation result accepted
  Created: Fri Jul 21 14:34:15 2006
Total 1 displayed, Up 1, Down 0
```


Meaning The sample output from **R1** for the **show mpls lsp extensive** command shows that **lsp1** is up with an active primary path that is up **(*Primary via-r2 State: Up)**, and a standby secondary path that is also up (**Standby via-r7 State: Up)**. Both paths have 75 Mbps of bandwidth, which is not double-counted because the **adaptive** statement ensures that new and old paths are recognized as belonging to the same LSP **lsp1**, as shown in the following sample output for the **show rsvp session detail** command. You can also use the **show rsvp interface** command to show the reserved and available bandwidth.

Sample Output

```
user@R1> show rsvp session detail
Ingress RSVP: 2 sessions

192.168.5.1
  From: 192.168.1.1 , LSPstate: Up, ActiveRoute: 0
   LSPname: lsp1, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 102736
   Resv style: 1 SE,  Label in: -, Label out: 102736
  Time left:    -, Since: Fri Jul 21 14:34:16 2006
  Tspec: rate 75Mbps size 75Mbps peak Infbps m 20 M 1500
   Port number: sender 1 receiver 60167 protocol 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.12.14 (fe-0/1/0.0) 7 pkts
  RESV rcvfrom: 10.0.12.14 (fe-0/1/0.0) 7 pkts
  Explct route: 10.0.12.14 10.0.24.14 10.0.45.2
  Record route: <self> 10.0.12.14 10.0.24.14 10.0.45.2

192.168.5.1
  From: 192.168.1.1 , LSPstate: Up, ActiveRoute: 0
  LSPname: lsp1, LSPpath: Secondary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 102608
   Resv style: 1 SE , Label in: -, Label out: 102608
  Time left:    -, Since: Fri Jul 21 14:34:45 2006
  Tspec: rate 75Mbps size 75Mbps peak Infbps m 20 M 1500
   Port number: sender 2 receiver 60167 protocol 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.17.14 (fe-0/1/1.0) 5 pkts
  RESV rcvfrom: 10.0.17.14 (fe-0/1/1.0) 5 pkts
  Explct route: 10.0.17.14 10.0.27.1 10.0.24.14 10.0.49.2 10.0.59.1
  Record route: <self> 10.0.17.14 10.0.27.1 10.0.24.14 10.0.49.2 10.0.59.1
Total 2 displayed, Up 2, Down 0
[...Output truncated...]
```

Meaning The sample output from **R1** for the **show rsvp session detail** command shows two RSVP sessions for **lsp1**. Both sessions originate on **R1** (**192.168.1.1**) and end in **R5** (**192.168.5.1**). The first session is for the primary path and the second session is for the secondary path. Both paths are in the SE reservation style. The port number is the protocol ID and sender/receiver port used in this RSVP session. In the port number field, the primary session shows **sender 1**, while the secondary session shows **sender 2**, indicating that two senders are using the LSP tunnel.

Rerouting the LSP Tunnel for the SE Reservation Style

An LSP tunnel may need to be rerouted due to conditions based on administrative policy, for example, when a more optimal route becomes available, when a resource fails along the LSP, or when a failed resource is reactivated. The SE reservation style allows a smooth transition from an old LSP to a new LSP with the make-before-break operation. This style also permits the old and new LSPs to share a single reservation over links they have in common, preventing double-counting of resources.

Related Topics

- Establish the Initial LSP Tunnel on page 60
- Reroute an LSP Tunnel on page 61
- *Junos Feature Guide*
- *Junos MPLS Applications Configuration Guide*
- Semeria, Chuck. *RSVP Signaling Extensions for MPLS Traffic Engineering.* White paper. 2002
- Semeria, Chuck. *IP Dependability: Network Link and Node Protection*. White paper. 2002
- RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*

Establish the Initial LSP Tunnel

The ingress router uses the Path message to request that the egress router set up the initial LSP tunnel with the SE reservation style. When establishing the initial LSP tunnel, the ingress and egress routers perform the following actions:

1. The ingress router includes the following in the initial Path message:
   - A LSP Tunnel IPv4 Session object that contains the following:
     - IPv4 address of the egress node.
     - Tunnel ID that remains constant for the life of the LSP tunnel between the ingress and egress routers.
     - Extended Tunnel ID that identifies the ingress router's IPv4 address.
   - The SE reservation style in the Session Attribute object.
   - A Sender Template object that contains the following:
     - The IPv4 address of the sender (ingress) node.
     - A LSP ID that can change in the future, when the LSP needs to be rerouted, allowing the ingress router to appear as a different sender so it can share resources with itself (see the **LSP ID** field of the LSP Tunnel IPv4 C-type extension for the Sender Template and Filter Spec objects).
2. Upon receipt of the Path message, the egress router sends a Resv message with an SE reservation style toward the ingress router.

3. When the ingress router receives the Resv message, the initial LSP tunnel is established with an SE reservation style.

Related Documentation

- Rerouting the LSP Tunnel for the SE Reservation Style on page 60
- Reroute an LSP Tunnel on page 61

Reroute an LSP Tunnel

When the ingress router attempts to reroute an exiting LSP tunnel to increase the bandwidth or change the path, it transmits a new Path message. During the reroute operation, the ingress router must appear as two different senders to the RSVP session. This is achieved by including a new LSP ID in the Sender Template object and the Filter Spec object. The ingress and egress routers perform the following actions:

1. The ingress router includes the following in the new Path message:
   - An Explicit Route object (ERO) for the new LSP tunnel.
   - The existing LSP Tunnel IPv4 Session object to identify the LSP that will be rerouted.
   - A new LSP ID and a new Sender Template object, ensuring that the ingress router appears as a different sender to the RSVP session.
2. The ingress router transmits the new Path message toward the egress router, continuing to use the old LSP tunnel to forward traffic and continuing to refresh the original PATH message (make-before-break).
3. The egress router responds to the new Path message with a Resv message that contains a number of RSVP objects, including:
   - A Label object to support the upstream on-demand label distribution process
   - An SE reservation Style object

   

   NOTE: On links that are not shared by the old and new LSP tunnels, the new Path/Resv message pair is treated as a new conventional LSP. However, on links that are traversed by both the old and new LSP tunnels, the LSP Tunnel IPv4 Session object and SE reservation style allow the new LSP tunnel to establish so that it shares resources with the old LSP tunnel, eliminating the double-counting problem on shared links.
4. The ingress router begins to use the new LSP tunnel after it receives the new Resv message.
5. The ingress router sends a Path Tear message to remove the old LSP tunnel from intermediate routers.

Related Topics

For additional information about MPLS fast reroute and MPLS protection methods, see the following:

CHAPTER 5

Load Balancing in an MPLS Network

In an Multiprotocol Label Switched (MPLS) network, load balancing is the process of distributing traffic equally across label switched paths. When you have added several LSPs to the same egress router, the default behavior of the Junos OS is to select the LSP with the lowest metric to carry all traffic. If all of the LSPs have the same metric, one of the LSPs is selected at random and all traffic is forwarded over it. You can change this default behavior by configuring load balancing on an ingress router, allowing the Junos OS to distribute the traffic equally across LSPs.

The terms *node* and *router* are used interchangeably throughout this book.

Checklist for Load Balancing in an MPLS Network

The checklist for load balancing provides the steps and commands to load balance traffic across an MPLS Network. The checklist includes links to an overview of load balancing as implemented in the Junos OS, more detailed information about the commands to configure and verify load balancing, and to network examples of load balancing. The network examples include using a hash-key and bandwidth to load balance. (See Table 10 on page 63.)

Table 10: Checklist for Load Balancing in an MPLS Network

Tasks	Command or Action
"Load Balancing Overview" on page 65	The overview includes when to configure load balancing, methods of load balancing and the parameters of load balancing. The following load-balancing options are also included: • IPv4 address family (INET) in the hash key • MPLS labels and IP payload in the hash key • LSP bandwidth
"Configuring and Verifying Load Balancing" on page 67	

Table 10: Checklist for Load Balancing in an MPLS Network *(continued)*

Tasks	Command or Action
1. Define a Load-Balancing Policy on page 67	[edit] edit policy-options [edit policy-options] set policy-statement *policy-name* then load-balance per-packet show commit
2. Apply the Load-Balancing Policy to the Forwarding Table on page 68	[edit] edit routing-options [edit routing-options] set forwarding-table export *policy-name* show commit
3. Verify That Load Balancing Is Working on page 69	show configuration show route show route forwarding-table show mpls lsp statistics monitor interface traffic clear mpls lsp statistics clear interface statistics
"Example: Load-Balanced MPLS Network" on page 72	
"Router Configurations for the Load-Balanced MPLS Network" on page 73	show configuration | no-more
"Using Hash-Key Load Balancing for LSP Traffic" on page 83	
1. Configuring MPLS Labels and IP Payload to Load-Balance LSP Traffic on page 84	[edit] edit forwarding-options hash-key [edit forwarding-options hash-key] set family mpls label-1 set family mpls label-2 set family mpls payload ip show commit
2. Configuring the IPv4 Address Family to Load-Balance LSP Traffic on page 86	[edit] edit forwarding-options hash-key [edit forwarding-options hash-key] set family inet layer-3 set family inet layer-4 show commit
"Hash Key Network Examples" on page 88	
1. Example: Load-Balancing a Network with Aggregated Interfaces on page 88	

Table 10: Checklist for Load Balancing in an MPLS Network *(continued)*

Tasks	Command or Action
a. Verifying the Operation of Load Balancing with Aggregated Interfaces on page 89	**show configuration forwarding-options** **show interfaces statistics *interface-name* detail** **show mpls lsp statistics**
b. Router Configurations for the Aggregated Interfaces Network on page 93	**show configuration | no-more**
2. Example: Load-Balancing a Network Using INET in the Hash Key on page 100	
a. Verifying the Operation of INET Load Balancing on page 101	**show configuration** **show route forwarding-table destination *destination*** **show route** **monitor interface traffic** **show mpls lsp statistics**
b. Router Configurations for the INET Load-Balanced Network on page 103	**show configuration | no-more**
"Using Bandwidth to Unevenly Load-Balance RSVP LSPs" on page 113	
1. Configure Bandwidth to Unevenly Load-Balance Traffic on page 115	**[edit]** **edit protocols mpls** **[edit protocols mpls]** **set label-switched-path *lsp-path-name* bandwidth *bps* show** **[edit protocols rsvp]** **set load-balance bandwidth** **show** **commit**
2. Verify the Operation of Uneven Bandwidth Load Balancing on page 116	**show route protocol rsvp detail** **show mpls lsp statistics**
3. Router Configurations for Bandwidth Load Balancing on page 118	**show configuration | no-more**
"Traffic Flows Before Load Balancing" on page 120	**show route | find mpls** **monitor interface traffic** **show mpls lsp statistics**

Load Balancing Overview

In an MPLS network, load balancing is generally configured on an ingress router. The load-balancing configuration distributes traffic equally across LSPs with a hash algorithm that selects the next-hop destination and installs it into the forwarding table for the active route of the LSP. Whenever the next hop changes in any way, the hash algorithm changes the next-hop address.

Use load balancing when you have many LSPs with equal-cost next hops going out different interfaces to the same destination. For example, ingress router **R1** has four LSPs configured to egress router **R3**, transiting **R2**. All four LSPs have the same metric from **R2** to **R3**, but exit **R2** from different interfaces. Without load balancing configured, one of the LSPs is selected at random and all traffic is forwarded over it. With load balancing configured, traffic is balanced evenly across all four LSPs.

Depending on the version of the Internet Processor ASIC in the routing platform, a different method is used to load-balance traffic. Routing platforms with an Internet Processor I ASIC use a round-robin method, sending each packet over a different link. The round-robin method results in a good traffic balance at the cost of potential out-of-order packet arrival, which is not good for TCP.

Routing platforms with the Internet Processor II ASIC divide traffic into individual traffic flows across up to 16 next hops, keeping each individual flow on a single interface. A flow is comprised of packets with the following identical parameters:

- Source IP address
- Destination IP address
- Protocol
- Source port number
- Destination port number
- Source interface index
- Type of service (ToS)

For example, if there are 60,000 prefixes in a routing table, and 6 links between two routers, 10,000 prefixes will go across each link. In the core of the network, the law of averages produces good traffic load balancing. However, at the edge of the network, where there may not be a large number of prefixes, traffic may not be well balanced.

To summarize, an Internet Processor I ASIC spreads packets with the same parameters across multiple equal-cost next hops; while an Internet Processor II ASIC sends packets with the same parameters to the same next hop, since they are in the same flow. The Junos OS command to turn on load balancing uses the action **load-balance per-packet**, which is misnamed in relation to the Internet Processor II ASIC. On the Internet Processor II ASIC, this command actually enables per-flow load balancing.

To configure load balancing, include a policy statement at the **[edit policy-options]** hierarchy level. This policy statement must be applied as an export policy at the **[edit forwarding-options]** hierarchy level. For more information, see “Configuring and Verifying Load Balancing” on page 67.

Load-Balancing Options

After you configure load balancing, if the outbound traffic across equal-cost next hops is not balanced to your satisfaction, you can provide additional information to identify traffic flows and balance traffic more evenly or unevenly, depending on your requirements. You can provide additional information to alter the way load balancing works in the following ways:

- Include MPLS labels and IP payload statements in the hash-key configuration at the **[edit forwarding-options hash-key]** hierarchy level to evenly balance traffic across LSPs and aggregated interfaces. For more information, see "Configuring MPLS Labels and IP Payload to Load-Balance LSP Traffic" on page 84.
- Include the IPv4 address family (INET) in the hash-key configuration at the **[edit forwarding-options hash-key]** hierarchy level to ensure that traffic is evenly load-balanced across LSPs and interfaces, and that packets in the same flow are sent out through the same interface. For more information, see "Configuring the IPv4 Address Family to Load-Balance LSP Traffic" on page 86.
- Include bandwidth at the **[edit protocols mpls label-switched-path *lsp-name*]** and the **[edit protocols rsvp]** hierarchy levels to change the number of prefixes carried by an LSP and thereby create an uneven load balance for different LSPs. For more information, see "Using Bandwidth to Unevenly Load-Balance RSVP LSPs" on page 113.

Configuring and Verifying Load Balancing

Purpose Load balancing is configured on the ingress router and uses the hash algorithm to distribute traffic equally across paths. The hash algorithm is designed to distribute packets to prevent any single link from being saturated. Before you can configure and verify load balancing in an MPLS network, you must have all the necessary MPLS components and protocols configured correctly. For information on configuring an MPLS network, see the *Junos MPLS Network Operations Guide*.

Keep the following information in mind when you configure load balancing:

- The **load-balance per packet** policy is configured on an ingress router with more than one LSP configured to the same egress router.
- Load balancing offers no guarantee of equal distribution of traffic over equal-cost links, nor does it guarantee that increasing the number of Internet flows will create a better hash distribution.

To configure and verify load balancing, follow these steps:

1. Define a Load-Balancing Policy on page 67
2. Apply the Load-Balancing Policy to the Forwarding Table on page 68
3. Verify That Load Balancing Is Working on page 69

Define a Load-Balancing Policy

Purpose On the ingress or transit router, you can include a policy statement that performs load balancing on all routes. For information on including a policy statement that performs load balancing on specific routes, see "Configuring Per-Packet Load Balancing" in the *Junos Routing Protocols Configuration Guide*.

Action On the ingress or transit router, to define a load-balancing policy for all routes, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```
[edit]
user@host# edit policy-options
```

2. Define the load-balance policy and action:

```
[edit policy-options]
user@host# set policy-statement policy-name then load-balance per-packet
```

3. Verify and commit the configuration:

```
user@host# show
user@host# commit
```

Sample Output

```
user@R6> edit
Entering configuration mode

[edit]
user@R6# edit policy-options

[edit policy-options]
user@R6# set policy-statement load-balance-traffic then load-balance per-packet

[edit policy-options]
user@R6# show
policy-statement load-balance-traffic {
    then {
        load-balance per-packet;
    }
}

[edit policy-options]
user@R6# commit
commit complete
```

Meaning The sample output from ingress router **R6** shows the process for configuring load balancing. On an Internet Processor I ASIC, packets with the same parameters are spread across multiple equal-cost next hops; while an Internet Processor II ASIC sends packets with the same parameters to the same next hop, since they are in the same flow. The Junos OS command to turn on load balancing uses the action **load-balance per-packet**, which is misnamed in relation to the Internet Processor II ASIC. On the Internet Processor II ASIC, this command actually enables per-flow load balancing.

Apply the Load-Balancing Policy to the Forwarding Table

Purpose Apply the policy configured in Step 1 to routes exported from the routing table to the forwarding table.

Action To apply a load-balancing policy to the forwarding table, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```
[edit]
user@host# edit routing-options
```

2. Define a load-balance per packet action:

```
[edit routing-options]
user@host#set forwarding-table export policy-name
```


3. Verify and commit the configuration:

```
user@host# show
user@host# commit
```

Sample Output

```
[edit]
user@R6# edit routing-options

[edit routing-options]
user@R6# set forwarding-table export load-balance-traffic

[edit routing-options]
user@R6# show
static {
[...Output truncated...]
}
router-id 192.168.6.1;
autonomous-system 65432;
forwarding-table {
    export load-balance-traffic;
}

[edit routing-options]
user@R6# commit
commit complete
```

Meaning The sample output shows the process for applying a load-balancing policy to export routes from the routing table to the forwarding table.

Verify That Load Balancing Is Working

Purpose After configuring load balancing, check that traffic is load-balanced equally across paths. In this section, the command output reflects the load-balancing configuration of the example network shown in Figure 10 on page 72. The **clear** commands are used to reset LSP and interface counters to zero so that the values reflect the operation of the load-balancing configuration.

Action To verify load balancing across interfaces and LSPs, use the following command on the ingress router:

```
user@host# show configuration
```

To verify load balancing across interfaces and LSPs, use the following commands on a transit router:

```
user@host# show route
user@host# show route forwarding-table
user@host# show mpls lsp statistics
user@host# monitor interface traffic
user@host# clear mpls lsp statistics
user@host# clear interface statistics
```

Sample Output The following sample output is for the configuration on ingress router **R1**:

```
user@R1> show configuration | no-more
[...Output truncated...]
```

```
routing-options {
    [...Output truncated...]
     forwarding-table {
         export lbpp;
    }
}
[...Output truncated...]
policy-options {
      policy-statement lbpp {
          then {
               load-balance per-packet;
          }
    }
}
```

Meaning The sample output for the **show configuration** command on ingress router **R1** shows that load balancing is correctly configured with the **lbpp** policy statement. Also, the **lbpp** policy is exported into the forwarding table at the **[edit routing-options]** hierarchy level.

Sample Output The following sample output is from transit router R2:

```
user@R2> show route 192.168.0.1 terse

inet.0: 25 destinations, 27 routes (25 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 192.168.0.1/32     0  10          3             so-0/0/1.0
                                                 >so-0/0/2.0
[...Output truncated...]
```

Meaning The sample output for the **show route** command issued on transit router **R2** shows the two equal-cost paths (**so-0/0/1** and **so-0/0/2)** through the network to the loopback address to **R0** (**192.168.0.1**). Even though the right angle bracket (>) usually indicates the active route, in this instance it does not, as shown in the following four sample outputs.

Sample Output The following sample output is from transit router R2:

```
user@R2> monitor interface traffic

R2                                  Seconds: 65                 Time: 11:41:14

Interface    Link  Input packets        (pps)     Output packets        (pps)
 so-0/0/0      Up              0          (0)                  0          (0)
 so-0/0/1      Up            126          (0)             164659       (2128)
 so-0/0/2      Up          85219       (1004)             164598       (2128)
 so-0/0/3      Up              0          (0)                  0          (0)
 fe-0/1/0      Up         328954       (4265)              85475       (1094)
 fe-0/1/1      Up              0          (0)                  0          (0)
 fe-0/1/2      Up              0          (0)                  0          (0)
 fe-0/1/3      Up              0          (0)                  0          (0)
[...Output truncated...]
```

Meaning The sample output for the **monitor interface traffic** command issued on transit router **R2** shows that output traffic is evenly distributed across the two interfaces **so-0/0/1** and **so-0/0/2**.

Sample Output The following sample output is from transit router R2:


```
user@R2> show mpls lsp statistics
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 5 sessions
To              From            State     Packets            Bytes LSPname
192.168.0.1     192.168.1.1     Up          87997    17951388 lsp1
192.168.0.1     192.168.1.1     Up          87997    17951388 lsp2
192.168.0.1     192.168.1.1     Up          87997    17951388 lsp3
192.168.0.1     192.168.1.1     Up          87997    17951388 lsp4
192.168.6.1     192.168.0.1     Up              0           0 r0-r1
Total 5 displayed, Up 5, Down 0
```

Meaning The sample output for the **show mpls lsp statistics** command issued on transit router **R2** shows that output traffic is evenly distributed across the four LSPs configured on ingress router **R6**.

Sample Output The following sample output is from transit router R2:

```
user@R2> show route forwarding-table destination 10.0.90.14
Routing table: inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
10.0.90.12/30      user     0                    ulst 262144    6
                                                 ucst    345    5 so-0/0/1.0
                                                 ucst    339    2 so-0/0/2.0
```

Meaning The sample output for the **show route forwarding-table destination** command issued on transit router **R2** shows **ulst** in the **Type** field, which indicates that load balancing is working. The two unicast (**ucst)** entries in the **Type** field are the two next hops for the LSPs.

Sample Output The following sample output is from transit router R2:

```
user@R2> show route forwarding-table | find mpls
Routing table: mpls
MPLS:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    dscd    38     1
0                  user     0                    recv    37     3
1                  user     0                    recv    37     3
2                  user     0                    recv    37     3
100112             user     0                   Swap 100032       so-0/0/1.0
100128             user     0                   Swap 100048       so-0/0/1.0
100144             user     0 10.0.12.13        Swap 100096       fe-0/1/0.0
100160             user     0                   Swap 100112       so-0/0/2.0
100176             user     0                   Swap 100128       so-0/0/2.0
```

Meaning The sample output for the **show route forwarding-table | find mpls** command issued on transit router R2 shows the MPLS routing table that contains the labels received and used by this router to forward packets to the next-hop router. This routing table is used mostly on transit routers to route packets to the next router along an LSP. The first three labels in the **Destination** column (Label 0, Label 1, and Label 2) are automatically entered

by MPLS when the protocol is enabled. These labels are reserved MPLS labels defined in RFC 3032. Label 0 is the IPv4 explicit null label. Label 1 is the MPLS equivalent of the IP Router Alert label, and Label 2 is the IPv6 explicit null label.

The remaining five labels in the **Destination** column are nonreserved labels that the router uses to forward traffic, and the last column **Netif**, shows the interfaces used to send the labeled traffic. For nonreserved labels, the second **Type** column shows the operation performed on matching packets. In this example, all non-reserved packets are swapped for outgoing packet labels. For example, packets with the label **100112** have their label swapped for **100032** before they are pushed out of interface **so-0/0/1.0**.

Example: Load-Balanced MPLS Network

When you configure several RSVP LSPs to the same egress router, the LSP with the lowest metric is selected and carries all traffic. If all of the LSPs have the same metric, one of the LSPs is selected at random and all traffic is forwarded over it. To distribute traffic equally across all LSPs, you can configure load balancing on the ingress or transit routers, depending on the type of load balancing configured.

Figure 10 on page 72 illustrates an MPLS network with four LSPs configured to the same egress router (**R0**). Load balancing is configured on ingress router **R1**. The example network uses Open Shortest Path First (OSPF) as the interior gateway protocol (IGP) with OSPF area **0.0.0.0**. An IGP is required for the Constrained Shortest Path First (CSPF) LSP, which is the default for the Junos OS. In addition, the example network uses a policy to create BGP traffic.

Figure 10: Load-Balancing Network Topology

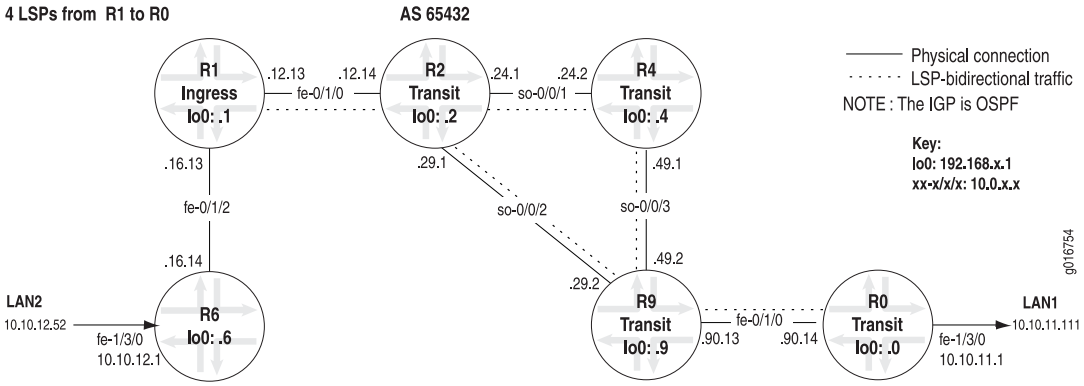

The network shown in Figure 10 on page 72 consists of the following components:

- A full-mesh interior BGP (IBGP) topology, using AS 65432
- MPLS and RSVP enabled on all routers
- A send-statics policy on routers **R1** and **R0** that allows a new route to be advertised into the network

- Four unidirectional LSPs between **R1** and **R0**, and one reverse direction LSP between **R0** and **R1**, which allows for bidirectional traffic
- Load balancing configured on ingress router **R1**

The network shown in Figure 10 on page 72 is a BGP full-mesh network. Since route reflectors and confederations are not used to propagate BGP learned routes, each router must have a BGP session with every other router running BGP.

For complete configurations for all routers in the example MPLS network, see "Router Configurations for the Load-Balanced MPLS Network" on page 73.

For a description of the situation before and after load balancing is configured in the network to use all four LSPs to forward traffic, see "Traffic Flows Before Load Balancing" on page 120.

Router Configurations for the Load-Balanced MPLS Network

Purpose The configurations in this topic are for the six routers in the example network illustrated in Figure 10 on page 72.

Action To display the configuration of a router, use the following Junos OS CLI operational mode command:

```
user@host> show configuration | no-more
```

Sample Output 1 The following configuration output is for edge router **R6**.

```
user@R6> show configuration | no-more
[...Output truncated...]
interfaces {
    fe-0/1/2 {
        unit 0 {
            family inet {
                address 10.0.16.14/30;
            }
             family mpls;  #MPLS enabled on relevant interfaces
        }
    }
    fe-1/3/0 {
        unit 0 {
            family inet {
                address 10.10.12.1/24;
            }
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.70.148/21;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.6.1/32;
```

```
            }
        }
    }
}
routing-options {
    static {
[...Output truncated...]
     router-id 192.168.6.1; #Manually configured RID
    autonomous-system 65432; #Full mesh IBGP
    }
}
protocols {
    rsvp {
        interface fe-0/1/2.0;
        interface fxp0.0 {
            disable;
        }
    }
    mpls {
        interface fe-0/1/2.0;
        interface fxp0.0 {
            disable;
        }
    }
    bgp {
        group internal {
            type internal;
            local-address 192.168.6.1;
            neighbor 192.168.1.1;
            neighbor 192.168.2.1;
            neighbor 192.168.4.1;
            neighbor 192.168.9.1;
            neighbor 192.168.0.1;
        }
    }
     ospf  { #IGP enabled
        traffic-engineering;
        area 0.0.0.0 {
            interface fe-0/1/2.0;
            interface fe-1/3/0.0;
            interface  lo0.0 {
                passive; #Ensures protocols do not run over this interface
            }
        }
    }
}
```

Sample Output 2 The following configuration output is for ingress router **R1**.

```
user@R1> show configuration | no-more
[...Output truncated...]
interfaces {
    fe-0/1/0 {
        unit 0 {
            family inet {
                address 10.0.12.13/30;
            }
             family mpls; #MPLS enabled on relevant interfaces
        }
    }
    fe-0/1/2 {
```



```
        unit 0 {
            family inet {
                address 10.0.16.13/30;
            }
            family mpls;
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.70.143/21;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.1.1/32;
            }
        }
    }
}
routing-options {
    static {
        [...Output truncated...]
     route 100.100.1.0/24 reject; #Static route for send-statics policy
}
     router-id 192.168.1.1; #Manually configured RID
     autonomous-system 65432; #Full mesh IBGP
    forwarding-table {
         export lbpp; #Routes exported to forwarding table
    }
}
protocols {
    rsvp {
        interface fe-0/1/0.0;
        interface fe-0/1/2.0;
        interface fxp0.0 {
            disable;
        }
    }
    mpls {
         label-switched-path lsp 1 { #First LSP
              to 192.168.0.1;  # Destination of the LSP
              install 10.0.90.14/32 active; # The prefix is installed in the
            primary via-r4;               # inet.0 routing table
        }
         label-switched-path lsp2 {
            to 192.168.0.1;
            install 10.0.90.14/32 active;
            primary via-r2;
        }
         label-switched-path lsp3 {
            to 192.168.0.1;
            install 10.0.90.14/32 active;
            primary via-r2;
        }
         label-switched-path lsp4 {
            to 192.168.0.1;
            install 10.0.90.14/32 active;
            primary via-r4;
```

```
            }
             path via-r2 { #Primary path to spread traffic across interfaces
                10.0.29.2 loose;
            }
             path via-r4  {
                10.0.24.2 loose;
            }
            interface fe-0/1/0.0;
            interface fe-0/1/2.0;
            interface fxp0.0 {
                disable;
            }
        }
        bgp {
             export send-statics; #Allows advertising of a new route
            group internal {
                type internal;
                local-address 192.168.1.1;
                neighbor 192.168.2.1;
                neighbor 192.168.4.1;
                neighbor 192.168.9.1;
                neighbor 192.168.6.1;
                neighbor 192.168.0.1;
            }
        }
         ospf  { #IGP enabled
            traffic-engineering;
            area 0.0.0.0 {
                interface fe-0/1/0.0;
                interface fe-0/1/2.0;
                 interface  lo0.0 {
                     passive; #Ensures protocols do not run over this interface
                }
            }
        }
    }
    policy-options {  #Load balancing policy
        policy-statement lbpp {
            then {
                load-balance per-packet;
            }
        }
        policy-statement send-statics {  #Static route policy
            term statics {
                from {
                    route-filter 100.100.1.0/24 exact;
                }
                then accept;
            }
        }
    }
```

Sample Output 3 The following configuration output is for transit router **R2**.

```
user@R2> show configuration | no-more
[...Output truncated...]
interfaces {
    so-0/0/1 {
        unit 0 {
            family inet {
                address 10.0.24.1/30;
```



```
            }
             family mpls; #MPLS enabled on relevant interfaces
        }
    }
    so-0/0/2 {
        unit 0 {
            family inet {
                address 10.0.29.1/30;
            }
            family mpls;
        }
    }
    fe-0/1/0 {
        unit 0 {
            family inet {
                address 10.0.12.14/30;
            }
            family mpls;
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.70.144/21;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.2.1/32;
            }
        }
    }
}
routing-options {
    static {
        [...Output truncated...]
     router-id 192.168.2.1; #Manually configured RID
     autonomous-system 65432; #Full mesh IBGP
    }
}
protocols {
    rsvp {
        interface so-0/0/1.0;
        interface fe-0/1/0.0;
        interface so-0/0/2.0;
        interface fxp0.0 {
            disable;
        }
    }
    mpls {
        interface fe-0/1/0.0;
        interface so-0/0/1.0;
        interface so-0/0/2.0;
        interface fxp0.0 {
            disable;
        }
    }
    bgp {
        group internal {
```

```
            type internal;
            local-address 192.168.2.1;
            neighbor 192.168.1.1;
            neighbor 192.168.4.1;
            neighbor 192.168.9.1;
            neighbor 192.168.6.1;
            neighbor 192.168.0.1;
        }
    }
     ospf  { #IGP enabled
        traffic-engineering;
        area 0.0.0.0 {
            interface fe-0/1/0.0;
            interface so-0/0/1.0;
            interface so-0/0/2.0;
            interface  lo0.0  {
                 passive; #Ensures protocols do not run over this interface
            }
        }
    }
}
```

Sample Output 4 The following configuration output is for transit router **R4**.

```
user@R4> show configuration | no-more
[...Output truncated...]
interfaces {
    so-0/0/1 {
        unit 0 {
            family inet {
                address 10.0.24.2/30;
            }
             family mpls;  # MPLS enabled on relevant interfaces
        }
    }
    so-0/0/3 {
        unit 0 {
            family inet {
                address 10.0.49.1/30;
            }
            family mpls;
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.70.146/21;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.4.1/32;
            }
        }
    }
}
routing-options {
    static {
        [...Output truncated...]
```



```
     router-id 192.168.4.1; #Manually configured RID
     autonomous-system 65432; #Full mesh IBGP
}
protocols {
    rsvp {
        interface so-0/0/1.0;
        interface so-0/0/3.0;
        interface fxp0.0 {
            disable;
        }
    }
    mpls {
        interface so-0/0/1.0;
        interface so-0/0/3.0;
        interface fxp0.0 {
            disable;
        }
    }
    bgp {
        group internal {
            type internal;
            local-address 192.168.4.1;
            neighbor 192.168.1.1;
            neighbor 192.168.2.1;
            neighbor 192.168.9.1;
            neighbor 192.168.6.1;
            neighbor 192.168.0.1;
        }
    }
     ospf  { #IGP enabled
        traffic-engineering;
        area 0.0.0.0 {
            interface so-0/0/1.0;
            interface so-0/0/3.0;
            interface  lo0.0  {
                 passive; #Ensures protocols do not run over this interface
            }
        }
    }
}
```

Sample Output 5 The following configuration output is for transit router **R9**.

```
user@R9> show configuration | no-more
[...Output truncated...]
interfaces {
    so-0/0/2 {
        unit 0 {
            family inet {
                address 10.0.29.2/30;
            }
             family mpls; #MPLS enabled on relevant interfaces
        }
    }
    so-0/0/3 {
        unit 0 {
            family inet {
                address 10.0.49.2/30;
            }
            family mpls;
        }
```

```
    }
    fe-0/1/0 {
        unit 0 {
            family inet {
                address 10.0.90.13/30;
            }
            family mpls;
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.69.206/21;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.9.1/32;
            }
        }
    }
}
routing-options {
    static {
        [...Output truncated...]
     router-id 192.168.9.1; #Manually configured RID
     autonomous-system 65432; #Full mesh IBGP
}
protocols {
    rsvp {
        interface so-0/0/2.0;
        interface so-0/0/3.0;
        interface fe-0/1/0.0;
        interface fxp0.0 {
            disable;
        }
    }
    mpls {
        interface so-0/0/2.0;
        interface so-0/0/3.0;
        interface fe-0/1/0.0;
        interface fxp0.0 {
            disable;
        }
    }
    bgp {
        group internal {
            type internal;
            local-address 192.168.9.1;
            neighbor 192.168.1.1;
            neighbor 192.168.2.1;
            neighbor 192.168.4.1;
            neighbor 192.168.0.1;
            neighbor 192.168.6.1;
        }
    }
     ospf  { #IGP enabled
        traffic-engineering;
        area 0.0.0.0 {
```



```
            interface so-0/0/2.0;
            interface so-0/0/3.0;
            interface fe-0/1/0.0;
            interface  lo0.0 {
                 passive; #Ensures protocols do not run over this interface
            }
        }
    }
}
```

Sample Output 6 The following configuration output is for egress router **R0**.

```
user@R0> show configuration | no-more
[...Output truncated...]
interfaces {
    fe-0/1/0 {
        unit 0 {
            family inet {
                address 10.0.90.14/30;
            }
             family mpls; #MPLS enabled on relevant interfaces
        }
    }
    fe-1/3/0 {
        unit 0 {
            family inet {
                address 10.10.11.1/24;
            }
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.69.207/21;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.0.1/32;
            }
        }
    }
}
routing-options {
    static {
        [...Output truncated...]
         route 100.100.10.0/24 reject; #Static route for send-statics policy
    }
     router-id 192.168.0.1; #Manually configured RID
     autonomous-system 65432; #Full mesh IBGP
}
protocols {
    rsvp {
        interface fe-0/1/0.0;
        interface fe-1/3/0.0;
        interface fxp0.0 {
            disable;
        }
    }
    mpls {
```

```
            label-switched-path r0-r6 {
                to 192.168.6.1;
            }
            interface fe-0/1/0.0;
            interface fe-1/3/0.0;
            interface fxp0.0 {
                disable;
            }
        }
        bgp {
            group internal {
                type internal;
                local-address 192.168.0.1;
                 export send-statics; #Allows advertising of a new route
                neighbor 192.168.9.1;
                neighbor 192.168.6.1;
                neighbor 192.168.1.1;
                neighbor 192.168.2.1;
                neighbor 192.168.4.1;
            }
        }
         ospf  { #IGP enabled
            traffic-engineering;
            area 0.0.0.0 {
                interface fe-0/1/0.0;
                interface fe-1/3/0.0;
                interface  lo0.0 {
                     passive; #Ensures protocols do not run over this interface
                }
            }
        }
    }
    policy-options {
         policy-statement send-statics {
            term statics {
                from {
                    route-filter 100.100.10.0/24 exact;
                }
                then accept;
            }
        }
    }
```

Meaning Sample Outputs 1 through 6 show the base interfaces, routing options, protocols, and policy options configurations for all six routers in the example network illustrated in Figure 10 on page 72.

All routers in the network have MPLS, RSVP, and BGP enabled. OSPF is configured as the IGP, and relevant interfaces have basic IP information and MPLS support.

In addition, all routers have the router ID (RID) configured manually at the **[edit routing-options]** hierarchy level to avoid duplicate RID problems. The **passive** statement is included in the OSPF configuration to ensure that protocols are not run over the loopback (**lo0**) interface and that the loopback (**lo0**) interface is advertised correctly throughout the network.

Sample Outputs 1, 3, 4, and 5 for **R6**, **R2**, **R4**, and **R9** show the base configuration for transit label-switched routers. The base configuration includes all interfaces enabled for

MPLS, the RID manually configured, and the relevant protocols (RSVP, MPLS, BGP, and OSPF).

Sample Output 2 from ingress router **R1** shows the base configuration plus four LSPs (**lsp1** through **lsp4**) configured to **R0**. The four LSPs are configured with different primary paths that specify a loose hop through **R4** for **lsp1** and **lsp4**, and through **R2** for **lsp2** and **lsp3**.

To create traffic, **R1** has a static route (**100.100.1.0/24**) configured at the **[edit routing-options static route]** hierarchy level. The prefix is included in the send-statics policy at the **[edit policy-options send statics]** hierarchy level so the routes can become BGP routes.

In addition, on the ingress router **R1**, load balancing is configured using the **per-packet** option, and the policy is exported at the **[edit routing-options forwarding-table]** hierarchy level.

Sample Output 6 from egress router **R0** shows one LSP (**r0-r6**) to **R6** used to create bidirectional traffic. OSPF requires bidirectional LSP reachability before it will advertise the LSP into the IGP. Although the LSP is advertised into the IGP, no hello messages or routing updates occur over the LSP—only user traffic is sent over the LSP. The router uses its local copy of the IGP database to verify bidirectional reachability.

In addition, **R0** has a static route (**100.100.10.0/24**) configured at the **[edit routing-options static route]** hierarchy level. The prefix is included in the send-statics policy at the **[edit policy-options send statics]** hierarchy level so the routes can become BGP routes.

Using Hash-Key Load Balancing for LSP Traffic

Purpose If the outbound traffic across equal-cost next hops is not well balanced after you have load balancing configured, you can configure the hash key to provide additional information to further identify traffic flows and balance traffic more evenly. The hash key is configured at the **[edit forwarding-options]** hierarchy level on ingress and transit routers, depending on your network configuration.

Within the hash key, you can configure the IPv4 address family (INET) or the MPLS protocol family. Typically for the best results, you configure the INET on the ingress router and the MPLS protocol family on the transit router. If a router happens to be both an ingress and transit router, you can configure both the INET and the MPLS protocol family. However, the INET will only be used when the router is acting as an ingress router and the MPLS protocol family will only be used when the router is acting as a transit router. You can configure only the INET on the ingress router, and check that the results are what is intended. Similarly, you can configure only MPLS labels on the transit router, and check the results.

In addition, the MPLS protocol family is most useful in configurations with aggregated interfaces, although it can be used on transit routers with regular (non-aggregated) interfaces.

To use the hash key to load-balance LSP traffic, follow these steps:

1. Configuring MPLS Labels and IP Payload to Load-Balance LSP Traffic on page 84
2. Configuring the IPv4 Address Family to Load-Balance LSP Traffic on page 86

Configuring MPLS Labels and IP Payload to Load-Balance LSP Traffic

If the outbound traffic across equal-cost next hops is not well balanced after you have load balancing configured, you can use MPLS labels and IP payload to provide additional information to further identify traffic flows and balance traffic more evenly, particularly between aggregated interfaces. With an aggregated interface, when you configure load balancing using the **per-packet** statement, the Junos OS uses the first MPLS label in the hash algorithm to determine the next hop for the LSP. This behavior can result in an uneven distribution of traffic for aggregated interfaces.

You configure MPLS labels on a transit router because the transit router uses MPLS labels to forward traffic. Configuring the first two MPLS labels and the IP header is useful in the following circumstances:

- If there are many circuit cross-connect (CCC) MPLS LSPs using remote interface switching over an aggregated interface, configuring the first label can load-balance traffic between the component links of an aggregated interface. However, in other circumstances, such as an RSVP LSP, there is no benefit in configuring the first MPLS label by itself because load balancing using the **per-packet** statement uses the first label by default.
- If there are many CCC MPLS LSPs using remote interface switching over an aggregated interface with Martini Layer 1 VPN or Layer 2 VPN traffic, configuring the second MPLS label can load-balance traffic between component links of an aggregated interface.
- If there are CCC MPLS LSPs using remote interface switching over an aggregated interface with Layer 3 VPN traffic, Layer 2 VPN, or Martini Layer 2 VPN translational cross-connect (TCC) traffic, configuring the first and second MPLS labels and IP payload can balance traffic between component links of an aggregated interface.

Essentially, load balancing is similar across platforms even though there are slight differences between platforms. On M-series platforms, only Label 1 and the IP payload are used in the hash-key algorithm. On T-series platforms and the M320, all three labels (Label 1, Label 2, and IP payload) are used in the hash-key algorithm. However, there is no harm done if you configure all three labels on an M-series router. The router simply ignores Label 2.

Before you use MPLS labels and IP payload to load-balance traffic, you must have the **load-balance per-packet** statement configured at the **[edit policy-options]** hierarchy level and that policy applied as an export policy at the **[edit forwarding-options]** hierarchy level. For more information about configuring load balancing, see “Configuring and Verifying Load Balancing” on page 67.

Action To configure the hash key to load-balance LSP traffic, follow these steps:

1. Ensure that you have load balancing configured; see “Configuring and Verifying Load Balancing” on page 67.
2. In configuration mode, go to the following hierarchy level:

   ```
   [edit]
   user@host# edit forwarding-options hash-key
   ```

3. Depending on your network configuration, include a combination of MPLS labels to include in the configuration:

   ```
   [edit forwarding-options hash-key]
   user@host# set family mpls label-1
   user@host# set family mpls label-2
   user@host# set family mpls payload ip
   ```

   

   NOTE: The configuration of all three statements together can be used on T-series and M320 routing platforms only. If you configure all three statements on an M-series router, only label-1 and the IP payload are used in the hash key.

4. Verify and commit the configuration:

   ```
   user@host# show
   user@host# commit
   ```

Sample Output

```
user@R2> edit
Entering configuration mode

[edit]
user@R2# edit forwarding-options hash-key

[edit forwarding-options hash-key]
user@R2# set family mpls label-1

[edit forwarding-options hash-key]
user@R2# set family mpls label-2

[edit forwarding-options hash-key]
user@R2# set family mpls payload ip

[edit forwarding-options hash-key]
user@R2# show
family mpls {
    label-1;
    label-2;
    payload {
        ip;
    }
}

[edit forwarding-options hash-key]
user@R2# commit
commit complete
```

Meaning The sample output shows the configuration of all three MPLS labels and verification that the configuration is correct.

Configuring the IPv4 Address Family to Load-Balance LSP Traffic

Purpose If the outbound traffic across equal-cost next hops is not well balanced after you have load balancing configured, you can use the IPv4 address family (INET) to provide additional information to identify traffic flows and balance traffic more evenly. You configure the INET or port data on an ingress router. Configuring port data is useful if you are using TCP or UDP. However, it may not be useful to include port data when you are using protocols that are not associated with a Layer 4 port, for example, Layer 2 VPNs, generic routing encapsulation (GRE) tunneling, or Internet Control Message Protocol (ICMP).

To configure port data, you include the **layer-3** or **layer-4** options under the **family-inet** statement at the **[edit forwarding-options hash-key]** hierarchy level. When you include the **layer-4** option, you must also include the **layer-3** statement. If you omit the **layer-3** statement, the management process removes the **hash-key** statement from the configuration and the router works as if you specified **layer-3**.

If you specify only the **layer-3** statement in the configuration, the router uses the incoming interface index as well as the following Layer 3 information in the packet header to load-balance:

- Source IP address
- Destination IP address
- Protocol

If you include both the **layer-3** and **layer-4** statements, the router uses the following Layer 3 and Layer 4 information to load-balance:

- Source IP address
- Destination IP address
- Protocol
- Source port number
- Destination port number
- Incoming interface index

The router recognizes packets in which all of these Layer 3 and Layer 4 parameters are identical, and ensures that these packets are sent out through the same interface. This prevents problems that might otherwise occur with packets arriving at their destination out of their original sequence.

Before you use port data to send packets through the same interface, you must have the **load-balance per-packet** statement configured at the **[edit policy-options]** hierarchy level and that policy applied as an export policy at the **[edit forwarding-options]** hierarchy

level. For more information about configuring load balancing, see “Configuring and Verifying Load Balancing” on page 67.

Action To configure the hash key with port data, follow these steps:

1. Ensure that you have load balancing configured, see “Configuring and Verifying Load Balancing” on page 67.
2. In configuration mode, go to the following hierarchy level:

   ```
   [edit]
   user@host# edit forwarding-options hash-key
   ```

3. Include Layer 3 (IP) data in the hash key:

   ```
   [edit forwarding-options hash-key]
   user@host# set family inet layer-3
   ```

4. Include Layer 4 TCP or UDP data in the hash key:

   ```
   [edit forwarding-options hash-key]
   user@host# set family inet layer-4
   ```

5. Verify and commit the configuration:

   ```
   user@host# show
   user@host# commit
   ```

Sample Output

```
user@R6> edit
Entering configuration mode

[edit]
user@R1# edit forwarding-options hash-key

[edit forwarding-options hash-key]
user@R1# set family inet layer-3

[edit forwarding-options hash-key]
user@R1# set family inet layer-4

[edit forwarding-options hash-key]
user@R1# show
family inet {
    layer-3;
    layer-4;
}

[edit forwarding-options hash-key]
user@R1# commit
commit complete
```

Meaning The sample output shows both the **layer-3** and **layer-4** options included in the hash key. Including both options provides additional information to identify traffic flows and balance traffic more evenly.

Hash Key Network Examples

Depending on the fields used in the hash-key algorithm and your network requirements, you can fine-tune the way traffic is load-balanced across your network. For example, if your network supports a large number of uses on routers running Network Address Translation (NAT) or Port Address Translation (PAT), the flows will be similar at Layer 3, so adding both Layer 3 and Layer 4 to the hash key can provide better load balancing. However, if a core router in your network is supporting tens of thousands of unrelated flows that vary significantly in source or destination addresses and incoming interfaces, including only Layer 3 in the hash key would probably result in a good distribution of traffic. With some exceptions, the more fields included in the hash-key algorithm, the greater the chance that traffic is unique, resulting in an optimal balance of traffic.

The following network examples illustrate various ways of using the hash key to load-balance traffic in different types of networks:

- Example: Load-Balancing a Network with Aggregated Interfaces on page 88
- Example: Load-Balancing a Network Using INET in the Hash Key on page 100

Example: Load-Balancing a Network with Aggregated Interfaces

Purpose With an aggregated interface, when you configure load balancing using the **per-packet** statement, the Junos OS uses the first MPLS label in the hash algorithm to determine the next hop for the LSP. This behavior can result in an uneven distribution of traffic for aggregated interfaces.

This example describes load balancing using an LDP tunneled over RSVP on a network comprised of M-series and T-series routers with aggregated interfaces. Figure 11 on page 88 illustrates the network used in this topic.

Figure 11: Aggregated Interfaces Network Topology

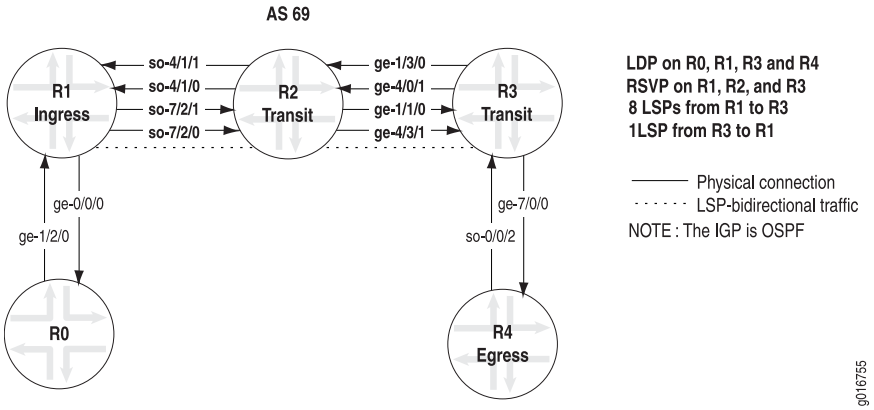

The network topology in Figure 11 on page 88 illustrates a router-only network with aggregated SONET and Ethernet interfaces that consists of the following components:

- BGP configured on PE routers **R0** and **R4**
- LDP running on **R0**, **R1**, **R3**, and **R4**
- RSVP running on **R1**, **R2**, and **R3**
- LSPs set up from **R1** to **R3**, and **R3** to **R1**
- Aggregated interfaces on **R1**, **R2**, and **R3**
- The hash key configured on transit router **R2**
- Load balancing configured on **R1**

With the hash key configuration on **R2**, outbound traffic for the aggregated interface varies in terms of Label 1, Label 2, or IP payload. This variance in traffic should result in the equal distribution of traffic across different physical links of the aggregated interface.

The following information is included in this example:

- Verifying the Operation of Load Balancing with Aggregated Interfaces on page 89
- For the configuration output of all routers in this network, see “Router Configurations for the Aggregated Interfaces Network” on page 93

Verifying the Operation of Load Balancing with Aggregated Interfaces

Purpose On an M-series or T-series platform, when you configure only the first two labels and the labels vary between traffic flows, traffic may be distributed. However, if there is not much variation between the two labels, traffic may not be distributed equally across aggregated interfaces.

The following output illustrates two situations. The first example output shows a situation in which traffic is not balanced across interfaces because there is not enough variation between the two configured labels. The second example output shows a situation in which traffic is balanced; all three MPLS labels are configured and the third label has enough variation to yield good load-balancing results.

Action To verify the operation of the hash key, enter the following Junos OS CLI operational mode commands:

```
user@R2> show configuration forwarding-options
user@R2> show interfaces statistics interface-name detail
user@R2> show mpls lsp statistics
```

Sample Output

```
user@R2> show configuration forwarding-options
hash-key {
    family mpls {
        label-1;
        label-2;
    } # The IP payload option is missing.
}
```

```
user@R2> show interfaces statistics ae0 detail
Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 289, Generation: 129
  Link-level type: Ethernet, MTU: 1514, Speed: 2000mbps, Loopback: Disabled,
```

```
Source filtering: Disabled,
  Flow control: Disabled, Minimum links needed: 1
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: 00:90:69:0f:07:f0, Hardware address: 00:90:69:0f:07:f0
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :               102162                 560 bps
   Output bytes  :            166100728             30744472 bps
   Input  packets:                 1259                    0 pps
   Output packets:              2442317                56515 pps
  Label-switched interface (LSI) traffic statistics:
   Input  bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards:
 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
 0
  Ingress queues: 8 supported, 8 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 best-effort                    0                    0                    0

    1 expedited-fo                   0                    0                    0

    2 assured-forw                   0                    0                    0

    3 network-cont                   0                    0                    0

  Egress queues: 8 supported, 8 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 best-effort              2440822              2440822                    0

    1 expedited-fo                   0                    0                    0

    2 assured-forw                   0                    0                    0

    3 network-cont                1225                 1225                    0


  Logical interface ae0.0 (Index 66) (SNMP ifIndex 290) (Generation 131)
    Flags: SNMP-Traps Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :          1259          0        102162          560
        Output:       2441888      56515     166056858     30744472
    Link:
       ge-1/1/0.0
        Input :          1259          0        102162          560
        Output:       1488498       56515      101217864      30744272
       ge-4/3/1.0
        Input :             0          0             0            0
        Output:        953235           0      64822700          200
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      ge-1/1/0.0                0           0            0            0
      ge-4/3/1.0                0           0            0            0
    Protocol inet, MTU: 1500, Generation: 128, Route table: 0
```

```
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.35.200.4/30, Local: 10.35.200.5, Broadcast: 10.35.200.7,

Generation: 147
    Protocol iso, MTU: 1497, Generation: 128, Route table: 0
      Flags: Is-Primary
    Protocol mpls, MTU: 1488, Generation: 128, Route table: 0
      Flags: Is-Primary
```

Meaning This sample shows that the hash key at the **[edit forwarding-options]** hierarchy level does not include the IP payload, indicating that the distribution of traffic is varying only by IP payload. The output for the **show interfaces statistics** command shows that traffic is not evenly distributed between links **ge-1/1/0** and **ge-4/3/1** of the aggregated interface **ae0**, probably due to the missing IP payload label.

Sample Output

```
user@R2> show configuration forwarding-options
haey {
    family mpls {
        label-1;
        label-2;
         payload {
            ip;
        }
    }
}
```

```
user@R2> show interfaces statistics ae0 detail
Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 185, SNMP ifIndex: 289, Generation: 186
  Link-level type: Ethernet, MTU: 1514, Speed: 2000mbps, Loopback: Disabled,
Source filtering: Disabled,
  Flow control: Disabled, Minimum links needed: 1
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: 00:90:69:0f:07:f0, Hardware address: 00:90:69:0f:07:f0
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :              1000775                    0 bps
   Output bytes  :             79662734             30743104 bps
   Input  packets:                13273                    0 pps
   Output packets:              1168916                56512 pps
  Label-switched interface (LSI) traffic statistics:
   Input  bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards:
 0, Resource errors: 0
  Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors: 0

  Ingress queues: 8 supported, 8 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 best-effort                    0                    0                    0

    1 expedited-fo                   0                    0                    0

    2 assured-forw                   0                    0                    0
```

```
    3 network-cont                  0                    0                    0

  Egress queues: 8 supported, 8 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 best-effort             1186217              1186219                    0

    1 expedited-fo                  0                    0                    0

    2 assured-forw                  0                    0                    0

    3 network-cont              13057                13057                    0


  Logical interface ae0.0 (Index 71) (SNMP ifIndex 292) (Generation 137)
    Flags: SNMP-Traps Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :         13273          0       1000775            0
        Output:       1168916      56512      79662734     30743104
    Link:
      ge-1/1/0.0 #Packets are evenly distributed across aggregated interfaces
        Input :         13273          0       1000775            0
        Output:        610927      28256      41716580     15371728
      ge-4/3/1.0
        Input :             0          0             0            0
        Output:        557989      28256      37946154     15371376
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      ge-1/1/0.0                0           0            0            0
      ge-4/3/1.0                0           0            0            0
    Protocol inet, MTU: 1500, Generation: 146, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.35.200.4/30, Local: 10.35.200.5, Broadcast: 10.35.200.7,
 Generation: 151
    Protocol iso, MTU: 1497, Generation: 147, Route table: 0
      Flags: Is-Primary
    Protocol mpls, MTU: 1488, Generation: 148, Route table: 0
      Flags: Is-Primary
```

Meaning This sample output shows the configuration of the hash key and the interface statistics for the aggregated interface **ae0**. The hash-key configuration specifies the labels used for outbound traffic on different physical links of the aggregated interface. In this case, two labels and IP payload are included in the configuration. The sample output for the **show interface statistics** command shows the outgoing traffic rate, which is evenly distributed between links **ge-1/1/0** and **ge-4/3/1** of aggregated interface **ae0**. However, an even distribution may not always be the case because it depends on a lot of factors, which can be defined at the **[edit forwarding-options]** hierarchy level.

Sample Output

```
user@R2> show mpls lsp statistics
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 9 sessions
```



```
To              From            State     Packets          Bytes LSPname
10.255.70.186   10.255.71.199   Up           1430         101943 to_R1_from_R3
10.255.71.199   10.255.70.186   Up          81745        5558550 to_R3_from_R1
10.255.71.199   10.255.70.186   Up          81748        5558768 to_R3_from_R1_1
10.255.71.199   10.255.70.186   Up          81760        5559492 to_R3_from_R1_2
10.255.71.199   10.255.70.186   Up         153259       10421488 to_R3_from_R1_3
10.255.71.199   10.255.70.186   Up         163509       11118573 to_R3_from_R1_4
10.255.71.199   10.255.70.186   Up         163453       11114666 to_R3_from_R1_5
10.255.71.199   10.255.70.186   Up         163450       11114554 to_R3_from_R1_6
10.255.71.199   10.255.70.186   Up         132785        9029356 to_R3_from_R1_7
Total 9 displayed, Up 9, Down 0
```

Meaning This sample output shows that there are nine LSPs transiting **R2**. All are up and passing varying amounts of traffic.

Router Configurations for the Aggregated Interfaces Network

Purpose The configurations in this topic are for the five routers in the example network illustrated in Figure 11 on page 88.

Action To display the configuration of a router, use the following Junos OS CLI operational mode command:

```
user@host> show configuration | no-more
```

Sample Output 1

```
user@R0> show configuration | no-more
[...Output truncated...]
interfaces {
    ge-1/2/0 {
        unit 0 {
            family inet {
                address 10.35.1.1/30;
            }
            family iso;
            family mpls;
        }
    }
}
routing-options {
    autonomous-system 69;
}
protocols {
    }
    mpls {
        traffic-engineering bgp-igp-both-ribs;
        interface all;
    }
    bgp {
        group int {
            type internal;
            local-address 10.255.71.197;
            family inet {
                any;
            }
            family inet-vpn {
                any;
            }
            neighbor 10.255.70.79;
```

```
        }
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface lo0.0;
                passive
            interface all;
            interface fxp0.0 {
                disable;
            }
        }
    }
    ldp {
        interface all;
        interface lo0.0
            passive
       }
    }
}
policy-options {
    policy-statement all_routes {
        then accept;
    }
}
```

Sample Output 2

```
user@R1> show configuration | no-more
[...Output truncated...]
chassis {
    redundancy {
        graceful-switchover {
            enable;
        }
    }
    aggregated-devices {
        sonet {
            device-count 1;
        }
    }
}
interfaces {
    ge-0/0/0 {
        unit 0 {
            family inet {
                address 10.35.1.2/30;
            }
            family iso;
            family mpls;
        }
    }
    so-7/2/0 {
        sonet-options {
            aggregate as0;
        }
    }
    so-7/2/1 {
        sonet-options {
            aggregate as0;
        }
    }
    as0 {
```

```
        aggregated-sonet-options {
            minimum-links 1;
        }
        unit 0 {
            family inet {
                address 10.35.200.1/30;
            }
            family iso;
            family mpls;
        }
    }
}
routing-options {
    autonomous-system 69;
    forwarding-table {
        export pplb;
    }
}
protocols {
    rsvp {
        interface all;
    }
    mpls {
        traffic-engineering bgp-igp-both-ribs;
        label-switched-path to_R3_from_R1 {
            to 10.255.71.199;
            ldp-tunneling;
        }
        label-switched-path to_R3_from_R1_1 {
            to 10.255.71.199;
            ldp-tunneling;
        }
        label-switched-path to_R3_from_R1_2 {
            to 10.255.71.199;
            ldp-tunneling;
        }
        label-switched-path to_R3_from_R1_3 {
            to 10.255.71.199;
            ldp-tunneling;
        }
        label-switched-path to_R3_from_R1_4 {
            to 10.255.71.199;
            ldp-tunneling;
        }
        label-switched-path to_R3_from_R1_5 {
            to 10.255.71.199;
            ldp-tunneling;
        }
        label-switched-path to_R3_from_R1_6 {
            to 10.255.71.199;
            ldp-tunneling;
        }
        label-switched-path to_R3_from_R1_7 {
            to 10.255.71.199;
            ldp-tunneling;
        }
        interface all;
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
```

```
            interface lo0.0;
                passive
            interface all;
            interface fxp0.0 {
                disable;
            }
        }
    }
    ldp {
        interface ge-0/0/0.0;
        interface lo0.0;
    }
}
policy-options {
    policy-statement pplb {
        then {
            load-balance per-packet;
        }
    }
}
```

Sample Output 3

```
user@R2> show configuration | no-more
[...Output truncated...]
chassis {
    redundancy {
        graceful-switchover {
            enable;
        }
    }
    aggregated-devices {
        ethernet {
            device-count 1;
        }
        sonet {
            device-count 3;
        }
    }
}
interfaces {
    ge-1/1/0 {
        gigether-options {
            802.3ad {
                ae0;
            }
        }
    }
    so-4/1/0 {
        sonet-options {
            aggregate as0;
        }
    }
    so-4/1/1 {
        sonet-options {
            aggregate as0;
        }
    }
    ge-4/3/1 {
        gigether-options {
            802.3ad {
                ae0;
            }
```

```
        }
    }
    ae0 {
        aggregated-ether-options {
            minimum-links 1;
        }
        unit 0 {
            family inet {
                address 10.35.200.5/30;
            }
            family iso;
            family mpls;
        }
    }
    as0 {
        aggregated-sonet-options {
            minimum-links 1;
        }
        unit 0 {
            family inet {
                address 10.35.200.2/30;
            }
            family iso;
            family mpls;
        }
    }
}
forwarding-options {
    hash-key {
        family mpls {
            label-1;
            label-2;
            label-3;
        }
    }
}
routing-options {
    autonomous-system 69;
    forwarding-table {
        export pplb;
    }
}
protocols {
    rsvp {
        interface all;
    }
    mpls {
        traffic-engineering bgp-igp-both-ribs;
        interface all;
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface lo0.0;
            interface all;
            interface fxp0.0 {
                disable;
            }
        }
    }
}
```

```
policy-options {
    policy-statement pplb {
        then {
            load-balance per-packet;
        }
    }
}
```

Sample Output 4

```
user@R3> show configuration | no-more
[...Output truncated...]
chassis {
    redundancy {
        graceful-switchover {
            enable;
        }
    }
    aggregated-devices {
        ethernet {
            device-count 1;
        }
        sonet {
            device-count 2;
        }
    }
}
interfaces {
    ge-1/3/0 {
        gigether-options {
            802.3ad {
                ae0;
            }
        }
    }
    ge-4/0/1 {
        gigether-options {
            802.3ad {
                ae0;
            }
        }
    }
    ge-7/0/0 {
        unit 0 {
            family inet {
                address 10.35.1.53/30;
            }
            family iso;
            family mpls;
        }
    }
    ae0 {
        aggregated-ether-options {
            minimum-links 1;
        }
        unit 0 {
            family inet {
                address 10.35.200.6/30;
            }
            family iso;
            family mpls;
        }
    }
```

```
}
routing-options {
    autonomous-system 69;
}
protocols {
    rsvp {
        interface all;
    }
    mpls {
        traffic-engineering bgp-igp-both-ribs;
        label-switched-path to_R1_from_R3 {
            to 10.255.70.186;
            ldp-tunneling;
        }
        interface all;
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface lo0.0;
            interface all;
            interface fxp0.0 {
                disable;
            }
        }
    }
    ldp {
        interface ge-7/0/0.0;
        interface lo0.0;
    }
}
```

Sample Output 5

```
user@R4> show configuration | no-more
[...Output truncated...]
interfaces {
    ge-0/3/0 {
        unit 0 {
            family inet {
                address 10.35.1.54/30;
            }
            family iso;
            family mpls;
        }
    }
}
routing-options {
    autonomous-system 69;
}
protocols {
    mpls {
        traffic-engineering bgp-igp-both-ribs;
        interface all;
    }
    bgp {
        group int {
            type internal;
            local-address 10.255.70.79;
            family inet {
                any;
            }
            family inet-vpn {
```

```
                any;
            }
        }
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface lo0.0;
            interface all;
            interface fxp0.0 {
                disable;
            }
        }
    }
    ldp {
        interface all;
    }
}
```

Meaning Sample Outputs 1 through 5 show the configuration of all routers in the example network shown in Figure 11 on page 88.

Example: Load-Balancing a Network Using INET in the Hash Key

Purpose The IPv4 address family (INET) provides additional information to identify traffic flows and balance traffic more evenly. You configure the INET or port data on an ingress router. Configuring port data is useful if you are using TCP or UDP. However, it may not be useful to include port data when you are using protocols that are not associated with a Layer 4 port, for example, Layer 2 VPNs, GRE tunneling, or ICMP.

The following network example shows the process for verifying the operation of the hash key configuration of port data as described in "Configuring the IPv4 Address Family to Load-Balance LSP Traffic" on page 86.

Figure 12: INET Network Topology

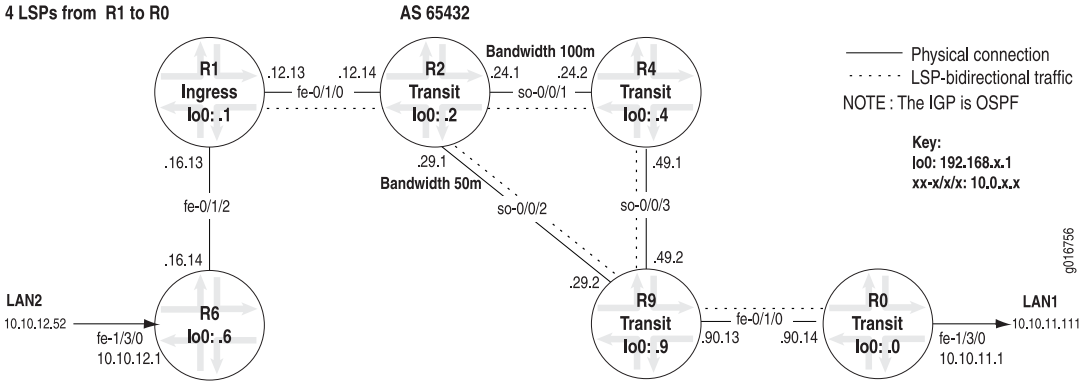

The network topology in Figure 11 on page 88 illustrates a router-only network with SONET and Ethernet interfaces that consists of the following components:

- A full-mesh interior BGP (IBGP) topology, using AS 65432
- MPLS and RSVP enabled on all routers

- A send-statics policy on routers **R1** and **R0** that allows a new route to be advertised into the network
- Four unidirectional LSPs between **R1** and **R0**, and one reverse direction LSP between **R0** and **R1**, which allows for bidirectional traffic
- Load balancing configured on the ingress router **R1**
- The hash key using port data configured on **R1**
- Bandwidth configured on the SONET interfaces on **R2**

In addition, the example network uses Open Shortest Path First (OSPF) as the interior gateway protocol (IGP) with OSPF area **0.0.0.0**. An IGP is required for the Constrained Shortest Path First (CSPF) LSP, which is the default for the Junos OS. Also, the example network uses a policy to create BGP traffic.

The following information is included in this example:

- Verifying the Operation of INET Load Balancing on page 101
- Router Configurations for the INET Load-Balanced Network on page 103

Verifying the Operation of INET Load Balancing

Purpose Verifying the operation of the hash key configuration of port data on the ingress and transit routers in your network.

Action On the ingress router, to verify the operation of the hash key, enter the following Junos OS CLI operational mode commands:

```
user@host> show configuration
user@host# show route forwarding-table destination destination
```

On the transit router, to verify the operation of the hash key, enter the following Junos OS CLI operational mode commands:

```
user@host> show route
user@host> monitor interface traffic
user@host> show mpls lsp statistics
```

Sample Output The following sample output is for ingress router R1:

```
user@R1> show configuration forwarding-options
hash-key {
     family inet { #Port data configuration
        layer-3;
        layer-4;
    }
}

user@R1> show configuration routing-options
static {
[...Output truncated...]
autonomous-system 65432;
forwarding-table {
     export lbpp; #Load balancing policy applied
```

```
}

user@R1> show configuration policy-options
policy-statement lbpp { #Load balancing policy defined
    then {
        load-balance per-packet;
    }
}
[...Output truncated...]
```

Meaning The sample output from ingress router **R1** for the three **show configuration** commands (**forwarding-options**, **routing-options**, and **policy-options**) shows that load balancing is correctly configured for the INET hash key and the load-balancing policy (**lbpp**).

Sample Output The following sample output is for ingress router R1:

```
user@R1> show route forwarding-table destination 10.0.90.14
Routing table: inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
10.0.90.14/32      user     0                     ulst 262144  2
                              10.0.12.14        Push 100688       fe-0/1/0.0
                              10.0.12.14        Push 100656       fe-0/1/0.0
                              10.0.12.14        Push 100672       fe-0/1/0.0
                              10.0.12.14        Push 100704       fe-0/1/0.0
```

Meaning The sample output from ingress router **R1** for the **show route forwarding-table destination** command shows unilist (**ulst)** in the **Type** field, indicating that load balancing is working. In this case, the **Type** field shows the operation performed on packets. The push operation adds a new label to the top of the packet before the packets are pushed out of interface **fe-0/1/0.0**.

Sample Output The following sample output is for transit router R2:

```
user@R2> show route 10.0.90.14
inet.0: 25 destinations, 27 routes (25 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.0.90.12/30      *[OSPF/10] 03:06:04, metric 3
                       via so-0/0/1.0
                    > via so-0/0/2.0
```

Meaning The sample output from transit router **R2** for the **show route** command shows two OSPF routes to the destination interface on egress router **R0**. Even though the route with the greater than sign (**>**) is the selected route, traffic will be balanced across both interfaces, as shown in the output for the following **show route forwarding-table** and **monitor traffic** commands.

Sample Output The following sample output is for transit router R2:

```
user@R2> show route forwarding-table destination 10.0.90.14
Routing table: inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
10.0.90.12/30      user     0                     ulst 262144  6
```



```
                                ucst     345     5 so-0/0/1.0
                                ucst     339     2 so-0/0/2.0
```

Meaning The sample output from transit router **R2** for the **show route forwarding-table destination** command shows unilist (**ulst)** in the **Type** field, indicating that load balancing is working. A packet sent to this next hop (**R2**) goes to any next hop in the unicast (**ucst**) list, **so-0/0/1.0** and **so-0/0/2.0**.

Sample Output The following sample output is for transit router R2:

```
user@R2> monitor interface traffic

R2                                 Seconds: 123                 Time: 21:28:29

Interface    Link  Input packets        (pps)     Output packets        (pps)
 so-0/0/0      Up               0          (0)                  0          (0)
 so-0/0/1      Up              95          (0)              50012           (1)
 so-0/0/2      Up          100132         (19)              50217           (0)
 so-0/0/3      Up               0          (0)                  0          (0)
 fe-0/1/0      Up          100127         (17)             100128          (1)
 fe-0/1/1      Up               0          (0)                  0          (0)
 fe-0/1/2      Up               0          (0)                  0          (0)
 fe-0/1/3      Up               0          (0)                  0          (0)
 [...Output truncated...]
```

Meaning The sample output from transit router **R2** for the **monitor interface traffic** command shows that traffic is balanced across interfaces **so-0/0/1.0** and **so-0/0/2.0.**

Sample Output The following sample output is for transit router R2:

```
user@R2> show mpls lsp statistics
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 5 sessions
To              From            State     Packets            Bytes LSPname
192.168.0.1     192.168.1.1     Up          24874          2188912 lsp1
192.168.0.1     192.168.1.1     Up          24471          2153448 lsp2
192.168.0.1     192.168.1.1     Up          25613          2253944 lsp3
192.168.0.1     192.168.1.1     Up          25042          2203696 lsp4
192.168.1.1     192.168.0.1     Up              0                0 r0-r1
Total 5 displayed, Up 5, Down 0
```

Meaning The sample output from transit router R2 for the show mpls lsp statistics command shows that traffic is balanced across the four LSPs (lsp1, lsp2, lsp3, and lsp4) transiting **R2**.

Router Configurations for the INET Load-Balanced Network

Purpose The configurations in this topic are for the six routers in the example network illustrated in Figure 12 on page 100.

Action To display a router configuration, use the following Junos OS CLI operational mode command:

```
user@host> show configuration | no-more
```

Sample Output 1 The following sample output is for edge router **R6**:

```
user@R6> show configuration | no-more
interfaces {
     fe-0/1/2  { #Interface connected to R1
        unit 0 {
            family inet {
                address 10.0.16.14/30;
            }
             family mpls; #MPLS enabled on relevant interfaces
        }
    }
    fe-1/3/0 {
        unit 0 {
            family inet {
                address 10.10.12.1/24;
            }
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.70.148/21;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.6.1/32;
            }
        }
    }
}
routing-options {
    static {
        [...Output truncated...]
    }
     router-id 192.168.6.1; #Manually configured RID
     autonomous-system 65432; #Full mesh IBGP
}
protocols {
    rsvp {
        interface fe-0/1/2.0;
        interface fxp0.0 {
            disable;
        }
    }
    mpls {
        interface fe-0/1/2.0;
        interface fxp0.0 {
            disable;
        }
    }
    bgp {
```



```
        group internal {
            type internal;
            local-address 192.168.6.1;
            neighbor 192.168.1.1;
            neighbor 192.168.2.1;
            neighbor 192.168.4.1;
            neighbor 192.168.9.1;
            neighbor 192.168.0.1;
        }
    }
     ospf  { #IGP enabled
        traffic-engineering;
        area 0.0.0.0 {
            interface fe-0/1/2.0;
            interface fe-1/3/0.0;
            interface  lo0.0  {
                 passive; #Ensures protocols do not run over this interface
            }
        }
    }
}
```

Sample Output 2 The following sample output is for ingress router R1:

```
user@R1> show configuration | no-more
interfaces {
     fe-0/1/0 { #Connected to R2
        unit 0 {
            family inet {
                address 10.0.12.13/30;
            }
             family mpls; #MPLS enabled on relevant interfaces
        }
    }
     fe-0/1/2 { #Connected to R6
        unit 0 {
            family inet {
                address 10.0.16.13/30;
            }
            family mpls;
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.70.143/21;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.1.1/32;
            }
        }
    }
}
forwarding-options {
    hash-key {
         family inet  { INET/port data
             layer-3;
```

```
                layer-4;
            }
        }
    }
    routing-options {
        static {
            [...Output truncated...
            }
             route 100.100.1.0/24 reject; #Static route for send-statics policy
        }
         router-id 192.168.1.1; #Manually configured RID
         autonomous-system 65432; #Full mesh IBGP
        forwarding-table {
             export lbpp; #Routes exported to forwarding table
        }
    }
    protocols {
        rsvp {
            interface fe-0/1/0.0;
            interface fe-0/1/2.0;
            interface fxp0.0 {
                disable;
            }
        }
        mpls {
            label-switched-path lsp1  { #First LSP
                 to 192.168.0.1; # Destination of the LSP
                 install 10.0.90.14/32 active; # The prefix is installed in the
                primary via-r4;               # inet.0 routing table
            }
            label-switched-path lsp2 {
                to 192.168.0.1;
                install 10.0.90.14/32 active;
                primary via-r2;
            }
            label-switched-path lsp3 {
                to 192.168.0.1;
                install 10.0.90.14/32 active;
                primary via-r2;
            }
            label-switched-path lsp4 {
                to 192.168.0.1;
                install 10.0.90.14/32 active;
                primary via-r4;
            }
             path via-r2  { #Primary path to spread traffic across interfaces
                10.0.29.2 loose;
            }
             path via-r4 {
                10.0.24.2 loose;
            }
            interface fe-0/1/0.0;
            interface fe-0/1/2.0;
            interface fxp0.0 {
                disable;
            }
        }
        bgp {
             export send-statics; #Allows advertising of a new route
            group internal {
                type internal;
```



```
            local-address 192.168.1.1;
            neighbor 192.168.2.1;
            neighbor 192.168.4.1;
            neighbor 192.168.9.1;
            neighbor 192.168.6.1;
            neighbor 192.168.0.1;
        }
    }
     ospf  { #IGP enabled
        traffic-engineering;
        area 0.0.0.0 {
            interface fe-0/1/0.0;
            interface fe-0/1/2.0;
            interface lo0.0 {
                passive;
            }
        }
    }
}
policy-options {
     policy-statement lbpp { #Load balancing policy
          then {
               load-balance per-packet;
        }
    }
     policy-statement send-statics { #Static route policy
        term statics {
            from {
                route-filter 100.100.1.0/24 exact;
            }
            then accept;
        }
    }
}
```

Sample Output 3 The following sample output is for transit router R2:

```
user@R2> show configuration | no-more
interfaces {
     so-0/0/1  { #Connected to R4
        unit 0 {
             bandwidth 100m; #Bandwidth to ensure equal-cost paths
            family inet {
                address 10.0.24.1/30;
            }
             family mpls; #MPLS enabled on relevant interfaces
        }
    }
     so-0/0/2  { #Connected to R9
        unit 0 {
             bandwidth 50m; #Bandwidth to ensure equal-cost paths
            family inet {
                address 10.0.29.1/30;
            }
            family mpls;
        }
    }
     fe-0/1/0  { Connected to R1
        unit 0 {
            family inet {
                address 10.0.12.14/30;
```

```
            }
            family mpls;
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.70.144/21;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.2.1/32;
            }
        }
    }
}
forwarding-options {
    hash-key {
         family mpls { #MPLS labels configuration
            label-1;
            label-2;
            payload {
                ip;
            }
        }
    }
}
routing-options {
    static {
        [...Output truncated...]
    }
    router-id 192.168.2.1;
    autonomous-system 65432;
    forwarding-table {
         export lbpp; #Routes exported into forwarding table
    }
}
protocols {
    rsvp {
        interface so-0/0/1.0;
        interface fe-0/1/0.0;
        interface so-0/0/2.0;
        interface fxp0.0 {
            disable;
        }
    }
    mpls {
        interface fe-0/1/0.0;
        interface so-0/0/1.0;
        interface so-0/0/2.0;
        interface fxp0.0 {
            disable;
        }
    }
    bgp {
        group internal {
            type internal;
            local-address 192.168.2.1;
```

```
            neighbor 192.168.1.1;
            neighbor 192.168.4.1;
            neighbor 192.168.9.1;
            neighbor 192.168.6.1;
            neighbor 192.168.0.1;
        }
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface fe-0/1/0.0;
            interface so-0/0/1.0;
            interface so-0/0/2.0;
            interface lo0.0 {
                passive;
            }
        }
    }
}
policy-options {
    policy-statement  lbpp { #Load balancing policy exported in forwarding table

        then {
            load-balance per-packet;
        }
    }
}
```

Sample Output 4 The following sample output is for transit router R4:

```
user@R4> show configuration | no-more
interfaces {
    so-0/0/1 {
        unit 0 {
            family inet {
                address 10.0.24.2/30;
            }
            family mpls;
        }
    }
    so-0/0/3 {
        unit 0 {
            family inet {
                address 10.0.49.1/30;
            }
            family mpls;
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.70.146/21;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.4.1/32;
            }
        }
```

```
    }
}
routing-options {
    static {
        [...Output truncated...]
    }
    router-id 192.168.4.1;
    autonomous-system 65432;
}
protocols {
    rsvp {
        interface so-0/0/1.0;
        interface so-0/0/3.0;
        interface fxp0.0 {
            disable;
        }
    }
    mpls {
        interface so-0/0/1.0;
        interface so-0/0/3.0;
        interface fxp0.0 {
            disable;
        }
    }
    bgp {
        group internal {
            type internal;
            local-address 192.168.4.1;
            neighbor 192.168.1.1;
            neighbor 192.168.2.1;
            neighbor 192.168.9.1;
            neighbor 192.168.6.1;
            neighbor 192.168.0.1;
        }
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface so-0/0/1.0;
            interface so-0/0/3.0;
            interface lo0.0 {
                passive;
            }
        }
    }
}
```

Sample Output 5 The following sample output is for transit router R9:

```
user@R9> show configuration | no-more
interfaces {
    so-0/0/2 {
        unit 0 {
            family inet {
                address 10.0.29.2/30;
            }
            family mpls;
        }
    }
    so-0/0/3 {
        unit 0 {
```



```
            family inet {
                address 10.0.49.2/30;
            }
            family mpls;
        }
    }
    fe-0/1/0 {
        unit 0 {
            family inet {
                address 10.0.90.13/30;
            }
            family mpls;
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.69.206/21;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.9.1/32;
            }
        }
    }
}
routing-options {
    static {
        [...Output truncated...]
    }
    router-id 192.168.9.1;
    autonomous-system 65432;
}
protocols {
    rsvp {
        interface so-0/0/2.0;
        interface so-0/0/3.0;
        interface fe-0/1/0.0;
        interface fxp0.0 {
            disable;
        }
    }
    mpls {
        interface so-0/0/2.0;
        interface so-0/0/3.0;
        interface fe-0/1/0.0;
        interface fxp0.0 {
            disable;
        }
    }
    bgp {
        group internal {
            type internal;
            local-address 192.168.9.1;
            neighbor 192.168.1.1;
            neighbor 192.168.2.1;
            neighbor 192.168.4.1;
            neighbor 192.168.0.1;
```

```
            neighbor 192.168.6.1;
        }
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface so-0/0/2.0;
            interface so-0/0/3.0;
            interface fe-0/1/0.0;
            interface lo0.0 {
                passive;
            }
        }
    }
}
```

Sample Output 6 The following sample output is for egress router **R0**:

```
user@R0> show configuration | no-more
interfaces {
    so-0/0/2 {
        unit 0 {
            family inet {
                address 10.0.29.2/30;
            }
            family mpls;
        }
    }
    so-0/0/3 {
        unit 0 {
            family inet {
                address 10.0.49.2/30;
            }
            family mpls;
        }
    }
    fe-0/1/0 {
        unit 0 {
            family inet {
                address 10.0.90.13/30;
            }
            family mpls;
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.69.206/21;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.9.1/32;
            }
        }
    }
}
routing-options {
    static {
```



```
            [...Output truncated...]
        }
        router-id 192.168.9.1;
        autonomous-system 65432;
    }
    protocols {
        rsvp {
            interface so-0/0/2.0;
            interface so-0/0/3.0;
            interface fe-0/1/0.0;
            interface fxp0.0 {
                disable;
            }
        }
        mpls {
            interface so-0/0/2.0;
            interface so-0/0/3.0;
            interface fe-0/1/0.0;
            interface fxp0.0 {
                disable;
            }
        }
        bgp {
            group internal {
                type internal;
                local-address 192.168.9.1;
                neighbor 192.168.1.1;
                neighbor 192.168.2.1;
                neighbor 192.168.4.1;
                neighbor 192.168.0.1;
                neighbor 192.168.6.1;
            }
        }
        ospf {
            traffic-engineering;
            area 0.0.0.0 {
                interface so-0/0/2.0;
                interface so-0/0/3.0;
                interface fe-0/1/0.0;
                interface lo0.0 {
                    passive;
                }
            }
        }
    }
```

Using Bandwidth to Unevenly Load-Balance RSVP LSPs

Purpose With RSVP LSPs, load-balancing LSP traffic using bandwidth allows uneven load balancing across multiple external links that have varying amounts of available bandwidth. When you use bandwidth to load-balance an RSVP LSP, the distribution of traffic is proportional to the bandwidth configuration of each LSP. You configure load balancing at the **[edit protocols rsvp]** hierarchy level on the ingress router.

For uneven load balancing using bandwidth to work, you must have at least two equal-cost LSPs toward the same egress router and at least one of the LSPs must have a bandwidth value configured at the **[edit protocols mpls label-switched-path *lsp-path-name*]** hierarchy level. If no LSPs have bandwidth configured, equal distribution load balancing is

performed. If only some LSPs have bandwidth configured, the LSPs without any bandwidth configured do not receive any traffic.

Keep the following information in mind when you use the **load-balance** statement at the **[edit protocols rsvp]** hierarchy level:

- The behavior of currently running LSPs is not altered. To force the currently running LSPs to use the new behavior, issue the **clear mpls lsp** command.
- The **load-balance** statement at the **[edit protocols rsvp]** hierarchy level only applies to ingress LSPs that have a policy with the **load-balancing per-packet** statement configured.
- For Differentiated Services-aware traffic-engineered LSPs, the bandwidth of an LSP is calculated by summing the bandwidth of all of the class types.

Before you can use bandwidth to unevenly load-balance LSP traffic, you must have the following configured on the ingress router:

- A policy with the **load-balance per-packet** statement at the **[edit policy-options]** hierarchy level and that policy applied as an export policy at the **[edit forwarding-options]** hierarchy level. For more information about configuring load balancing, see “Configuring and Verifying Load Balancing” on page 67.
- Bandwidth configured for each LSP at the **[edit protocols mpls label-switched-path *lsp-path-name*]** hierarchy level. For more information on configuring LSP bandwidth, see the *Junos MPLS Applications Configuration Guide*.

Figure 13 on page 114 illustrates a network configured with RSVP bandwidth.

Figure 13: RSVP Bandwidth Network

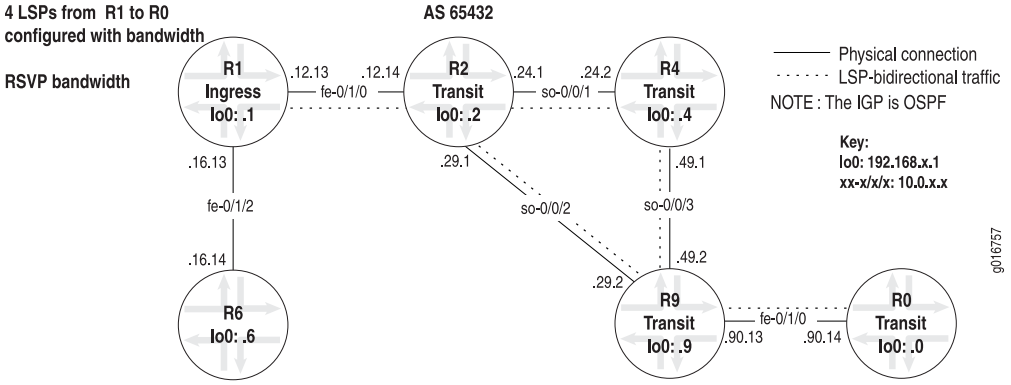

The network topology in Figure 13 on page 114 illustrates a router-only network with SONET and Ethernet interfaces that consists of the following components:

- A full-mesh IBGP topology, using AS 65432
- MPLS and RSVP enabled on all routers
- A send-statics policy on routers **R1** and **R0** that allows a new route to be advertised into the network
- Four unidirectional LSPs configured with uneven bandwidth between **R1** and **R0**

- One reverse direction LSP between **R0** and **R1**, which allows for bidirectional traffic
- Load balancing configured on the ingress router **R1**
- RSVP bandwidth configured on the ingress router **R1**

In addition, the example network uses OSPF as the IGP with OSPF area **0.0.0.0**. An IGP is required for the CSPF LSP, which is the default for the Junos OS. Also, the example network uses a policy to create BGP traffic. For the full configuration of routers in this network, see "Router Configurations for Bandwidth Load Balancing" on page 118.

The following information is included in this example:

Configure Bandwidth to Unevenly Load-Balance Traffic

Purpose Configuring bandwidth to unevenly load-balance traffic is performed in three stages. The first stage enables a load-balancing policy, the second stage configures the LPS bandwidth, and the third stage enables RSVP load balancing.

Action To configure bandwidth to unevenly load-balance RSVP LSPs, follow these steps:

1. Ensure that you have load balancing configured: see "Configuring and Verifying Load Balancing" on page 67.
2. Configure LSP bandwidth. In configuration mode, go to the following hierarchy level:

   ```
   [edit]
   user@host# edit protocols mpls
   ```

3. Configure the LSP bandwidth:

   ```
   [edit protocols mpls]
   user@host# set label-switched-path lsp-path-name bandwidth bps
   ```

4. Verify the configuration:

   ```
   [edit protocols mpls]
   user@host# show
   ```

5. Configure RSVP bandwidth. Go to the following hierarchy level:

   ```
   [edit]
   user@host# edit protocols rsvp
   ```

6. Configure the bandwidth statement:

   ```
   [edit protocols rsvp]
   user@host# set load-balance bandwidth
   ```

7. Verify and commit the configuration:

   ```
   user@host# show
   user@host# commit
   ```

Sample Output

```
user@R1> edit
Entering configuration mode
```

```
[edit]
user@R1# edit protocols mpls

[edit protocols mpls]
user@R1# set label-switched-path lsp1 bandwidth 10m

[edit protocols mpls]
user@R1# show
label-switched-path lsp1 {
    to 192.168.0.1;
    install 10.0.90.14/32 active;
    bandwidth 10m;
    primary via-r4;

[edit protocols mpls]
user@R1# top

[edit]
user@R1# edit protocols rsvp

[edit protocols rsvp]
user@R1# set load-balance bandwidth

[edit protocols rsvp]
user@R1# show
load-balance bandwidth;
interface fe-0/1/2.0;
interface fxp0.0 {
    disable;
}

[edit protocols rsvp]
user@R1# commit
commit complete
```

Meaning The sample output shows the configuration of LSP bandwidth and RSVP bandwidth on ingress router **R1**. The sample output shows only one LSP configured with bandwidth, however, for RSVP bandwidth to work, you must have at least two equal-cost LSPs toward the same egress router and at least one of the LSPs must have a bandwidth value configured. If no LSPs have bandwidth configured, equal-distribution load balancing is performed. If only some LSPs have bandwidth configured, the LSPs without any bandwidth configured do not receive any traffic.

Verify the Operation of Uneven Bandwidth Load Balancing

Purpose When a router is performing unequal-cost load balancing between LSPs paths, the **show route detail** command displays a balance field associated with each next hop being used.

Action To verify that an RSVP LSP is unevenly load-balanced, use the following Junos OS CLI operational mode commands:

```
user@host> show route protocol rsvp detail
user@host> show mpls lsp statistics
```

Sample Output

```
user@R1> show route protocol rsvp detail
```

```
inet.0: 25 destinations, 25 routes (25 active, 0 holddown, 0 hidden)
10.0.90.14/32 (1 entry, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7
                Next-hop reference count: 7
                Next hop: 10.0.12.14 via fe-0/1/0.0 weight 0x1  balance 10%
                 Label-switched-path lsp1
                Label operation: Push 100768
                Next hop: 10.0.12.14 via fe-0/1/0.0 weight 0x1  balance 20%
                 Label-switched-path lsp2
                Label operation: Push 100736
                Next hop: 10.0.12.14 via fe-0/1/0.0 weight 0x1  balance 30%,
      selected
                 Label-switched-path lsp3
                Label operation: Push 100752
                Next hop: 10.0.12.14 via fe-0/1/0.0 weight 0x1  balance 40%
                 Label-switched-path lsp4
                Label operation: Push 100784
                State: <Active Int>
                Local AS: 65432
                Age: 8:03       Metric: 4
                Task: RSVP
                Announcement bits (2): 0-KRT 4-Resolve tree 1
                AS path: I
inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
192.168.0.1/32 (1 entry, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7
                Next-hop reference count: 7
                Next hop: 10.0.12.14 via fe-0/1/0.0 weight 0x1 balance 10%
                Label-switched-path lsp1
                Label operation: Push 100768
                Next hop: 10.0.12.14 via fe-0/1/0.0 weight 0x1 balance 20%
                Label-switched-path lsp2
                Label operation: Push 100736
                Next hop: 10.0.12.14 via fe-0/1/0.0 weight 0x1 balance 30%
                Label-switched-path lsp3
                Label operation: Push 100752
                Next hop: 10.0.12.14 via fe-0/1/0.0 weight 0x1 balance 40%,
selected
                Label-switched-path lsp4
                Label operation: Push 100784
                State: <Active Int>
                Local AS: 65432
                Age: 8:03       Metric: 4
                Task: RSVP
                Announcement bits (1): 1-Resolve tree 1
                AS path: I

user@R1> show mpls lsp statistics
Ingress LSP: 4 sessions
To              From            State     Packets            Bytes LSPname
192.168.0.1     192.168.1.1     Up          10067           845628 lsp1
192.168.0.1     192.168.1.1     Up          20026          1682184 lsp2
192.168.0.1     192.168.1.1     Up          29796          2502864 lsp3
192.168.0.1     192.168.1.1     Up          40111          3369324 lsp4
Total 4 displayed, Up 4, Down 0

Egress LSP: 1 sessions
To              From            State     Packets            Bytes LSPname
192.168.1.1     192.168.0.1     Up             NA               NA r0-r1
```

```
Total 1 displayed, Up 1, Down 0

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

Meaning The sample output from ingress router **R1** shows that traffic is distributed according to the LSP bandwidth configuration, as indicated by the **Balance: xx%** field. For example, **lsp1** has 10 Mbps of bandwidth configured, as reflected in the **Balance: 10%** field.

Router Configurations for Bandwidth Load Balancing

Purpose The configuration in this topic is for ingress router **R1** in the example network illustrated in Figure 13 on page 114. The configuration for the other five routers in the network are the same as those found in "Router Configurations for the Load-Balanced MPLS Network" on page 73.

Action To display a router configuration, use the following Junos OS CLI operational mode command:

```
user@host> show configuration | no-more
```

Sample Output

```
user@R1> show configuration | no-more
[....Output truncated...]
interfaces {
    fe-0/1/0 {
        unit 0 {
            family inet {
                address 10.0.12.13/30;
            }
            family mpls;
        }
    }
    fe-0/1/2 {
        unit 0 {
            family inet {
                address 10.0.16.13/30;
            }
            family mpls;
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.70.143/21;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.1.1/32;
            }
        }
    }
}
routing-options {
    static {
```

```
        [...Output truncated...]
        }
        route 100.100.1.0/24 reject;
    }
    router-id 192.168.1.1;
    autonomous-system 65432;
    forwarding-table {
        export lbpp;
    }
}
protocols {
    rsvp {
         load-balance bandwidth; #RSVP bandwidth load balancing
        interface fe-0/1/0.0;
        interface fe-0/1/2.0;
        interface fxp0.0 {
            disable;
        }
    }
    mpls {
        label-switched-path lsp1 {
            to 192.168.0.1;
            install 10.0.90.14/32 active;
             bandwidth 10m; #Bandwidth configured for each LSP
            primary via-r4;
        }
        label-switched-path lsp2 {
            to 192.168.0.1;
            install 10.0.90.14/32 active;
             bandwidth 20m; #Bandwidth configured for each LSP
            primary via-r2;
        }
        label-switched-path lsp3 {
            to 192.168.0.1;
            install 10.0.90.14/32 active;
             bandwidth 30m; #Bandwidth configured for each LSP
            primary via-r2;
        }
        label-switched-path lsp4 {
            to 192.168.0.1;
            install 10.0.90.14/32 active;
             bandwidth 40m; #Bandwidth configured for each LSP
            primary via-r4;
        }
        path via-r2 {
            10.0.29.2 loose;
        }
        path via-r4 {
            10.0.24.2 loose;
        }
        interface fe-0/1/0.0;
        interface fe-0/1/2.0;
        interface fxp0.0 {
            disable;
        }
    }
    bgp {
        export send-statics;
        group internal {
            type internal;
            local-address 192.168.1.1;
```

```
            neighbor 192.168.2.1;
            neighbor 192.168.4.1;
            neighbor 192.168.9.1;
            neighbor 192.168.6.1;
            neighbor 192.168.0.1;
        }
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface fe-0/1/0.0;
            interface fe-0/1/2.0;
            interface lo0.0 {
                passive;
            }
        }
    }
}
policy-options {
    policy-statement lbpp { Load balancing policy
        then {
            load-balance per-packet;
        }
    }
    policy-statement send-statics {
        term statics {
            from {
                route-filter 100.100.1.0/24 exact;
            }
            then accept;
        }
    }
}
```

Meaning The sample output shows the configuration for the ingress router **R1** in the example network illustrated in Figure 13 on page 114. The configuration for the other five routers in the network is the same as those found in "Router Configurations for the Load-Balanced MPLS Network" on page 73.

Traffic Flows Before Load Balancing

Purpose The following sample output illustrates the details to look for when you issue different **show** commands to check if traffic is balanced. The following output is before load balancing is configured and is taken from transit router R2 in the network shown in Figure 10 on page 72.

Action To check the distribution of traffic across interfaces and LSPs, use the following CLI operational mode commands:

```
user@host> show route | find mpls
user@host> monitor interface traffic
user@host> show mpls lsp statistics
```

Sample Output 1

```
user@R2> show route | find mpls

mpls.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```



```
0                  *[MPLS/0] 1d 00:12:08, metric 1
                      Receive
1                  *[MPLS/0] 1d 00:12:08, metric 1
                      Receive
2                  *[MPLS/0] 1d 00:12:08, metric 1
                      Receive
100112             *[RSVP/7] 13:10:36, metric 1
                    > via so-0/0/1.0, label-switched-path lsp1
100128             *[RSVP/7] 13:01:08, metric 1
                    > via so-0/0/1.0, label-switched-path lsp4
100144             *[RSVP/7] 00:26:49, metric 1
                    > to 10.0.12.13 via fe-0/1/0.0, label-switched-path r0-r6
100160             *[RSVP/7] 00:23:25, metric 1
                    > via so-0/0/2.0, label-switched-path lsp2
100176             *[RSVP/7] 00:23:25, metric 1
                    > via so-0/0/2.0, label-switched-path lsp3
```

Sample Output 2

```
user@R2> monitor interface traffic

R2                                 Seconds: 89                 Time: 14:33:09

Interface    Link  Input packets        (pps)     Output packets        (pps)
 so-0/0/0      Up              0          (0)                  0          (0)
 so-0/0/1   Up      90    (1)       91    (1)
 so-0/0/2   Up     118    (1)     100122    (0)
 so-0/0/3      Up              0          (0)                  0          (0)
 fe-0/1/0      Up         100119          (0)                115          (0)
 fe-0/1/1      Up              0          (0)                  0          (0)
 fe-0/1/2      Up              0          (0)                  0          (0)
 fe-0/1/3      Up              0          (0)                  0          (0)
 [...Output truncated...]
```

Sample Output 3

```
user@R2> show mpls lsp statistics
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 5 sessions
To              From            State     Packets            Bytes LSPname
192.168.0.1     192.168.6.1     Up              0                0 lsp2
192.168.0.1   192.168.6.1   Up    112026    22853304 lsp1
192.168.0.1     192.168.6.1     Up              0                0 lsp3
192.168.0.1     192.168.6.1     Up              0                0 lsp4
192.168.6.1     192.168.0.1     Up              0                0 r0-r6
Total 5 displayed, Up 5, Down 0
```

Meaning Sample Outputs 1 through 3 from transit router **R2** show that traffic is not balanced across LSPs or interfaces.

Sample Output 1 for the **show route** command shows that all LSPs have the same metric (**1**) to the destination, even though they are traversing different interfaces. **lsp1** and **lsp4** are using **so-0/0/1**, while **lsp2** and **lsp3** are using **so-0/0/2**.

Sample Output 2 for the **monitor interface traffic** command shows that traffic is not evenly balanced across interfaces **so-0/0/1** and **so-0/0/2**. Almost all traffic is going out **so-0/0/2**.

Sample Output 3 for the **show mpls lsp statistics** command shows that traffic across LSPs is not balanced. All traffic is going over **lsp1**.

Related Topics

For additional information about MPLS fast reroute and MPLS protection methods, see the following:

- *Junos Feature Guide*
- *Junos MPLS Applications Configuration Guide*
- Semeria, Chuck. *RSVP Signaling Extensions for MPLS Traffic Engineering.* White paper. 2002
- Semeria, Chuck. *IP Dependability: Network Link and Node Protection*. White paper. 2002
- RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*

The Junos OS uses the load-balancing function across different protocols and features. For information about other types of load balancing, see the following:

- Option: Optimizing VPLS Traffic Flows, *Junos Feature Guide*
- Protocol-Independent Load Balancing for Layer 3 VPNs, *Junos VPNs Configuration Guide*
- Load Balancing Among Multiple Monitoring Interfaces, *Junos Services Interfaces Configuration Guide*

PART 2

Case Studies

CHAPTER 6

Troubleshooting Fast Reroute

This case study describes a problem establishing Fast Reroute (FRR) link protection in a Multiprotocol Label Switching (MPLS)-based VPN. Specifically, FRR requires a load balancing policy for the correct installation of routes in the forwarding table and fast local repair. The principles and solution used in this case study apply to all forms of local protection. For an overview of local protection, see "Local Protection Checklist" on page 23.

The chapter includes a brief summary of the FRR problem within the context of an MPLS-based VPN, an example network scenario, and commands to troubleshoot and resolve the problem.

The troubleshooting process described in this case study should not be followed rigidly; it is a basis from which you can develop your own process to suit your particular situation.

Troubleshooting Fast Reroute Checklist

Problem This checklist provides the steps and commands to troubleshoot a problem establishing Fast Reroute (FRR) link protection in a Multiprotocol Label Switching (MPLS)-based VPN. Specifically, FRR requires a load balancing policy for the correct installation of routes in the forwarding table and fast local repair. The principles and solution used in this case study apply to all forms of local protection. For an overview of local protection, see "Local Protection Checklist" on page 23.The checklist provides links to a brief summary of the FRR problem within the context of an MPLS-based VPN, an example network scenario, and more details about commands used to troubleshoot and resolve the problem. (See Table 11 on page 125

The troubleshooting process described in this case study should not be followed rigidly; it is a basis from which you can develop your own process to suit your particular situation.

Solution

Table 11: Troubleshooting Fast Reroute Checklist

Tasks	Command or Action
"Fast Reroute Problem Overview" on page 126	

Table 11: Troubleshooting Fast Reroute Checklist *(continued)*

Tasks	Command or Action
"Symptom" on page 127	Local repair is taking about one second) to complete, which is slow. show route forwarding-table extensive
"Cause" on page 128	The forwarding table does not include the necessary next-hops to support local repair.
"Troubleshooting Commands" on page 128	**show configuration routing-instances** *routing-instance-name* **show bgp summary instance routing** -*instance-name* **show configuration protocols mpls** **show mpls lsp ingress** **show rsvp session ingress** **show rsvp session ingress detail** **show route table** *table destination* **detail** **show route forwarding-table vpn** *vpn* **destination** *destination* **extensive**
"Solution" on page 132	Enable load-balancing and ensure that multiple next-hop forwarding table entries appear in the forwarding table for each destination. **show configuration policy-options** **show configuration routing-options** **show route forwarding-table vpn** *vpn* **destination** *destination* **extensive**
"Conclusion" on page 134	A load balancing policy is required for link protection to work effectively. The principles are the same for the configuration of the **fast reroute** and the **node-link-protection** statements.
"Router Configurations" on page 134	show configuration | no-more

Fast Reroute Problem Overview

Problem Incorrect configuration is a common mistake when trying to establish protection for an MPLS LSP. Protection with either fast reroute or link protection requires a **per-packet load-balance policy** exported at the **[edit routing-options forwarding-table]** hierarchy level. Correctly configured protection for an MPLS LSP results in two next-hop forwarding table entries per destination, either an incoming MPLS label or an IP destination. For information on configuring FRR, see "MPLS FRR Protection Overview" on page 3.

Figure 14 on page 127 illustrates a network topology with link protection and load balancing enabled to ensure that routes are correctly placed in to the forwarding table.

Figure 14: Fast Reroute Problem Network

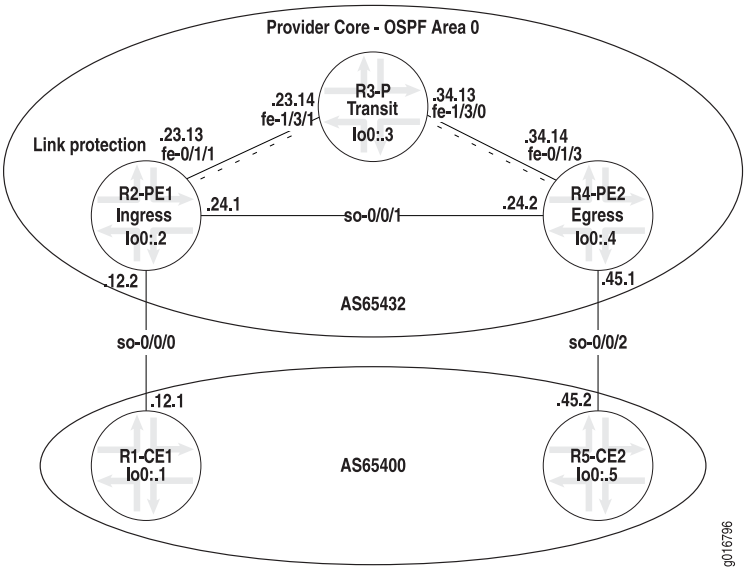

The network shown in Figure 14 on page 127 illustrates an MPLS-based VPN with traffic protection and load balancing, consisting of the following:

- All physical interfaces addresses are from the 10.0.x.x/30 address space.
- All loopback addresses are from the 192.168.x.1/32 block.
- The IGP is a single-area (Area 0) OSPF.
- RSVP is deployed as the MPLS signaling protocol between PE routers.
- LSPs (**r2-r4** and **r4-r2**) established between PE routers.
- MP-IBGP mesh between PE routers, loopback peering, and VPN-IPv4 NLRI.
- CE-PE link running EBGP.
- Full-mesh Layer 3 VPN between CE1 and CE2.
- Traffic protection for the link between the PE1 and P routers.
- Load balancing on PE1.

The overall goal of this network is to provide point-to-point connectivity between the two CE routers and traffic protection in the core of the network.

Symptom In the network shown in Figure 14 on page 127, the external symptom is that local repair is taking about one second to complete, which is slow. Use the **show route forwarding-table vpn vpn-a destination** command to check that the correct routes are included in the forwarding table. In the example output below, there is only one route installed in the forwarding table, when for fast local repair, there should be multiple next hops installed.

Sample Output

```
user@R2-PE1> show route forwarding-table vpn vpn-a destination 192.168.5.1 extensive

Routing table: vpn-a.inet [Index 2]
Internet:
```

```
Destination:  192.168.5.0/24
  Route type: user
  Route reference: 0                    Route interface-index: 0
  Flags: sent to PFE, prefix load balance
  Next-hop type: indirect               Index: 262142   Reference: 2
  Next-hop type: Push 100160
   Next-hop interface: so-0/0/1.0 #Only one next hop in the forwarding table.
```

Cause Slow local repair is caused by the forwarding table not including the necessary next-hops to support local repair. The forwarding table shows only a single next-hop, when local repair requires additional next-hops for fast recovery.

Troubleshooting Commands The Junos OS includes commands that are useful when troubleshooting a problem. This topic provides a brief description of each command followed by sample output, and a discussion of the output in relation to the problem.

The following commands can be used when troubleshooting a fast reroute error in an MPLS-VPN network:

```
user@R2-PE1> show configuration routing-instances vpn-a
user@R2-PE1> show configuration routing-options
user@R2-PE1> show bgp summary instance vpn-a
user@R2-PE1> show configuration protocols mpls
user@R2-PE1> show mpls lsp ingress
user@R2-PE1> show rsvp session ingress
user@R2-PE1> show rsvp session ingress detail
user@R2-PE1> show route table vpn-a 192.168.5.1 detail
user@R2-PE1> show route forwarding-table vpn vpn-a destination 192.168.5.1 extensive
```

Sample Output The **show configuration *statement-path*** command is used to display a specific configuration hierarchy; in this case, to verify the correct configuration of a specific routing instance named **vpn-a**.

```
user@R2-PE1> show configuration routing-instances vpn-a
instance-type  vrf ;
interface  so-0/0/0.0 ;
vrf-target {
    import target:65432:100;
    export target:65432:100;
}
protocols {
    bgp {
        group CE1 {
            type external;
             peer-as 65400;
             neighbor 10.0.12.1 ;
        }
    }
}
```

Meaning The sample output for the **show configuration** command shows the current running configuration of the specific routing instance named **vpn-a** configured on the ingress PE1 router. The **vpn-a** instance configuration has a VRF table that supports EBGP routing on

the PE-CE link (**so-0/0/0.0**). This interface is the correct interface for the CE1-PE1 link in the network topology shown in Figure 14 on page 127.

The VRF instance is linked to a VFR target community configured at the [edit policy-options] hierarchy level, allowing advertising of L3 VPN routes between PE routers. (See the PE1 configuration in "Router Configurations" on page 134 for the policy options configuration.) The import statement places, into the vpn-a.inet.0 table, all received L3 VPN MP-BGP routes tagged with the correct target community. The export statement advertises and tags all routes in the vpn-a.inet.o table with the listed target community to all MP-BGP peers.

The BGP protocols configuration within the routing instance applies the BGP import and export policies to the exchange of BGP routes on the PE-CE routing instance.

Sample Output The **show bgp summary** command is used to display summary information about BGP and its neighbors to determine if routes are received from peers in the autonomous system (AS). In this case, information for the specified instance **vpn-a** is displayed.

```
user@R2-PE1> show bgp summary instance vpn-a
Groups: 1 Peers: 1 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
vpn-a.inet.0          11          7          0          0          0          0
Peer               AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Damped...
10.0.12.1    65400    2471    2473    0    0  20:35:20 Establ
  vpn-a.inet.0: 5/5/0
```

Meaning The sample output for the **show bgp summary instance vpn-a** command shows that the peering session between the CE1 and PE1 routers is established, indicating that the peers are exchanging update messages.

Sample Output The **show configuration *statement-path*** command is used to display a specific configuration hierarchy; in this case, the MPLS hierarchy.

```
user@R2-PE1> show configuration protocols mpls
label-switched-path r2-r4   {
    to 192.168.4.1;
     link-protection  ;
     primary direct  ;
}
path direct {
    10.0.24.2 strict;
}
interface all;
interface fxp0.0 {
    disable;
}
```

Meaning The sample output for the **show configuration protocols mpls** command shows the current running MPLS configuration on the ingress PE1 router. The configuration include the LSP **r2-r4**, link protection, and the strict primary path **direct**.

Sample Output The **show mpls lsp** command is used to display summarized information about the configured and active LSPs on a router; in this case, the command shows only the ingress LSPs on the ingress PE1 router.

```
user@R2-PE1> show mpls lsp ingress
Ingress LSP: 1 sessions
To              From            State Rt ActivePath       P     LSPname
192.168.4.1     192.168.2.1     Up     0 direct           *     r2-r4
Total 1 displayed, Up 1, Down 0
```

Meaning The sample output for the **show mpls lsp ingress** command shows that the ingress LSP **r2-r4** is up and following the configured path **direct**.

Sample Output The **show rsvp session** command is used to display summarized information about active RSVP sessions on a router; in this case, the command shows summarized information about ingress RSVP sessions on the PE1 router

```
user@R2-PE1> show rsvp session ingress
Ingress RSVP: 2 sessions
To              From            State   Rt Style Labelin Labelout LSPname
192.168.4.1     192.168.2.1     Up       0  1  SE        -        3  r2-r4
192.168.4.1     192.168.2.1     Up       0  1  SE        -   100064
Bypass->10.0.24.2
Total 2 displayed, Up 2, Down 0
```

Meaning The sample output for the **show rsvp session ingress** command shows two RSVP sessions are up; the main LSP **r2-r4** and a bypass path protecting the main LSP. Both RSVP sessions are in the Shared Explicit (**SE**) style, creating a shared reservation among for the two paths.

Sample Output The **show rsvp session ingress detail** command is used to display more detailed information about the two ingress RSVP sessions on the PE1 router.

```
user@R2-PE1> show rsvp session ingress detail
Ingress RSVP: 2 sessions

192.168.4.1
   From: 192.168.2.1, LSPstate: Up , ActiveRoute: 0
   LSPname: r2-r4,  LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 SE, Label in: -, Label out: 3
  Time left:    -, Since: Fri Mar  9 14:05:03 2007
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 63395 protocol 0
  Link protection desired
   Type: Link protected LSP
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.24.2 (so-0/0/1.0) 2008 pkts
  RESV rcvfrom: 10.0.24.2 (so-0/0/1.0) 2006 pkts
  Explct route: 10.0.24.2
  Record route: <self> 10.0.24.2

192.168.4.1
```



```
  From: 192.168.2.1, LSPstate: Up, ActiveRoute: 0
   LSPname: Bypass->10.0.24.2
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 100064
  Resv style: 1 SE, Label in: -, Label out: 100064
  Time left:    -, Since: Fri Mar  9 14:05:58 2007
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 63396 protocol 0
   Type: Bypass LSP
     Number of data route tunnel through: 1
    Number of RSVP session tunnel through: 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.23.14 (fe-0/1/1.0) 2001 pkts
  RESV rcvfrom: 10.0.23.14 (fe-0/1/1.0) 1736 pkts
   Explct route: 10.0.23.14 10.0.34.14
  Record route: <self> 10.0.23.14 10.0.34.14
Total 2 displayed, Up 2, Down 0
```

Meaning The sample output for the **show rsvp session ingress detail** command shows the RSVP session for the ingress LSP and the bypass path, which appears as a separate RSVP ingress session for the protected interface **10.0.24.2**. The bypass path is automatically generated. By default, the name appears as **Bypass >** ***interface-address,*** where the interface address is the next downstream router's interface (**10.0.24.2**). The explicit route **10.0.23.14 10.0.34.14** for the session shows **R3** as the transit node and **R4** as the egress node.

Sample Output The **show route table***routing-table-name* command is used to display information about a particular routing table. In this case, the **vpn-a.inet.0** routing table.

```
user@R2-PE1> show route table vpn-a 192.168.5.1 detail
vpn-a.inet.0: 9 destinations, 13 routes (9 active, 0 holddown, 0 hidden)
192.168.5.0/24 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 192.168.4.1:4
                Next-hop reference count: 11
                Source: 192.168.4.1
                Next hop: via so-0/0/1.0 weight 0x1, selected
                Label-switched-path r2-r4
                Label operation: Push 100160
                Next hop: 10.0.23.14 via fe-0/1/1.0 weight 0x8001
                Label-switched-path r2-r4
                Label operation: Push 100160, Push 100064(top)
                Protocol next hop: 192.168.4.1
                Push 100160
                Indirect next hop: 8791000 262142
                State: <Secondary Active Int Ext>
                Local AS: 65432 Peer AS: 65432
                Age: 1d 5:22:31         Metric2: 1
                Task: BGP_65432.192.168.4.1+2056
                Announcement bits (1): 0-KRT
                AS path: 65400 I
                Communities: target:65432:100
                VPN Label: 100160
                Localpref: 100
```

```
Router ID: 192.168.4.1
Primary Routing Table bgp.l3vpn.0
```

Meaning The sample output for the **show route table vpn-a 192.168.5.1 detail** command shows routes associated with the remote PE-CE location as indicated by the loopback address of the PE2 router **192.168.5.1**. In this case, there are different next hops with unequal weights (**0x1** and **0x8001**) associated with the remote location. For correct traffic protection, those two routes must appear in the forwarding table.

Sample Output The **show route forwarding-table** command displays the route entries in the kernel's forwarding table. This is the version of the forwarding table in the Routing Engine. The Routing Engine copies this table to the Packet Forwarding Engine. In this case, the set of routes installed in the forwarding table to verify that the routing protocol process (rpd) has relayed the correct information to the forwarding table for the specified destination.

```
user@R2-PE1> show route forwarding-table vpn vpn-a destination 192.168.5.1
extensive
Routing table: vpn-a.inet [Index 2]
Internet:

Destination:  192.168.5.0/24
  Route type: user
  Route reference: 0                    Route interface-index: 0
  Flags: sent to PFE, prefix load balance
  Next-hop type: indirect               Index: 262142   Reference: 2
  Next-hop type: Push 100160
  Next-hop interface: so-0/0/1.0
```

Meaning The sample output for the **show route forwarding-table vpn vpn-a destination 192.168.5.1 extensive** command shows only one next hop **so-0/0/1.0** is installed in the forwarding table, indicating that the information in the forwarding table is not correct. We would expect to see the same paths installed in the forwarding table as appear in the routing table in the output for the **show route table vpn-a 192.168.5.1 detail**.

Solution The solution is to enable load-balancing and ensure that multiple next-hop forwarding table entries appear in the forwarding table for each destination. The forwarding-table entries can be an incoming MPLS label or an IP destination.

A load-balancing policy applied to the forwarding-table is the same mechanism required for ECMP (equal-cost multipath) load-balancing to install multiple next-hops into the forwarding-table. The extra paths installed for local repair are not used for load-balancing, because the paths are differently weighted, as demonstrated in the sample output for the **show routing table** and the **show route forwarding-table** commands.

NOTE: The load-balancing policy must be applied to all provider (P) and provider-edge (PE) routers that are required to support local repair.

The following sample output shows an example load-balancing configuration and the commands used to verify that the required two next-hop entries appear in the forwarding table.

Sample Output Use the following two **show configuration *statement-path*** commands to display a specific configuration hierarchy; in this case, policy-options and routing-options.

```
user@R2-PE1> show configuration policy-options
policy-statement lbpf {
    then {
         load-balance per-packet ;
    }
}
[...Output truncated...]

user@R2-PE1> show configuration routing-options
static {
    [...Output truncated...]
    route 100.100.1.0/24 reject;
}
router-id 192.168.2.1;
route-distinguisher-id 192.168.2.1;
autonomous-system 65432;
forwarding-table {
    export lbpf ;
}
```

Meaning The sample output for the **show configuration policy-options** and **show configuration routing-options** commands shows the two parts required to configure a load balancing policy. The **lbpf** policy includes the **load-balance per-packet** statement. The policy is then applied at the **[edit routing options forwarding-table]** hierarchy level with the **export lbpf** statement. Enabling load balancing results in the export of routes from the routing table to the forwarding table, and a solution to the problem.

NOTE: The load-balance per-packet statement is named *per-packet* for historical reasons. When the Packet Forwarding Engine was an IP Processor-I (before Junos 4.0), Junos supported only per-packet load balancing. When the IP Processor-II was introduced the behavior was changed to per-flow load balancing without changing the statement.

Sample Output Use the **show route forwarding-table** command to 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. In this case, the option **vpn *vpn*** is used to display routing table entries for the specified VPN **vpn-a**.

```
user@R2-PE1> show route forwarding-table vpn vpn-a destination 192.168.5.1 extensive

Routing table: vpn-a.inet [Index 2]
Internet:

Destination:  192.168.5.0/24
  Route type: user
  Route reference: 0                    Route interface-index: 0
  Flags: sent to PFE
  Next-hop type: indirect               Index: 262142   Reference: 2
  Next-hop type: unilist                Index: 262146   Reference: 1
```

```
  Next-hop type: Push 100160
  Next-hop interface: so-0/0/1.0     Weight: 0x1
  Nexthop: 10.0.23.14
  Next-hop type: Push 100160, Push 100064(top)
  Next-hop interface: fe-0/1/1.0     Weight: 0x8001
```

Meaning The sample output for the **show route forwarding-table vpn vpn-a destination 192.168.5.1 extensive** command shows the correct two routes were relayed from the routing table to the forwarding table.

Conclusion In conclusion, a load balancing policy is required for link protection to work effectively. The principles are the same for the configuration of the **fast reroute** and the **node-link protection** statements.

Router Configurations The following output shows the configurations of all routers in the network. The **no-more** option entered after the pipe **(|)** prevents the output from being paginated if the output is longer than the length of the terminal screen.

Sample Output The following sample output is for the customer edge (CE) 1 router:

```
user@R1-CE1> show configuration | no-more
[...Output truncated...]
interfaces {
    so-0/0/0 {
        unit 0 {
            family inet {
                address 10.0.12.1/30;
            }
            family iso;
            family mpls;
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.70.143/21;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.1.1/32;
            }
        }
    }
}
routing-options {
    static {
        /* corperate and alpha net */
        route 172.16.0.0/12 {
            next-hop 192.168.71.254;
            retain;
            no-readvertise;
        }
        /* old lab nets */
        route 192.168.0.0/16 {
```



```
            next-hop 192.168.71.254;
            retain;
            no-readvertise;
        }
        route 0.0.0.0/0 {
            discard;
            retain;
            no-readvertise;
        }
        route 172.16.0.0/24 reject;
        route 172.16.1.0/24 reject;
        route 172.16.2.0/24 reject;
        route 172.16.3.0/24 reject;
        route 192.168.1.0/24 reject;
    }
    router-id 192.168.1.1;
    autonomous-system 65400;
}
protocols {
    bgp {
        group PE1 {
            type external;
            export stat;
            peer-as 65432;
            neighbor 10.0.12.2;
        }
    }
    ospf {
        traffic-engineering;
        export stat;
        area 0.0.0.0 {
            interface so-0/0/0.0;
            interface lo0.0 {
                passive;
            }
        }
    }
}
policy-options {
    policy-statement stat {
        term 1 {
            from protocol static;
            then accept;
        }
        term 2 {
            then reject;
        }
    }
}
```

Sample Output The following sample output is for the provider edge (PE) 1 ingress router :

```
user@R2-PE1> show configuration | no-more
[...Output truncated...]
interfaces {
    so-0/0/0 {
        description to-r1;
        unit 0 {
            family inet {
                address 10.0.12.2/30;
            }
```

```
            family iso;
            family mpls;
        }
    }
    so-0/0/1 {
        description to-r4;
        unit 0 {
            family inet {
                address 10.0.24.1/30;
            }
            family iso;
            family mpls;
        }
    }
    fe-0/1/1 {
        description to-r3;
        unit 0 {
            family inet {
                address 10.0.23.13/30;
            }
            family iso;
            family mpls;
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.70.144/21;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.2.1/32;
            }
        }
    }
}
routing-options {
    static {
        route 172.16.0.0/12 {
            next-hop 192.168.71.254;
            retain;
            no-readvertise;
        }
        route 192.168.0.0/16 {
            next-hop 192.168.71.254;
            retain;
            no-readvertise;
        }
        route 0.0.0.0/0 {
            discard;
            retain;
            no-readvertise;
        }
        route 100.100.1.0/24 reject;
    }
    router-id 192.168.2.1;
    route-distinguisher-id 192.168.2.1;
    autonomous-system 65432;
```

```
    forwarding-table {
        export lbpf;
    }
}
protocols {
    rsvp {
        interface fxp0.0 {
            disable;
        }
        interface all {
            link-protection;
        }
    }
    mpls {
        label-switched-path r2-r4 {
            to 192.168.4.1;
            link-protection;
            primary direct;
        }
        path via-r3 {
            10.0.23.14 strict;
            10.0.34.14 strict;
        }
        path direct {
            10.0.24.2 strict;
        }
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
    bgp {
        export send-statics;
        group ibgp {
            type internal;
            local-address 192.168.2.1;
            family inet {
                unicast;
            }
            family inet-vpn {
                unicast;
            }
            export next-hop-self;
            peer-as 65432;
            neighbor 192.168.4.1;
        }
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface lo0.0 {
                passive;
            }
            interface fe-0/1/1.0;
            interface so-0/0/1.0;
        }
    }
}
policy-options {
    policy-statement lbpf {
        then {
```

```
                load-balance per-packet;
            }
        }
        policy-statement next-hop-self {
            from route-type external;
            then {
                next-hop self;
            }
        }
        policy-statement send-statics {
            term statics {
                from {
                    route-filter 100.100.1.0/24 exact;
                }
                then accept;
            }
        }
        policy-statement vpna-export {
            term 1 {
                from protocol static;
                then {
                    community add vpna-target;
                    community add vpna-origin;
                    accept;
                }
            }
            term 2 {
                then reject;
            }
        }
        policy-statement vpna-import {
            term 1 {
                from {
                    protocol bgp;
                    community vpna-target;
                }
                then accept;
            }
            term 2 {
                then reject;
            }
        }
        community vpna-origin members origin:192.168.2.1:1;
        community vpna-target members target:65432:100;
    }
    routing-instances {
        vpn-a {
            instance-type vrf;
            interface so-0/0/0.0;
            vrf-target {
                import target:65432:100;
                export target:65432:100;
            }
            protocols {
                bgp {
                    group CE1 {
                        type external;
                        peer-as 65400;
                        neighbor 10.0.12.1;
                    }
                }
```

```
        }
    }
}
```

Sample Output The following sample output is for the provider (P) transit router:

```
user@R3-P> show configuration | no-more
[...Output truncated...]
interfaces {
    fe-1/3/0 {
        description to-r4;
        unit 0 {
            family inet {
                address 10.0.34.13/30;
            }
            family iso;
            family mpls;
        }
    }
    fe-1/3/1 {
        description to-r2;
        unit 0 {
            family inet {
                address 10.0.23.14/30;
            }
            family iso;
            family mpls;
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.70.145/21;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.3.1/32;
            }
            family iso {
                address 49.0004.1921.6800.3001.00;
            }
        }
    }
}
routing-options {
    static {
        /* corperate and alpha net */
        route 172.16.0.0/12 {
            next-hop 192.168.71.254;
            retain;
            no-readvertise;
        }
        /* old lab nets */
        route 192.168.0.0/16 {
            next-hop 192.168.71.254;
            retain;
            no-readvertise;
        }
```

```
        route 0.0.0.0/0 {
            discard;
            retain;
            no-readvertise;
        }
    }
    router-id 192.168.3.1;
    autonomous-system 65432;
}
protocols {
    rsvp {
        interface all {
            link-protection;
        }
        interface fxp0.0 {
            disable;
        }
    }
    mpls {
        icmp-tunneling;
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface lo0.0 {
                passive;
            }
            interface fxp0.0 {
                disable;
            }
            interface all;
        }
    }
}
```

Sample Output The following sample output is for the provider edge (PE) 2 ingress router :

```
user@R4-PE2> show configuration | no-more
[...Output truncated...]
interfaces {
    so-0/0/1 {
        description to-R2;
        unit 0 {
            family inet {
                address 10.0.24.2/30;
            }
            family iso;
            family mpls;
        }
    }
    so-0/0/2 {
        description to-R5-CE2;
        unit 0 {
            family inet {
                address 10.0.45.1/30;
            }
            family iso;
```



```
            family mpls;
        }
    }
    fe-0/1/3 {
        description to-R3-P;
        unit 0 {
            family inet {
                address 10.0.34.14/30;
            }
            family iso;
            family mpls;
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.70.146/21;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.4.1/32;
            }
        }
    }
}
routing-options {
    static {
        route 172.16.0.0/12 {
            next-hop 192.168.71.254;
            retain;
            no-readvertise;
        }
        route 192.168.0.0/16 {
            next-hop 192.168.71.254;
            retain;
            no-readvertise;
        }
        route 0.0.0.0/0 {
            discard;
            retain;
            no-readvertise;
        }
        route 100.100.4.0/24 reject;
    }
    router-id 192.168.4.1;
    route-distinguisher-id 192.168.4.1;
    autonomous-system 65432;
    forwarding-table {
        export lbpf;
    }
}
protocols {
    rsvp {
        interface fxp0.0 {
            disable;
        }
        interface all {
            link-protection;
```

```
            }
        }
        mpls {
            label-switched-path r4-r2 {
                to 192.168.2.1;
            }
            interface all;
            interface fxp0.0 {
                disable;
            }
        }
        bgp {
            export send-statics;
            group ibgp {
                type internal;
                local-address 192.168.4.1;
                family inet {
                    unicast;
                }
                family inet-vpn {
                    unicast;
                }
                export next-hop-self;
                peer-as 65432;
                neighbor 192.168.2.1;
            }
        }
        ospf {
            traffic-engineering;
            area 0.0.0.0 {
                interface lo0.0 {
                    passive;
                }
                interface fe-0/1/3.0;
                interface so-0/0/1.0;
            }
        }
    }
    policy-options {
        policy-statement lbpf {
            then {
                load-balance per-packet;
            }
        }
        policy-statement next-hop-self {
            from route-type external;
            then {
                next-hop self;
            }
        }
        policy-statement send-statics {
            term statics {
                from {
                    route-filter 100.100.4.0/24 exact;
                }
                then accept;
            }
        }
        policy-statement vpnb-export {
            term 1 {
                from protocol static;
```

```
            then {
                community add vpnb-target;
                community add vpnb-origin;
                accept;
            }
        }
        term 2 {
            then reject;
        }
    }
    policy-statement vpnb-import {
        term 1 {
            from {
                protocol bgp;
                community vpnb-target;
            }
            then accept;
        }
        term 2 {
            then reject;
        }
    }
    community vpnb-origin members origin:192.168.5.1:1;
    community vpnb-target members target:65432:100;
}
routing-instances {
    vpn-b {
        instance-type vrf;
        interface so-0/0/2.0;
        vrf-target {
            import target:65432:100;
            export target:65432:100;
        }
        protocols {
            bgp {
                group CE2 {
                    type external;
                    peer-as 65400;
                    neighbor 10.0.45.2;
                }
            }
        }
    }
}
```

Sample Output The following sample output is for the customer edge (CE) 2 router:

```
user@R5-CE2> show configuration | no-more
[...Output truncated...]
interfaces {
    so-0/0/2 {
        unit 0 {
            family inet {
                address 10.0.45.2/30;
            }
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.70.147/21;
```

```
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.5.1/32;
            }
            family iso {
                address 49.0004.1921.6800.5001.00;
            }
        }
    }
}
routing-options {
    graceful-restart;
    static {
        /* corperate and alpha net */
        route 172.16.0.0/12 {
            next-hop 192.168.71.254;
            retain;
            no-readvertise;
        }
        /* old lab nets */
        route 192.168.0.0/16 {
            next-hop 192.168.71.254;
            retain;
            no-readvertise;
        }
        route 0.0.0.0/0 {
            discard;
            retain;
            no-readvertise;
        }
        route 172.16.0.0/24 reject;
        route 172.16.1.0/24 reject;
        route 172.16.2.0/24 reject;
        route 172.16.3.0/24 reject;
        route 192.168.5.0/24 reject;
    }
    router-id 192.168.5.1;
    autonomous-system 65400;
}
protocols {
    bgp {
        group PE2 {
            type external;
            export stat;
            peer-as 65432;
            neighbor 10.0.45.1;
        }
    }
    ospf {
        traffic-engineering;
        export stat;
        area 0.0.0.0 {
            interface so-0/0/2.0;
            interface lo0.0 {
                passive;
            }
        }
```

```
    }
}
policy-options {
    policy-statement stat {
        term 1 {
            from protocol static;
            then accept;
        }
        term 2 {
            then reject;
        }
    }
}
```

CHAPTER 7

Troubleshooting Link Protection for Multiple Bypass LSPs Overview

This case study simulates a network problem with link protection for multiple bypass paths for Resource Reservation Protocol (RSVP)-signaled LSPs (LSPs). It includes a brief summary of link protection, an example network scenario, and commands to troubleshoot and resolve the problem.

The troubleshooting process described in this case study should not be followed rigidly; it is a basis from which you can develop your own process to suit your particular situation.

Troubleshooting Link Protection for Multiple Bypass LSPs Checklist

Problem This checklist provides steps and command to troubleshoot a network problem with link protection for multiple bypass paths for Resource Reservation Protocol (RSVP)-signaled LSPs (LSPs). The checklist includes links to a brief summary of link protection, an example network scenario, and more detailed information about the commands to troubleshoot and resolve the problem.

The troubleshooting process described in this case study should not be followed rigidly; it is a basis from which you can develop your own process to suit your particular situation. (See Table 12 on page 147.

Solution

Table 12: Troubleshooting Link Protection for Multiple Bypass LSPs Checklist

Tasks	Command or Action
"Troubleshooting Link Protection for Multiple Bypass LSPs" on page 148	
• Symptom on page 149	One bypass LSP is pre-signaled instead of two. **show mpls lsp bypass**
• Cause on page 149	The bandwidth reserved on the primary LSP is served by only one bypass path.

Table 12: Troubleshooting Link Protection for Multiple Bypass LSPs Checklist *(continued)*

Tasks	Command or Action
• Troubleshooting Commands on page 149	**show mpls lsp** **show mpls lsp bypass extensive** **show rsvp session ingress detail** **show rsvp interface** **show rsvp interface *type-fpc/pic/port* extensive** **show configuration *statement-path***
• Solution on page 156	**show configuration *statement-path*** **show mpls lsp bypass** **show mpls lsp bypass extensive**
• Conclusion on page 159	Multiple bypass paths are pre-signaled when the bandwidth values in the configuration require multiple bypass paths.
• Router Configurations on page 159	**show configuration | no-more**

Troubleshooting Link Protection for Multiple Bypass LSPs

Problem Link protection (many-to-one or facility backup) allows a router immediately upstream from a link failure to use an alternate interface to forward traffic to its downstream neighbor. This is accomplished by preestablishing a bypass path that is shared by all protected LSPs traversing the failed link. A single bypass path can safeguard a set of protected LSPs. When an outage occurs, the router immediately upstream from the link outage switches protected traffic to the bypass link, and then signals the link failure to the ingress router.

In this simulation, the network administrator mistakenly expects two bypass paths to be pre-signaled to protect four LSPs over two interfaces. However, because of the bandwidth configuration on both interfaces and the RSVP protocol, only one bypass path is pre-signaled. The second bypass path is not pre-signaled because the existing bandwidth reserved on the primary LSP is served by one bypass path.

Figure 15 on page 148 illustrates the network topology used in this case study.

Figure 15: Link Protection for Multiple Bypass LSPs Network

	R1	R2
at-0/1/2	10.0.12.5	10.0 .12.6
fe-0/1/0	10.0.12.13	10.0 .12.14
so-0/0/0	10.0.12.1	10.0.12.2

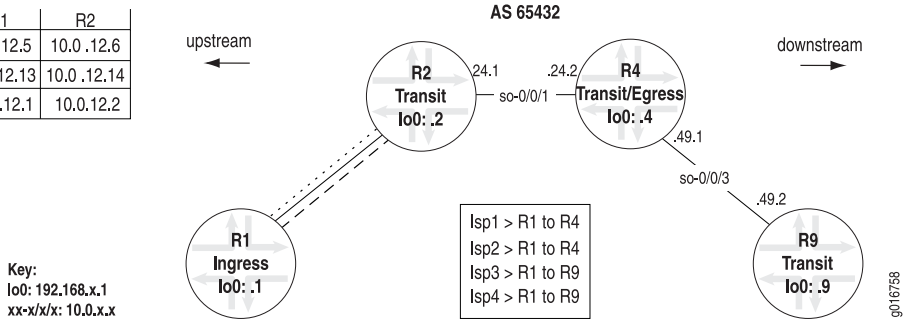

The MPLS network topology in Figure 15 on page 148 shows a router-only network with SONET, Fast Ethernet, and ATM interfaces that consists of the following components:

- A full-mesh internal BGP (IBGP) topology using AS 65432
- MPLS and RSVP are enabled on all routers
- A send-statics policy on routers **R1**, **R4**, and **R9** that allows a new route to be advertised into the network
- Six unidirectional LSPs between **R1** and **R4**, and **R1** and **R9**, with two LSPS running in the opposite direction to allow for bidirectional traffic
- Three interface connections between **R1** and **R4**, which allows for a primary LSP and two bypass paths on different interfaces
- Bandwidth configured for interfaces, RSVP, and LSPs

Sample configurations for all four routers in the network shown in Figure 15 on page 148 are provided at the end of this case study in “Router Configurations” on page 159.

Symptom In the network shown in Figure 15 on page 148, only one bypass LSP is pre-signaled instead of two, as shown in the following sample output.

Sample Output

```
user@R1> show mpls lsp bypass
Ingress LSP: 5 sessions
To              From            State   Rt Style Labelin Labelout LSPname
192.168.2.1   192.168.1.1   Up    0 1 SE    -    3 Bypass->10.0.12.14
Total 1 displayed, Up 1, Down 0

Egress LSP: 2 sessions
Total 0 displayed, Up 0, Down 0

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

Cause The cause of this problem is that bandwidth reserved for the primary LSPs is served by only one bypass path.

Troubleshooting Commands The Junos OS includes commands that are useful when troubleshooting a problem. This topic provides a brief description of each command, followed by sample output, and a discussion of the output in relation to the network shown in Figure 15 on page 148.

The following commands can be used when troubleshooting:

```
user@host> show mpls lsp
user@host> show mpls lsp bypass extensive
user@host> show rsvp session ingress detail
user@host> show rsvp interface
user@host> show rsvp interface type-fpc/pic/port extensive
user@host> show configuration statement-path
```

Sample Ouput

```
user@R1> show mpls lsp
Ingress LSP: 4 sessions
To              From            State Rt ActivePath       P     LSPname
192.168.4.1     192.168.1.1      Up      0 path1            *     lsp1
```

```
192.168.4.1     192.168.1.1      Up      0 path1            *     lsp2
192.168.9.1     192.168.1.1      Up     0 path1            *     lsp3
192.168.9.1     192.168.1.1      Up     0 path1            *     lsp4
Total 4 displayed, Up 4, Down 0

Egress LSP: 2 sessions
To              From            State   Rt Style Labelin Labelout LSPname
192.168.1.1     192.168.4.1     Up       0  1 FF       3        - r4-r1
192.168.1.1     192.168.9.1     Up       0  1 FF       3        - r9-r1
Total 2 displayed, Up 2, Down 0

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

What It Means The sample output of the **show mpls lsp** command shows that four LSPs originating from this router **R1** are up (ingress LSPs). The two LSPs originating at **R4** and **R9**, and terminating at **R1** are also up (egress LSPs). No LSPs are transiting this router (transit LSPs). In this case, all LSPs are up, indicating that the problem is not with the LSPs being in a down state.

Use the **show mpls lsp bypass extensive** command to display detailed information about LSPs used for protecting other LSPs (bypass LSPs).

Sample Output

```
user@R1> show mpls lsp bypass extensive
Ingress LSP: 5 sessions

192.168.2.1
  From: 192.168.1.1, LSPstate: Up , ActiveRoute: 0
   LSPname: Bypass->10.0.12.14 #This bypass path is from R1 to R2
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 SE, Label in: -, Label out: 3
  Time left:    -, Since: Fri Nov 10 08:29:27 2006
  Tspec: rate 100Mbps size 100Mbps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 45808 protocol 0
  Type: Bypass LSP
    Number of data route tunnel through: 4 #LSPs protected by this bypass path
    Number of RSVP session tunnel through: 0
   ActiveResv 4, PreemptionCnt 0, Update threshold 0%
  Subscription 100%,
  bc0 = ct0, StaticBW 100Mbps
  ct0: StaticBW 100Mbps, AvailableBW 60Mbps
    MaxAvailableBW 100Mbps = (bc0*subscription)
    ReservedBW [0] 40Mbps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7]0bps

  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.12.2 (so-0/0/0.0)  57 pkts
  RESV rcvfrom: 10.0.12.2 (so-0/0/0.0) 57 pkts
  Explct route: 10.0.12.2
  Record route: <self> 10.0.12.2
Total 1 displayed, Up 1, Down 0

Egress LSP: 2 sessions
Total 0 displayed, Up 0, Down 0
```



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

Meaning The sample output of the **show mpls lsp bypass extensive** command shows one bypass LSP (**Bypass->10.0.12.14**) from ingress router **R1** to transit router **R2**. All four of the ingress LSPs are protected by this single bypass path, as indicated by the **Number of data route tunnel through: 4** field. Interface **so-0/0/0.0** is the interface on which the bypass is pre-signaled. In this case study, the problem is that interface **at-0/1/2.0** is supposed to also have a pre-signaled bypass path, and the two LSPs should be protected by a bypass path on each interface (**so-0/0/0.0** and **at-0/1/2.0**).

Use the **show rsvp session ingress detail** command to display detailed information about RSVP sessions.

Sample Output

```
user@R1> show rsvp session ingress detail
Ingress RSVP: 5 sessions

192.168.2.1
  From: 192.168.1.1, LSPstate: Up, ActiveRoute: 0
  LSPname:  Bypass->10.0.12.14
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 SE, Label in: -, Label out: 3
  Time left:    -, Since: Fri Nov 10 08:29:27 2006
  Tspec: rate 100Mbps size 100Mbps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 45808 protocol 0
  Type: Bypass LSP
    Number of data route tunnel through: 4
    Number of RSVP session tunnel through: 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.12.2  (so-0/0/0.0) 60 pkts
  RESV rcvfrom: 10.0.12.2 (so-0/0/0.0) 61 pkts
  Explct route: 10.0.12.2
  Record route: <self> 10.0.12.2

192.168.4.1
  From: 192.168.1.1 , LSPstate: Up, ActiveRoute: 0
   LSPname: lsp1, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 101008
  Resv style: 1 SE, Label in: -, Label out: 101008
  Time left:    -, Since: Thu Nov  9 11:39:04 2006
  Tspec: rate 10Mbps size 10Mbps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 45673 protocol 0
   Link protection desired
   Type: Link protected LSP
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.12.14 (fe-0/1/0.0) 1880 pkts
  RESV rcvfrom: 10.0.12.14 (fe-0/1/0.0) 1838 pkts
  Explct route: 10.0.12.14 10.0.24.2
  Record route: <self> 10.0.12.14 10.0.24.2

192.168.4.1
  From: 192.168.1.1 , LSPstate: Up, ActiveRoute: 0
```

```
    LSPname: lsp2, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 101104
  Resv style: 1 SE, Label in: -, Label out: 101104
  Time left:    -, Since: Thu Nov  9 20:34:02 2006
  Tspec: rate 10Mbps size 10Mbps peak Infbps m 20 M 1500
  Port number: sender 2 receiver 45675 protocol 0
    Link protection desired
  Type: Link protected LSP
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.12.14 (fe-0/1/0.0) 1076 pkts
  RESV rcvfrom: 10.0.12.14 (fe-0/1/0.0) 1068 pkts
  Explct route: 10.0.12.14 10.0.24.2
  Record route: <self> 10.0.12.14 10.0.24.2

192.168.9.1
  From: 192.168.1.1 , LSPstate: Up, ActiveRoute: 0
  LSPname: lsp3, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 101120
  Resv style: 1 SE, Label in: -, Label out: 101120
  Time left:    -, Since: Thu Nov  9 20:34:02 2006
  Tspec: rate 10Mbps size 10Mbps peak Infbps m 20 M 1500
  Port number: sender 2 receiver 45685 protocol 0
  Link protection desired
  Type: Link protected LSP
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.12.14 (fe-0/1/0.0) 1080 pkts
  RESV rcvfrom: 10.0.12.14 (fe-0/1/0.0) 1072 pkts
  Explct route: 10.0.12.14 10.0.24.2 10.0.49.2
  Record route: <self> 10.0.12.14 10.0.24.2 10.0.49.2

192.168.9.1
  From: 192.168.1.1 , LSPstate: Up, ActiveRoute: 0
  LSPname: lsp4, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 101136
  Resv style: 1 SE, Label in: -, Label out: 101136
  Time left:    -, Since: Thu Nov  9 20:34:02 2006
  Tspec: rate 10Mbps size 10Mbps peak Infbps m 20 M 1500
  Port number: sender 2 receiver 45687 protocol 0
  Link protection desired
  Type: Link protected LSP
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.12.14 (fe-0/1/0.0) 1076 pkts
  RESV rcvfrom: 10.0.12.14 (fe-0/1/0.0) 1068 pkts
  Explct route: 10.0.12.14 10.0.24.2 10.0.49.2
  Record route: <self> 10.0.12.14 10.0.24.2 10.0.49.2
Total 5 displayed, Up 5, Down 0
```

Meaning The sample output of the **show rsvp session ingress detail** command shows five RSVP sessions originating at ingress router **R1**. Each session is up and each LSP is protected by the bypass path **Bypass->10.0.12.14** on interface **so-0/0/0.0**. In this case study, two of the LSPs should be protected by a second bypass path on interface **at-0/1/2.0**.

Use the show rsvp interface command to display the status of RSVP-enabled interfaces and packet statistics.

Sample Output

```
user@R1> show rsvp interface
RSVP interface: 3 active
                  Active Subscr- Static      Available   Reserved    Highwater
Interface   State resv   iption  BW          BW          BW          mark
at-0/2/1.0  Up         0   100%  50Mbps        50Mbps       0bps        0bps

fe-0/1/0.0  Up         4   100%  100Mbps       60Mbps       40Mbps      100Mbps

so-0/0/0.0  Up         1   100%  100Mbps     0bps         100Mbps      100Mbps
```

Meaning The sample output of the **show rsvp interface** command shows that all RSVP interfaces are up with four reservations on the Fast Ethernet interface (**fe-0/1/0.0**), one reservation on the SONET interface (**so-0/0/0.0)**, and no reservations on the ATM interface **at-0/2/1.0**. The total interface bandwidth (**Static BW**) is 100 Mbps on the Fast Ethernet and SONET interfaces, and only 50 Mbps on the ATM interface, indicating that the SONET interface is providing enough bandwidth to satisfy the requirements of the primary path of all four LSPs. Therefore, there is no need for a second bypass path on the ATM interface with this configuration.

Sample Output Use the **show rsvp interface *interface-name* extensive** command to display detailed information about a specific interface. The extensive option provides output for the latest 50 events on this interface.

```
user@R1>   show rsvp interface fe-0/1/0.0 extensive
 fe-0/1/0.0  Index 66, State Ena/Up
  NoAuthentication, NoAggregate, NoReliable, LinkProtection
  HelloInterval 9(second)
  Address 10.0.12.13
  ActiveResv 4, PreemptionCnt 0, Update threshold 10%
  Subscription 100%,
  bc0 = ct0, StaticBW 100Mbps
  ct0: StaticBW 100Mbps, AvailableBW 60Mbps
    MaxAvailableBW 100Mbps = (bc0*subscription)
     ReservedBW [0] 40Mbps [1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7]
0bps
    Protection: On, Bypass: 1, LSP: 4, Protected LSP: 0, Unprotected LSP: 4
       1 Nov 10 09:48:12 New bypass Bypass->10.0.12.14
      Bypass: Bypass->10.0.12.14, State: Up, Type: LP, LSP: 0, Backup: 0
       4 Nov 10 09:49:13 Record Route:  10.0.12.2
       3 Nov 10 09:49:13 Up
       2 Nov 10 09:49:13 CSPF: computation result accepted
       1 Nov 10 09:48:43 CSPF failed: no route toward 10.0.12.14[2 times]
```

Meaning The sample output of the **show rsvp interface *interface-name* extensive** command shows one bypass path protecting four LSPs. The bypass path is pre-signaled on **10.0.12.2**, which is the SONET interface **so-0/0/0.0**. Also, the total amount of bandwidth that RSVP is allowed to reserve is 100 Mbps. 40 Mbps are reserved with 60 Mbps available, indicating that there is more than enough bandwidth available to meet the needs of the four LSPs with one bypass path.

Sample Output Use the **show configuration *statement-path*** command to display a specific configuration hierarchy; for example, routing protocols.

```
user@R1> show configuration protocols rsvp
interface fe-0/1/0.0 {
    link-protection {
        bandwidth 100m;
        max-bypasses 2;
    }
}
interface so-0/0/0.0;
interface at-0/2/1.0;
interface fxp0.0 {
    disable;
}
```

Meaning The sample output of the show configuration protocols rsvp command shows that the Fast Ethernet interface is configured with link protection, 100 Mbps of bandwidth, and two bypass paths. In this case study, the amount of bandwidth may need to be adjusted until two bypass paths are pre-signaled. The sample output of the **show configuration protocols rsvp** command shows that the Fast Ethernet interface is configured with link protection, 100 Mbps of bandwidth, and two bypass paths. In this case study, the amount of bandwidth may need to be adjusted until two bypass paths are pre-signaled.

Sample Output Use the **show configuration *statement-path*** command to display a specific configuration hierarchy; for example, interfaces.

```
user@R1> show configuration interfaces
so-0/0/0 {
    unit 0 {
         bandwidth 100m;
        family inet {
            address 10.0.12.1/32;
        }
        family mpls;
    }
}
fe-0/1/0 {
    unit 0 {
        family inet {
            address 10.0.12.13/30;
        }
        family mpls;
    }
}
at-0/2/1 {
    atm-options {
        pic-type atm2;
        vpi 0;
    }
    unit 0 {
         bandwidth 50m;
        vci 0.128;
        family inet {
            address 10.0.12.5/32 {
                destination 10.0.12.6;
            }
        }
```



```
        family mpls;
    }
}
[...Output truncated...]
```

Meaning The sample output of the **show configuration interface** command shows that the SONET interface is configured with 100 Mbps, while the ATM interface is configured with 50 Mbps. In this case study, the amount of bandwidth for each interface may need to be adjusted until two bypass paths are pre-signaled.

Sample Output Use the **show configuration *statement-path*** command to display a specific configuration hierarchy; for example, routing protocols.

```
user@R1> show configuration protocols mpls
mpls {
        label-switched-path lsp1 {
            from 192.168.1.1;
            to 192.168.4.1;
             bandwidth 10m;
             link-protection;
            primary path1;
        }
        label-switched-path lsp2 {
            from 192.168.1.1;
            to 192.168.4.1;
             bandwidth 20m;
             link-protection;
            primary path1;
        }
        label-switched-path lsp3 {
            from 192.168.1.1;
            to 192.168.9.1;
             bandwidth 30m;
             link-protection;
            primary path1;
        }
        label-switched-path lsp4 {
            from 192.168.1.1;
            to 192.168.9.1;
             bandwidth 40m;
             link-protection;
            primary path1;
        }
        path path1 {
            10.0.12.14 strict;
        }
        interface fe-0/1/0.0;
        interface so-0/0/0.0;
        interface at-0/2/1.0;
        interface fxp0.0 {
            disable;
        }
    }
```

Meaning The sample output of the **show configuration protocols mpls** command shows that the four LSPs are configured with different bandwidth values. In this case study, the bandwidth value for each LSP may need to be adjusted until two bypass paths are pre-signaled.

Solution Adjust the bandwidth for the interfaces, RSVP link protection, and LSPs until two bypass LSPs are pre-signaled. In this case study, the bandwidth value for the SONET interface and the protected Fast Ethernet interface was reduced. An adjustment was not made to the bandwidth of the LSPs. The bandwidth adjustment described in this case study should not be followed rigidly; it is a basis from which you can develop your own process of adjusting bandwidth that suits your particular situation.

For information on adjusting the bandwidth for interfaces, see the *Junos Network Interfaces Configuration Guide*. For information on adjusting the bandwidth for link protection, see "Checklist for Load Balancing in an MPLS Network" on page 63. For information on adjusting the LSP bandwidth, see "Checklist for Path Protection" on page 9.

The Junos OS includes commands that are useful when verifying the solution to a problem. This topic provides a brief description of each command, followed by sample output, and a discussion of the output in relation to the network shown in Figure 15 on page 148.

You can use the following commands when verifying the solution to a problem:

```
user@host> show configuration statement-path
user@host> show mpls lsp bypass
user@host> show mpls lsp bypass extensive
```

Sample Output Use the **show configuration *statement-path*** command to display a specific configuration hierarchy; for example, interfaces.

```
user@R1>  show configuration interfaces
so-0/0/0  {
    unit 0 {
         bandwidth 50m;
        family inet {
            address 10.0.12.1/32;
        }
        family mpls;
    }
}
[...Output truncated...]
at-0/2/1  {
    atm-options {
        pic-type atm2;
        vpi 0;
    }
    unit 0 {
         bandwidth 50m;
        vci 0.128;
        family inet {
            address 10.0.12.5/32 {
                destination 10.0.12.6;
            }
        }
        family mpls;
    }
}
[...Output truncated...]
```

Meaning The sample output of the **show configuration interfaces** command shows that the bandwidth for the SONET interface has been adjusted down from 100 Mbps to 50 Mbps. This adjustment did not in itself result in two bypass paths coming up. A further adjustment

to the link protection bandwidth was necessary before two bypass paths were pre-signaled.

Sample Output Use the **show configuration *statement-path*** command to display a specific configuration hierarchy; for example, routing protocols.

```
user@R1> show configuration protocols
protocols {
    rsvp {
        interface fe-0/1/0.0 {
             link-protection {
                 bandwidth 50m;
                 max-bypasses 2;
             }
        }
        interface fe-0/1/2.0;
        interface so-0/0/0.0;
        interface at-0/2/1.0;
        interface fxp0.0 {
            disable;
        }
    }
[...Output truncated...]
```

Meaning The sample output of the **show configuration interfaces** command shows that the bandwidth for link protection on the Fast Ethernet interface has been adjusted down from 100 Mbps to 50 Mbps. The sample output of the **show configuration interfaces** command shows that the bandwidth for link protection on the Fast Ethernet interface has been adjusted down from 100 Mbps to 50 Mbps.

Sample Output Use the show mpls lsp bypass command to display information about LSPs used for protecting other LSPs (bypass LSPs).Use the show mpls lsp bypass command to display information about LSPs used for protecting other LSPs (bypass LSPs).

```
user@R1> show mpls lsp bypass
Ingress LSP: 6 sessions
To              From            State   Rt Style Labelin Labelout LSPname
192.168.2.1     192.168.1.1     Up       0  1 SE       -        3 Bypass->10.0.12.14
192.168.2.1     192.168.1.1     Up       0  1 SE       -        3 Bypass->10.0.12.14-1
Total 2 displayed, Up 2, Down 0

Egress LSP: 2 sessions
Total 0 displayed, Up 0, Down 0

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

Meaning The sample output of the **show configuration interfaces** command shows that two bypass LSPs are pre-signaled (**Up**), **10.0.12.14** and **10.0.12.14-1**, indicating that reducing the bandwidth of the link-protected interface and the SONET interface was successful. The bandwidth adjustment made in this case study may be different from the adjustment that is required in your network.

Sample Output Use the show mpls lsp bypass extensive command to display detailed information about LSPs used for protecting other LSPs (bypass LSPs). The **no-more** option entered after

the pipe (|) prevents the output from being paginated if the output is longer than the length of the terminal screen.

```
user@R1> show mpls lsp bypass extensive | no-more
Ingress LSP: 6 sessions

192.168.2.1
  From: 192.168.1.1, LSPstate: Up, ActiveRoute: 0
   LSPname: Bypass->10.0.12.14
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 SE, Label in: -, Label out: 3
  Time left:    -, Since: Thu Nov  9 17:47:17 2006
  Tspec: rate 50Mbps size 50Mbps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 45762 protocol 0
  Type: Bypass LSP
     Number of data route tunnel through: 2
    Number of RSVP session tunnel through: 0
  ActiveResv 2, PreemptionCnt 0, Update threshold 0%
  Subscription 100%,
  bc0 = ct0, StaticBW 50Mbps
  ct0: StaticBW 50Mbps, AvailableBW 0bps
    MaxAvailableBW 50Mbps = (bc0*subscription)
    ReservedBW [0] 50Mbps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7]0bps

  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.12.6  (at-0/2/1.0) 213 pkts
  RESV rcvfrom: 10.0.12.6 (at-0/2/1.0) 213 pkts
  Explct route: 10.0.12.6
  Record route: <self> 10.0.12.6

192.168.2.1
  From: 192.168.1.1, LSPstate: Up, ActiveRoute: 0
  LSPname: Bypass->10.0.12.14-1
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 SE, Label in: -, Label out: 3
  Time left:    -, Since: Thu Nov  9 17:47:51 2006
  Tspec: rate 50Mbps size 50Mbps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 45764 protocol 0
  Type: Bypass LSP
     Number of data route tunnel through: 2
    Number of RSVP session tunnel through: 0
  ActiveResv 2, PreemptionCnt 0, Update threshold 0%
  Subscription 100%,
  bc0 = ct0, StaticBW 50Mbps
  ct0: StaticBW 50Mbps, AvailableBW 0bps
    MaxAvailableBW 50Mbps = (bc0*subscription)
    ReservedBW [0] 50Mbps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7]0bps

  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.12.2 (so-0/0/0.0) 212 pkts
  RESV rcvfrom: 10.0.12.2 (so-0/0/0.0) 212 pkts
  Explct route: 10.0.12.2
  Record route: <self> 10.0.12.2
Total 2 displayed, Up 2, Down 0
```

```
Egress LSP: 2 sessions
Total 0 displayed, Up 0, Down 0

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

Meaning The sample output of the **show mpls lsp bypass extensive** command shows two bypass LSPs (**Bypass->10.0.12.14** and **Bypass->10.0.12.14-1**) from ingress router **R1** to transit router **R2**. All four of the ingress LSPs are protected by the two bypass paths, as indicated by the **Number of data route tunnel through: 2** field in the output for each bypass LSP. The SONET interface **so-0/0/0.0** and the ATM interface **at-0/1/2.0** are the interfaces on which the bypass paths are pre-signaled.

Conclusion In this simulation, the network administrator mistakenly expected two bypass paths to be pre-signaled when the bandwidth configuration on the interfaces and the RSVP protocol required only one bypass path. After troubleshooting the example network scenario, and adjusting the bandwidth for the interfaces and link protection in the RSVP protocol, the second bypass path was pre-signaled, and the problem resolved.

In conclusion, multiple bypass paths are pre-signaled when the bandwidth values in the configuration require multiple bypass paths.

Router Configurations Output that shows the configurations of all routers in the network. The **no-more** option entered after the pipe **(|)** prevents the output from being paginated if the output is longer than the length of the terminal screen.

Sample Output 1 The following sample output is for ingress router R1:

```
user@R1> show configuration | no-more
interfaces {
    so-0/0/0 {
        unit 0 {
             bandwidth 50m;
            family inet {
                address 10.0.12.1/32;
            }
            family mpls;
        }
    }
    fe-0/1/0 {
        unit 0 {
            family inet {
                address 10.0.12.13/30;
            }
            family mpls;
        }
    }
    at-0/2/1 {
        atm-options {
            pic-type atm2;
            vpi 0;
        }
        unit 0 {
            bandwidth 50m;
            vci 0.128;
```

```
            family inet {
                address 10.0.12.5/32 {
                    destination 10.0.12.6;
                }
            }
            family mpls;
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.70.143/21;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.1.1/32;
            }
        }
    }
}
routing-options {
    static {
        [...Output truncated...]
    }
    router-id 192.168.1.1;
    autonomous-system 65432;
}
protocols {
    rsvp {
        interface fe-0/1/0.0 {
             link-protection {
                bandwidth 50m;
                max-bypasses 2;
            }
        }
        interface fe-0/1/2.0;
        interface so-0/0/0.0;
        interface at-0/2/1.0;
        interface fxp0.0 {
            disable;
        }
    }
    mpls {
        label-switched-path lsp1 {
            from 192.168.1.1;
            to 192.168.4.1;
             bandwidth 10m;
             link-protection;
            primary path1;
        }
        label-switched-path lsp2 {
            from 192.168.1.1;
            to 192.168.4.1;
             bandwidth 20m;
             link-protection;
            primary path1;
        }
        label-switched-path lsp3 {
```



```
            from 192.168.1.1;
            to 192.168.9.1;
             bandwidth 30m;
             link-protection;
            primary path1;
        }
        label-switched-path lsp4 {
            from 192.168.1.1;
            to 192.168.9.1;
             bandwidth 40m;
             link-protection;
            primary path1;
        }
        path path1 {
            10.0.12.14 strict;
        }
        interface fe-0/1/0.0;
        interface so-0/0/0.0;
        interface at-0/2/1.0;
        interface fxp0.0 {
            disable;
        }
    }
    bgp {
        export send-statics;
        group internal {
            type internal;
            local-address 192.168.1.1;
            neighbor 192.168.2.1;
            neighbor 192.168.4.1;
            neighbor 192.168.9.1;
        }
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface fe-0/1/0.0;
            interface at-0/2/1.0;
            interface so-0/0/0.0;
            interface lo0.0 {
                passive;
            }
        }
    }
}
policy-options {
    policy-statement send-statics {
        term statics {
            from {
                route-filter 100.100.1.0/24 exact;
            }
            then accept;
        }
    }
}
```

Sample Output 2 The following sample output is for transit router R2:

```
user@R2> show configuration | no-more
interfaces {
    so-0/0/0 {
```

```
        unit 0 {
            family inet {
                address 10.0.12.2/30;
            }
            family mpls;
        }
    }
    so-0/0/1 {
        unit 0 {
            family inet {
                address 10.0.24.1/30;
            }
            family mpls;
        }
    }
    fe-0/1/0 {
        unit 0 {
            family inet {
                address 10.0.12.14/30;
            }
            family mpls;
        }
    }
    at-0/2/1 {
        atm-options {
            pic-type atm2;
            vpi 0;
        }
        unit 0 {
            vci 0.128;
            family inet {
                address 10.0.12.6/32 {
                    destination 10.0.12.5;
                }
            }
            family mpls;
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.70.144/21;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.2.1/32;
            }
        }
    }
}
routing-options {
    static {
        [...Output truncated...]
    }
    router-id 192.168.2.1;
    autonomous-system 65432;
}
protocols {
```

```
    rsvp {
        interface so-0/0/1.0;
        interface fe-0/1/0.0;
        interface so-0/0/0.0;
        interface at-0/2/1.0;
        interface fxp0.0;
    }
    mpls {
         interface fe-0/1/0.0;
        interface so-0/0/1.0;
         interface so-0/0/0.0;
         interface at-0/2/1.0;
        interface fxp0.0;
    }
    bgp {
        group internal {
            type internal;
            local-address 192.168.2.1;
            neighbor 192.168.1.1;
            neighbor 192.168.4.1;
            neighbor 192.168.9.1;
        }
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface fe-0/1/0.0;
            interface so-0/0/1.0;
            interface at-0/2/1.0;
            interface so-0/0/0.0;
            interface lo0.0;
        }
    }
}
```

Sample Output 3 The following sample output is for transit/egress router R4:

```
user@R4> show configuration | no-more
[...Output truncated...]
interfaces {
    so-0/0/1 {
        unit 0 {
            family inet {
                address 10.0.24.2/30;
            }
            family mpls;
        }
    }
    so-0/0/3 {
        unit 0 {
            family inet {
                address 10.0.49.1/30;
            }
            family mpls;
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.70.146/21;
            }
```

```
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.4.1/32;
            }
        }
    }
}
routing-options {
    static {
        [...Output truncated...]
    }
    router-id 192.168.4.1;
    autonomous-system 65432;
}
protocols {
    rsvp {
        interface so-0/0/1.0;
        interface so-0/0/3.0;
        interface fxp0.0 {
            disable;
        }
    }
    mpls {
         label-switched-path r4-r1 {
            to 192.168.1.1;
        }
        interface so-0/0/1.0;
        interface so-0/0/3.0;
        interface fxp0.0 {
            disable;
        }
    }
    bgp {
        group internal {
            type internal;
            local-address 192.168.4.1;
            neighbor 192.168.1.1;
            neighbor 192.168.2.1;
            neighbor 192.168.9.1;
        }
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface so-0/0/1.0;
            interface so-0/0/3.0;
            interface lo0.0 {
                passive;
            }
        }
    }
}
```

Sample Output 4 The following sample output is for egress router R9:

```
user@R9> show configuration | no-more
[...Output truncated...]
interfaces {
```

```
    so-0/0/3 {
        unit 0 {
            family inet {
                address 10.0.49.2/30;
            }
            family mpls;
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.69.206/21;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.9.1/32;
            }
        }
    }
}
routing-options {
    static {
        [...Output truncated...]
    }
    router-id 192.168.9.1;
    autonomous-system 65432;
}
protocols {
    rsvp {
        interface so-0/0/3.0;
        interface fxp0.0 {
            disable;
        }
    }
    mpls {
         label-switched-path r9-r1 {
            to 192.168.1.1;
        }
        interface so-0/0/3.0;
        interface fxp0.0 {
            disable;
        }
    }
    bgp {
        group internal {
            type internal;
            local-address 192.168.9.1;
            neighbor 192.168.1.1;
            neighbor 192.168.2.1;
            neighbor 192.168.4.1;
        }
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface so-0/0/3.0;
            interface lo0.0 {
                passive;
```

```
                }
            }
        }
    }
```

CHAPTER 8

Admission Control Errors When Fast Reroute is Configured

This case study describes a network interoperability issue between Juniper Networks routers and another vendor's equipment. When fast reroute is configured, admission control errors appear in the output of the Juniper Networks router. This chapter includes a brief summary of admission control errors, an example network scenario, and commands to troubleshoot and resolve the problem.

The troubleshooting process described in this case study should not be followed rigidly; it is a basis from which you can develop your own process to suit your particular situation.

Admission Control Errors When Fast Reroute is Configured

Problem This checklist provides the steps and commands to troubleshoot a case study about a network interoperability issue between Juniper Networks routers and another vendor's equipment. The checklist includes links to a brief summary of admission control errors, an example network scenario, and more detailed information about the commands used to troubleshoot and resolve the problem. The troubleshooting process described in this case study should not be followed rigidly; it is a basis from which you can develop your own process to suit your particular situation. (See Table 13 on page 167)

Solution

Table 13: Admission Control Errors When Fast Reroute is Configured Checklist

Tasks	Command or Action
"Troubleshooting Fast Reroute Admission Control Errors Overview" on page 168	
"Symptom" on page 169	**show mpls lsp ingress extensive**
"Cause" on page 170	Because of interoperability issues with another vendor's equipment, multiple Path messages from a Juniper Networks T640 routing platform are sent to the other vendor's equipment.

Table 13: Admission Control Errors When Fast Reroute is Configured Checklist *(continued)*

Tasks	Command or Action
"Troubleshooting Commands" on page 171	**show mpls lsp ingress extensive** **show configuration protocols mpls** **monitor start** *filename* **show log** *filename*
"Solution" on page 177	Configure adaptive LSPs.
"Router Configuration" on page 177	**show configuration | no-more**

Troubleshooting Fast Reroute Admission Control Errors Overview

Problem Admission control errors are not generated by Juniper Networks routers. These errors are sent from other vendor's equipment to a Juniper Networks router and appear in the log output of the **show mpls lsp extensive** command.

Admission control occurs on receipt of an RSVP Path message. When a new Path message is considered for admission, the bandwidth requested is compared with the bandwidth available at the priority specified in the **Setup Prio** field. If the requested bandwidth is not available, a PathErr message is returned with an Error Code of 01, admission control failure. See RFC 3209 for more details.

In this case study, the presenting problem is an admission control failure message in the output for the **show mpls lsp extensive** command. After the initial investigation, the available bandwidth is adjusted to accommodate the requested bandwidth. This action does not resolve the problem, and admission control failure messages continue to appear in the output for the **show mpls lsp extensive** command.

Upon further investigation, the admission control failure messages appear only when fast reroute is configured. When fast reroute is removed from the configuration, the admission control errors disappear. Fast reroute protection is required in the network configuration, indicating that removing fast reroute is not a viable solution. The problem is redefined as an interoperability issue and the investigation examines possible causes.

Figure 16 on page 169 illustrates a network topology that is representative of a situation in which interoperability issues cause an admission control error.

Figure 16: Admission Control Error Network

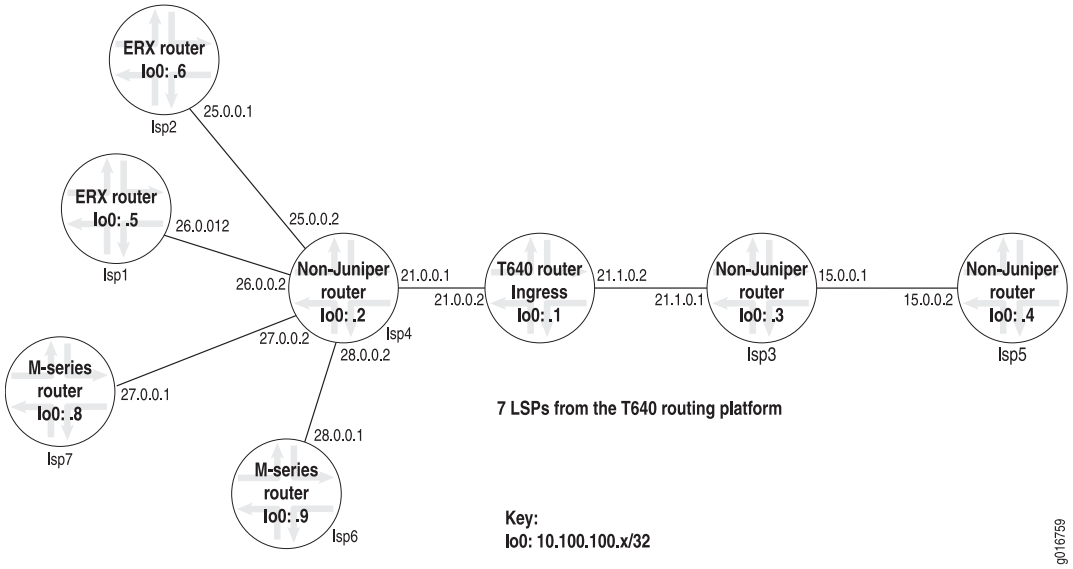

The MPLS network topology in Figure 16 on page 169 shows an Ethernet network of Juniper Networks and non-Juniper Networks equipment that consists of the following components:

- Seven LSPs originating from a T640 routing platform
- Four LSPs terminating at Juniper Networks equipment (**lsp1**, **lsp2**, **lsp6**, and **lsp7**)
- Three LSPs terminating in non-Juniper Networks equipment (**lsp3**, **lsp4**, and **lsp5**)
- All LSPs are transiting non-Juniper Networks equipment
- MPLS trace options is enabled on the T640 routing platform

A sample configuration for the T640 routing platform shown in Figure 16 on page 169 is provided at the end of this case study in “Router Configuration” on page 177.

Symptom In the network shown in Figure 16 on page 169, admission control failure messages appear in the output for the **show mpls lsp ingress extensive** command as shown in the following output.

Sample Output

```
user@T640> show mpls lsp ingress extensive

Ingress LSP:  7 sessions

10.100.100.2
  From: 10.100.100.1, State: Up, ActiveRoute: 0, LSPname: lsp4
  ActivePath: primary (primary)
  FastReroute desired
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   primary            State: Up
    Bandwidth: 10Mbps
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 10)
 21.0.0.1 S
```

```
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

           21.0.0.1
    1515 Aug  5 11:22:50 21.0.0.1:  Admission Control failure [642 times]
    1514 Aug  4 19:46:39 Fast-reroute Detour Up
    1513 Aug  4 19:46:39 21.0.0.1: Admission Control failure
    1512 Aug  4 19:46:39 Up
    1511 Aug  4 19:46:39 Down
    1510 Aug  4 19:46:36 21.0.0.1: Admission Control failure
    1509 Aug  4 19:46:36 Up
    1508 Aug  4 19:46:36 Down
    1507 Aug  4 19:46:30 21.0.0.1: Admission Control failure
    1506 Aug  4 19:46:27 Selected as active path
    1505 Aug  4 19:46:27 Record Route:  21.0.0.1
    1504 Aug  4 19:46:27 Up
[...Output truncated...]
Total 7 displayed,  Up 7, Down 0
```

Meaning The sample output for the **show mpls lsp ingress extensive** command is a snippet that shows one of the problem LSPs (**lsp4**). There are seven ingress LSPs (**7 sessions**) in the **Up** state (**Up 7**), even though at least four of the LSPs have admission control failure messages similar to this one. (See "Troubleshooting Commands" on page 171 for the output for all seven LSPs.) The LSPs with the admission control messages appear to be intermittently coming up and going down (flapping).

Cause The cause of the admission control failure errors appears to be that the other vendor's equipment cannot work with the RSVP Path message sent by the Junos OS. In the Fast Reroute (FRR) object, the Junos OS includes a legacy object and not the standard object. (See RFC 4090 for more information on FRR objects.)

The legacy object has a **flags** field value of **0x00**, which indicates that one-to-one (fast reroute) or facility backup are *not* required. The standard object includes a value of **1** or **2** in the **flags** field depending on the type of protection required. **0x01** indicates one-to-one (fast reroute) backup required, and **0x02** indicates facility backup (many-to-one) backup required.

The Junos OS recognizes both the legacy and standard forms of the fast-reroute object. At the moment, Junos OS sends out only the legacy form which does not have a flags field value (**0x00**). In this case, the **flags** field value should be **0x01** for one-to-one or fast reroute backup. (See Figure 17 on page 171.)

Figure 17: RSVP Duplicate Packets

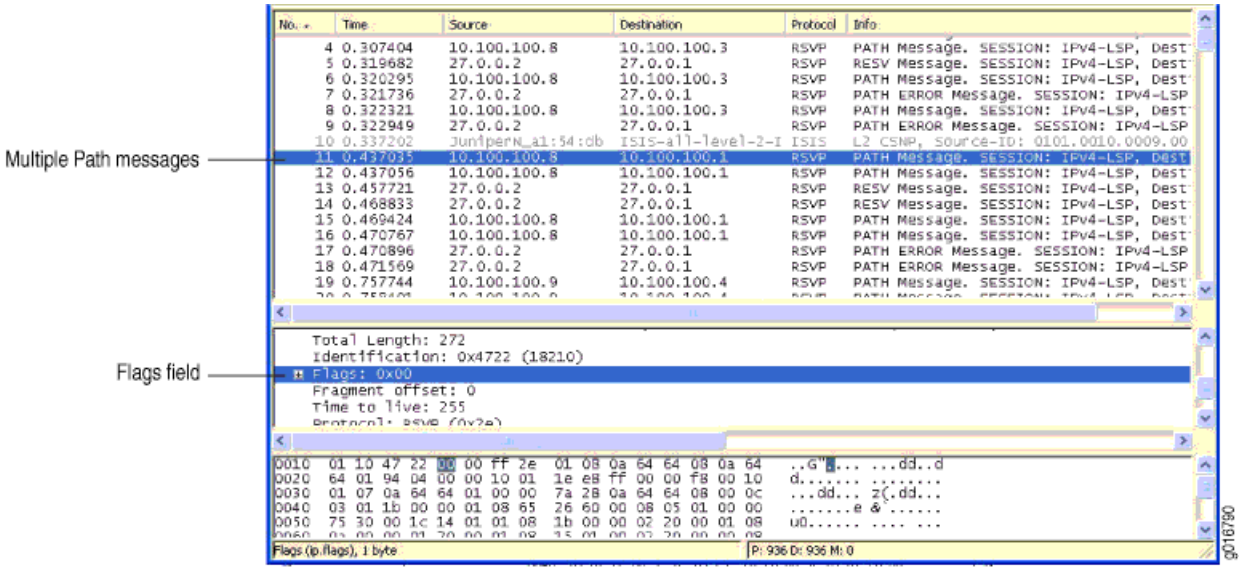

Figure 17 on page 171 shows multiple RSVP Path messages for the same destination with a **flags** field value of **0x00**, indicating that one-to-one or facility backup is *not* required.

Troubleshooting Commands

The Junos OS includes commands that are useful when troubleshooting a problem. This topic provides a brief description of each command, followed by sample output, and a discussion of the output in relation to the problem.

The following commands can be used when troubleshooting an admission control failure problem:

```
user@host> show mpls lsp ingress extensive
user@host> show configuration protocols mpls
user@host> monitor start filename
user@host> show log filename
```

NOTE: Before you use the monitor start and show log commands, you must configure trace options. For directions on configuring trace options for MPLS, see the *Junos MPLS Network Operations Guide*.

Sample Output

Use the show mpls lsp lsp-name ingress extensive command to display detailed information about LSPs configured on the ingress router.

```
user@T640> show mpls lsp ingress extensive

Ingress LSP: 7 sessions

10.100.100.5
  From: 10.100.100.1, State: Up, ActiveRoute: 0,  LSPname: lsp1
  ActivePath: primary (primary)
  LoadBalance: Random
```

```
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   primary            State: Up
    Bandwidth: 10Mbps
    SmartOptimizeTimer: 180
   Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

          21.0.0.1 31.0.0.1
   11 Aug  4 19:46:27 Selected as active path
   10 Aug  4 19:46:27 Record Route:  21.0.0.1 31.0.0.1
    9 Aug  4 19:46:27  Up
[...Output truncated...]

10.100.100.6
  From: 10.100.100.1, State: Up, ActiveRoute: 0,  LSPname: lsp2
  ActivePath: primary (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   primary            State: Up
    Bandwidth: 10Mbps
    SmartOptimizeTimer: 180
   Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

          21.0.0.1 25.0.0.1
   11 Aug  4 19:46:27 Selected as active path
   10 Aug  4 19:46:27 Record Route:  21.0.0.1 25.0.0.1
    9 Aug  4 19:46:27  Up
[...Output truncated...]

10.100.100.3
  From: 10.100.100.1, State: Up, ActiveRoute: 0,  LSPname: lsp3
  ActivePath: primary (primary
  FastReroute desired
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   primary            State: Up
    Bandwidth: 10Mbps
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
 33.0.0.1 S 21.1.0.1 S
   Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

          33.0.0.1(flag=1) 21.1.0.1
   608 Aug  4 19:46:30 21.0.0.1:  Admission Control failure
   607 Aug  4 19:46:27 Selected as active path
   606 Aug  4 19:46:27 Record Route:  21.0.0.1 10.0.0.1
   605 Aug  4 19:46:27 Up
   604 Aug  4 19:46:27 Originate Call
   603 Aug  4 19:46:27 CSPF: computation result accepted
   602 Aug  4 19:46:27 Clear Call
   601 Aug  4 19:46:27 Deselected as active
   600 Aug  3 10:59:13 Fast-reroute Detour Up
   599 Aug  3 10:59:13 Record Route:  33.0.0.1(flag=1) 21.1.0.1
   598 Aug  3 10:58:57 Record Route:  33.0.0.1 21.1.0.1
   597 Aug  3 10:58:57 Fast-reroute Detour Down
   596 Aug  1 11:14:38 Record Route:  33.0.0.1(flag=1) 21.1.0.1
   595 Aug  1 11:14:38 Fast-reroute Detour Up
   594 Aug  1 11:14:18 Record Route:  33.0.0.1 21.1.0.1
   593 Aug  1 11:14:18 Up
   592 Aug  1 11:14:18 Originate make-before-break call
   591 Aug  1 11:14:18 CSPF: computation result accepted
   590 Aug  1 11:14:18 21.0.0.1: Admission Control failure
```

```
   589 Aug  1 11:14:16 Fast-reroute Detour Down
   588 Aug  1 11:14:15 Record Route:  21.0.0.1 10.0.0.1
   587 Aug  1 11:14:15 Up
[...Output truncated...]

10.100.100.2
  From: 10.100.100.1, State: Up, ActiveRoute: 0,  LSPname: lsp4
  ActivePath: primary (primary)
  FastReroute desired
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   primary            State: Up
    Bandwidth: 10Mbps
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 10)
 21.0.0.1 S
   Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

         21.0.0.1
   1515 Aug  5 11:22:50 21.0.0.1:  Admission Control failure [642 times]
   1514 Aug  4 19:46:39  Fast-reroute Detour Up
   1513 Aug  4 19:46:39 21.0.0.1: Admission Control failure
[...Output truncated...]

10.100.100.4
  From: 10.100.100.1, State: Up, ActiveRoute: 0,  LSPname: lsp5
  ActivePath: primary (primary)
  FastReroute desired
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   primary            State: Up
    Bandwidth: 10Mbps
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
 33.0.0.1 S 21.1.0.1 S 15.0.0.2 S
   Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

         33.0.0.1(flag=9) 21.1.0.1(flag=1) 15.0.0.2
   572 Aug  4 19:47:39 Record Route:  33.0.0.1(flag=9) 21.1.0.1(flag=1) 15.0.0.2

   571 Aug  4 19:46:54 Record Route:  33.0.0.1(flag=9) 21.1.0.1 15.0.0.2
   570 Aug  4 19:46:54  Fast-reroute Detour Up
   569 Aug  4 19:46:30 Record Route:  33.0.0.1 21.1.0.1 15.0.0.2
   568 Aug  4 19:46:30 Up
   567 Aug  4 19:46:30 Originate make-before-break call
   566 Aug  4 19:46:30 CSPF: computation result accepted
   565 Aug  4 19:46:30 21.0.0.1:  Admission Control failure
   564 Aug  4 19:46:27 Selected as active path
[...Output truncated...]

10.100.100.9
  From: 10.100.100.1, State: Up, ActiveRoute: 0,  LSPname: lsp6
  ActivePath:  (primary)
  FastReroute desired
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary                   State: Up
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
 21.0.0.1 S 28.0.0.1 S
   Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):
```

```
          21.0.0.1 28.0.0.1
   219152 Aug  5 11:24:10 21.0.0.1:  Admission Control failure
   219151 Aug  5 11:24:10 Up
   219150 Aug  5 11:24:10 Down
[...Output truncated...]

10.100.100.8
  From: 10.100.100.1, State: Up, ActiveRoute: 0,  LSPname: lsp7
  ActivePath: primary (primary)
  FastReroute desired
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   primary            State: Up
    Bandwidth: 10Mbps
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
 21.0.0.1 S 27.0.0.1 S
   Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

          21.0.0.1 27.0.0.1
   71812 Aug  5 11:24:11 21.0.0.1:  Admission Control failure [2 times]
   71811 Aug  5 11:24:11  Fast-reroute Detour Up
   71810 Aug  5 11:24:11 Up
   71809 Aug  5 11:24:11 Down
[...Output truncated...]
Total 7 displayed, Up 7, Down 0
```

Meaning The sample output of the **show mpls lsp ingress extensive** command shows detailed information about the seven ingress LSPs on the T640 platform. All LSPs are up. Five LSPs (**lsp3**, **lsp4**, **lsp5**, **lsp6**, and **lsp7)** have admission control failure messages. Two LSPs (**lsp1** and **lsp2**) do not have admission control failure messages.

All LSPs shown in the network topology in Figure 16 on page 169 transit or terminate on non-Juniper Networks equipment. The question is, why do two LSPs (**lsp1** and **lsp2**) not have admission control errors.

Sample Output Use the show configuration *statement-path* command to display a specific configuration hierarchy; for example, routing protocols.

```
user@T640>
show configuration protocols mpls
traceoptions
 {
    file mpls;
    flag error;
}
label-switched-path  lsp1  {
    to 10.100.100.5;
    no-cspf;
    primary primary {
        bandwidth 10m;
        }
    }
label-switched-path  lsp2  {
    to 10.100.100.6;
    no-cspf;
    primary primary {
```



```
        bandwidth 10m;
        }
    }
label-switched-path  lsp3  {
    to 10.100.100.3;
    fast-reroute;
    primary primary {
        bandwidth 10m;
        }
    }
label-switched-path  lsp4 {
    to 10.100.100.2;
    fast-reroute;
    primary primary {
        bandwidth 10m;
        }
    }
label-switched-path  lsp5  {
    to 10.100.100.4;
    fast-reroute;
    primary primary {
        bandwidth 10m;
        }
    }
label-switched-path  lsp6  {
    to 10.100.100.9;
    no-cspf;
    fast-reroute;
    }
label-switched-path  lsp7  {
    to 10.100.100.8;
    fast-reroute;
    primary primary {
         bandwidth 10m;
         }
    }
path primary;
interface ge-1/0/2.0
interface ge-1/0/4.0
}
```

Meaning The sample output for the **show configuration protocols mpls** command shows the MPLS configuration. Included in the configuration are trace options, seven LSPs, a primary path, and interfaces. Trace options are configured to provide information to assist the investigation of the problem.

The first thing to notice about the MPLS configuration is that the two LSPs (**lsp1** and **lsp2**) do not have the **fast-reroute** statement included. Further investigation shows the following:

- **lsp1** transits non-Juniper Networks equipment, terminates in Juniper Networks equipment, fast reroute is not configured, and there are no admission control failure messages
- **lsp2** transits non-Juniper Networks equipment, terminates in Juniper Networks equipment, fast reroute is not configured, and there are no admission control failure messages

- **lsp3** transits and terminates in non-Juniper Networks equipment, fast reroute is not configured, and there are no admission control failure messages
- **lsp4** terminates in non-Juniper Networks equipment, fast reroute is configured, and there are admission control failure messages
- **lsp5** terminates in non-Juniper Networks equipment, fast reroute is configured, and there are admission control failure messages
- **lsp6** terminates in non-Juniper Networks equipment, fast reroute is configured, and there are admission control failure messages
- **lsp6** transits non-Juniper Networks equipment, terminates in an M-series routing platform, fast reroute is configured, and there are admission control failure messages
- **lsp7** transits non-Juniper Networks equipment, terminates in an M-series routing platform, fast reroute is configured, and there are admission control failure messages

When fast reroute is *not* configured, the LSPs transiting non-Juniper Networks equipment are free of admission control errors. The LSPs with FRR configured have admission control errors. Because all LSPs transit non-Juniper Networks equipment, it would appear that somehow the configuration of fast reroute is an issue for non-Juniper Networks equipment.

Use the show log *filename* command to display the contents of the specified log file. In this case, the log file mpls is configured at the [edit protocols mpls traceoptions] hierarchy level. When the log file is configured, you must issue the monitor start *filename* command to begin logging messages to the file.

Sample Output

```
user@host> monitor start mpls

user@T640> show log /var/log/T640/mpls
Aug  4 19:08:32 trace_on: Tracing to "/var/log/T640/mpls" started
[...Output truncated...]
Aug  4 19:08:32 Receive  PathErr  from 21.0.0.1 (27.0.0.2->0.0.0.0)  Admission
Control failure
Aug  4 19:08:32 mpls lsp  lsp6  primary 21.0.0.1: Admission Control failure[4
times]
Aug  4 19:08:32 task_timer_uset: timer MPLS_MPLS short wait fast <Touched> set
to interval 0.001000 at
Aug  4 19:08:32 task_timer_dispatch: calling MPLS_MPLS short wait fast, late by
0.014
Aug  4 19:08:32 task_timer_reset: reset MPLS_MPLS short wait fast
Aug  4 19:08:32 task_timer_dispatch: returned from MPLS_MPLS short wait fast,
rescheduled in 0
Aug  4 19:08:32 CCC xmit lsp lookup: lsp6 is not a transmit LSP
Aug  4 19:08:32 CCC xmit lsp lookup: lsp6 is not a transmit LSP
Aug  4 19:08:32 mpls lsp lsp6 primary 21.0.0.1: Routing problem,  subcode 0 [2
times]
[...Output truncated...]
```

Meaning The sample output of the **show log *filename*** command is a snippet from the log file that shows the path error (**PathErr**) message for **lsp6** with the admission control failure error and a **0** subcode. Subcode **0** is not one of the error codes (1, 2, or 3) defined in RFC 2205.

NOTE: For readability, some lines in the output that extend beyond 80 characters have been truncated.

Solution The initial solution was to adjust the bandwidth for the LSPs with the admission control failure messages. This approach was not effective because the problem was also an interoperability issue; the other vendor's equipment could not work with the RSVP Path message sent by the Junos OS which included a legacy object and not the standard object in the fast reroute object. For a discussion of the legacy and standard objects, see "Cause" on page 170.

The final solution was to configure all LSPs to be *adaptive*. An adaptive LSP sends out a standard form of the FRR object, and uses the SE RSVP reservation style, which in this case solves the interoperability issue. For information on configuring an adaptive LSP, see "Checklist for RSVP Reservation Styles" on page 51.

Router Configuration Output that shows the configurations of the ingress router in the network. The **no-more** option entered after the pipe **(|)** prevents the output from being paginated if the output is longer than the length of the terminal screen.

Sample Output The following sample output is for ingress router T640:

```
user@T640> show configuration | no-more
[...Output truncated...]
interfaces {
    ge-1/0/2 {
        unit 0 {
            family inet {
                address 21.0.0.2/30;
            }
            family iso;
            family mpls;
        }
    }
    ge-1/0/4 {
        unit 0 {
            family inet {
                address 21.1.0.2/30;
            }
            family iso;
            family mpls;
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 10.100.100.1/32;
            }
            family iso {
                address 01.0000.0101.0010.0001.00;
            }
        }
    }
}
```

```
protocols {
    rsvp {
        interface ge-1/0/2.0;
        interface ge-1/0/4.0;
    }
    mpls {
        traceoptions {
            file mpls;
            flag error;
        }
        label-switched-path lsp1 {
            to 10.100.100.5;
            no-cspf;
            primary primary {
                bandwidth 10m;
            }
        }
        label-switched-path lsp2 {
            to 10.100.100.6;
            no-cspf;
            primary primary {
                bandwidth 10m;
            }
        }
        label-switched-path lsp3 {
            to 10.100.100.3;
            fast-reroute;
            primary primary {
                bandwidth 10m;
            }
        }
        label-switched-path lsp4{
            to 10.100.100.2;
            fast-reroute;
            primary primary {
                bandwidth 10m;
            }
        }
        label-switched-path lsp5{
            to 10.100.100.4;
            fast-reroute;
            primary primary {
                bandwidth 10m;
            }
        }
        label-switched-path lsp6{
            to 10.100.100.9;
            no-cspf;
            fast-reroute;
        }
        label-switched-path lsp7{
            to 10.100.100.8;
            fast-reroute;
            primary primary {
                bandwidth 10m;
            }
        }
path primary;
interface ge-1/0/2.0;
interface ge-1/0/4.0;
}
```

```
    bgp {
        group ibgp {
            type internal;
            local-address 10.100.100.1;
            peer-as 2000;
            neighbor 10.100.100.2;
            neighbor 10.100.100.3;
            neighbor 10.100.100.4;
            neighbor 10.100.100.5;
            neighbor 10.100.100.6;
            neighbor 10.100.100.8;
            neighbor 10.100.100.9;
        }
    }
    isis {
        level 1 disable;
        interface ge-1/0/2.0;
        interface ge-1/0/4.0;
        interface lo0.0 {
            passive;
        }
    }
}
routing-options {
    autonomous-system 2000;
```

Meaning The sample output for the **show configuration** command shows the interfaces, protocols, and routing options configuration for the ingress router (T640) in the network shown in Figure 16 on page 169.

CHAPTER 9

Problem Establishing a GMPLS LSP

This case study describes a problem with establishing a Generalized Multiprotocol Label Switching (GMPLS) label-switched path (LSP). Specifically, the configuration of the data channel is incorrect because the configuration includes different interface types at both ends of the tunnel. The principles and solution used in this case study also apply to control channel configuration.

The chapter includes a brief summary of GRE tunnels within the context of GMPLS, an example network scenario, and commands to troubleshoot and resolve the problem.

The troubleshooting process described in this case study should not be followed rigidly; it is a basis from which you can develop your own process to suit your particular situation.

Problem Establishing a GRE Tunnel Checklist

Problem This checklist provides the links and commands for troubleshooting a case study about a problem establishing a Generalized Multiprotocol Label Switching (GMPLS) label-switched path (LSP). Specifically, the configuration of the data channel is incorrect because the configuration includes different interface types at both ends of the tunnel. The principles and solution used in this case study also apply to control channel configuration.

The checklist includes the links to a brief summary of GRE tunnels within the context of GMPLS, an example network scenario, and more detailed information about the commands used to troubleshoot and resolve the problem.

The troubleshooting process described in this case study should not be followed rigidly; it is a basis from which you can develop your own process to suit your particular situation. (See Table 14 on page 181

Solution

Table 14: Problem Establishing a GRE Tunnel Checklist

Tasks	Command or Action
"Troubleshooting GMPLS and GRE Tunnel" on page 182	
"Symptom" on page 184	**show mpls lsp** **show rsvp session**

Table 14: Problem Establishing a GRE Tunnel Checklist *(continued)*

Tasks	Command or Action
"Cause" on page 185	The cause of the problem with the GMPLS LSP is the configuration of different interface types at both ends of the GMPLS data channel.
"Troubleshooting Commands" on page 185	**show mpls lsp extensive** **show rsvp session detail** **show link-management peer** **show link-management te-link** **show configuration protocols mpls** **monitor start *filename*** **show log *filename***
"Solution" on page 190	Configure both ends of the data channel with the same switching type. **show configuration protocols link-management** **show mpls lsp** **show link-management te-link**
"Conclusion" on page 191	Both ends of a GMPLS data must be the same encapsulation or interface type.
"Router Configurations" on page 191	**show configuration | no-more**

Troubleshooting GMPLS and GRE Tunnel

Problem The logical control channel for GMPLS must be a point-to-point link and must have some form of IP reachability. On broadcast interfaces or when there are multiple hops between control channel peers, use a GRE tunnel for the control channel. For more detailed information on GMPLS and GRE tunnels see the *Junos MPLS Applications Configuration Guide* and the *Junos Feature Guide*.

A tunnel PIC is *not* required to configure a GRE tunnel for the GMPLS control channel. Instead, use the software-based **gre** interface, rather than the hardware-based **gr-*fpc/pic/port*** interface.

CAUTION: Due to restrictions to the software-based gre interface, the GMPLS control channel is the only supported use of the software-based gre interface. Any other use is expressly unsupported and might cause an application failure.

The following example shows a basic **gre** interface configuration. In this case, the tunnel source is the loopback address of the local router and the destination address is the loopback destination of the remote router. Traffic that has a next hop of the tunnel destination will use the tunnel. The tunnel is not automatically used by all the traffic passing through the interface. Only traffic with the tunnel destination as the next hop uses the tunnel.

Sample Output

```
user@R1> show configuration interfaces
[...Output truncated...]
gre {
    unit 0 {
        tunnel {
            source 10.0.12.13;
            destination 10.0.12.14;
        }
        family inet {
            address 10.35.1.6/30;
        }
        family mpls;
    }
}
```

Sample Output The following sample output for the show interfaces command shows the encapsulation type and header, the maximum speed, packets through the logical interface, the destination, and logical address.

```
user@R1> show interfaces gre
Physical interface: gre, Enabled, Physical link is Up
  Interface index: 10, SNMP ifIndex: 8
   Type: GRE, Link-level type: GRE,  MTU: Unlimited, Speed: Unlimited
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
    Input packets : 0
    Output packets: 0

  Logical interface gre.0 (Index 70) (SNMP ifIndex 47)
    Flags: Point-To-Point SNMP-Traps 0x4000
     IP-Header 10.0.12.14:10.0.12.13:47:df:64:0000000000000000
     Encapsulation: GRE-NULL
    Input packets : 171734
    Output packets: 194560
    Protocol inet, MTU: 1476
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
          Destination: 10.35.1.4/30, Local: 10.35.1.6, Broadcast: 10.35.1.7
    Protocol mpls, MTU: 1464
      Flags: None
```

The following are various requirements when you configure a GMPLS LSP using a GRE tunnel:

- The data channel must start and end on the same type of interface.
- The control channel can be a GRE tunnel that starts and ends on the same or different interface type.
- The GRE tunnel must be configured indirectly with the **peer-interface** *peer-name* statement at the **[edit protocol ospf]** hierarchy level.
- The GRE interface must be disabled at the **[edit protocols ospf]** and **[edit protocols rsvp]** hierarchy levels.
- Data and control channels must be defined correctly in the LMP configuration .
- It is optional to disable Constrained Shortest Path First (CSPF) with the **no-cspf** statement.

This case focuses on the incorrect configuration of the endpoints of the GRE tunnel. However, you can use a similar process and commands to diagnose other GRE tunnel problems. Figure 18 on page 184 illustrates a network topology with MPLS tunneled through a GRE interface.

Figure 18: GMPLS Network Topology

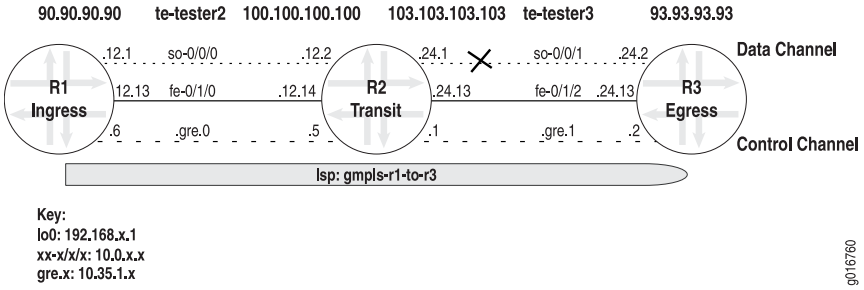

The MPLS network topology in Figure 18 on page 184 shows Juniper Networks routers configured with a GRE tunnel that consists of the following components:

- A strict GMPLS LSP path from the ingress router to the egress router.
- On the ingress router, CSPF disabled with the **no-cspf** statement at the [**edit protocol mpls label-switched-path *lsp-name***] hierarchy level.
- Traffic-engineering links and control channels within the **peer** statement at the [**edit protocols link-management**] hierarchy level on all routers.
- OSPF and OSPF traffic engineering configured on all routers.
- A reference to the **peer-interface** in both OSPF and RSVP on all routers.
- A switching-type problem between **R2** and **R3**.

Symptom The LSP in the network shown in Figure 18 on page 184 is down, as indicated by the output from the **show mpls lsp** and **show rsvp session** commands, which display very similar information. The **show mpls lsp** command shows all LSPs configured on the router, as well as all transit and egress LSPs. The **show rsvp session** command displays summary information about RSVP sessions. You can use either command to verify the state of the LSP. In this case, LSP **gmpls-r1-to-r3** is down (**Dn**).

Sample Output

```
user@R1> show mpls lsp
Ingress LSP: 1 sessions
To              From            State Rt ActivePath       P     LSPname
192.168.4.1   192.168.1.1   Dn   0 -          gmpls-r1-to-r3
Bidir
Total 1 displayed, Up 0, Down 1

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

user@R1>  show rsvp session
Ingress RSVP: 1 sessions
```

```
To              From            State   Rt Style Labelin Labelout LSPname
192.168.4.1   192.168.1.1   Dn    0 0 -    -    - gmpls-r1-to-r3
Bidir
Total 1 displayed, Up 0, Down 1

Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

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

Cause The cause of the problem with the GMPLS LSP is the configuration of different interface types at both ends of the GMPLS data channel.

Troubleshooting Commands The Junos OS includes commands that are useful when troubleshooting a problem. This topic provides a brief description of each command, followed by sample output, and a discussion of the output in relation to the problem.

You can use the following commands when troubleshooting a GMPLS problem:

```
user@host> show mpls lsp extensive
user@host> show rsvp session detail
user@host> show link-management peer
user@host> show link-management te-link
user@host> show configuration protocols mpls
user@host> monitor start filename
user@host> show log filename
```

Sample Output Use the show mpls lsp extensive command on transit router R1 to display detailed information about all LSPs transiting, terminating, and configured on the router.

```
user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

192.168.4.1
  From: 192.168.1.1, State: Dn, ActiveRoute: 0, LSPname: gmpls-r1-to-r3
  Bidirectional
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: SDH/SONET, Switching type: PSC-1, GPID: IPv4
  Primary   p1               State: Dn
    SmartOptimizeTimer: 180
    8 Dec 20 18:08:02 192.168.4.1: MPLS label allocation failure [3 times]
    7 Dec 20 18:07:53 Originate Call
    6 Dec 20 18:07:53 Clear Call
    5 Dec 20 18:07:53 Deselected as active
    4 Dec 20 18:06:13 Selected as active path
    3 Dec 20 18:06:13 Record Route:  100.100.100.100 93.93.93.93
    2 Dec 20 18:06:13 Up
    1 Dec 20 18:06:13 Originate Call
  Created: Wed Dec 20 18:06:12 2006
Total 1 displayed, Up 0, Down 1

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

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

Meaning The sample output for the **show mpls lsp extensive** command shows an error message (**MPLS label allocation failure)** in the log section of the output. This LSP event indicates that the MPLS protocol or the **family mpls** statement were not configured properly. When the LSP event is preceded by an IP address, the address is typically the router that has the MPLS configuration error. In this case, it appears that the router with the **lo0** address of **192.168.4.1** (**R3**) has an MPLS configuration error.

Sample Output Use the show rsvp session detail command to display detailed information about RSVP sessions.

```
user@R1> show rsvp session detail
Ingress RSVP: 1 sessions

192.168.4.1
  From: 192.168.1.1,  LSPstate: Dn, ActiveRoute: 0
   LSPname: gmpls-r1-to-r3, LSPpath: Primary
  Bidirectional, Upstream label in: 21253, Upstream label out: -
  Suggested label received: -,  Suggested label sent: 21253
  Recovery label received: -, Recovery label sent: -
   Resv style: 0 - , Label in: -, Label out: -
  Time left:    -, Since: Wed Dec 20 18:07:53 2006
  Tspec: rate 0bps size 0bps peak 155.52Mbps m 20 M 1500
  Port number: sender 2 receiver 46115 protocol 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 0
  PATH sentto: 10.35.1.5 (tester2) 3 pkts
   Explct route: 100.100.100.100 93.93.93.93
   Record route: <self> ...incomplete
Total 1 displayed, Up 0, Down 1

Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

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

Meaning The sample output for the **show rsvp session detail** command shows that LSP **gmpls-r1-to-r3** is down (**LSPstate: Dn**). The route record is incomplete, indicating a problem with the explicit route **100.100.100.100 93.93.93.93**. The address **100.100.100.100** is the data channel on **R2 so-0/0/0**, and the address **93.93.93.93** is the data channel on **R3**.

Sample Output Use the show link-management peer command to display MPLS peer link information.

```
user@R1> show link-management peer
Peer name: tester2, System identifier: 48428
  State: Up, Control address:  10.35.1.5
   Control-channel                  State
   gre.0                             Active
  TE links:
   tester2

user@R2> show link-management peer
Peer name: tester2, System identifier: 48428
```



```
 State: Up, Control address: 10.35.1.6
   Control-channel                  State
   gre.0                             Active
  TE links:
   te-tester2

Peer name: tester3 , System identifier: 48429
  State: Up , Control address: 10.35.1.2
   Control-channel                  State
   gre.1                             Active
  TE links:
   te-tester3

user@R3> show link-management peer
Peer name: tester3, System identifier: 48429
  State: Up, Control address: 10.35.1.1
   Control-channel                  State
   gre.0                             Active
  TE links:
   te-tester3
```

Meaning The sample output from all routers in the example network in Figure 18 on page 184 for the **show link-management peer** command shows that all control channels are up and active. A detailed analysis of the output shows the following information:

- Name of the peer, **tester2** or **tester3**, which is the same on neighboring routers for ease of troubleshooting.
- Internal identifier for the peer, **48428** for **tester2** and **48429** for **tester3**. The internal identifier is a range of values from 0 through 64,000.
- The state of the peer, which can be up or down. In this case, all peers are up.
- The address to which a control channel is established, for example, **10.35.1.5.**
- The state of the control channel, which can be up, down, or active.
- The traffic-engineered links that are managed by their peer, indicating that control channel **gre.0** is managed by **tester3**.

Sample Output Use the show link-management te-link command to display the resources used to set up Multiprotocol Label Switching (MPLS) traffic-engineered forwarding paths.

```
user@R1> show link-management te-link
TE link name:  tester2, State: Up
  Local identifier: 2005, Remote identifier: 21253, Local address: 90.90.90.90,
Remote address: 100.100.100.100,
  Encoding: SDH/SONET, Switching: PSC-1, Minimum bandwidth: 155.52Mbps, Maximum
bandwidth: 155.52Mbps, Total bandwidth: 155.52Mbps,
  Available bandwidth: 0bps
    Name        State Local ID  Remote ID      Bandwidth  Used   LSP-name
     so-0/0/0  Up   21253   21253      155.52Mbps   Yes   gmpls-r1-to-r3

user@R2> show link-management te-link
 TE link name:  te-tester2, State: Up
  Local identifier: 7002, Remote identifier: 22292, Local address: 100.100.100.100,
 Remote address: 90.90.90.90,
  Encoding: SDH/SONET, Switching: PSC-1, Minimum bandwidth: 155.52Mbps, Maximum
```

```
bandwidth: 155.52Mbps, Total bandwidth: 155.52Mbps,
  Available bandwidth: 0bps
    Name        State Local ID  Remote ID      Bandwidth  Used   LSP-name
    so-0/0/0    Up       21253      21253     155.52Mbps   Yes   gmpls-r1-to-r3
TE link name:  te-tester3, State: Up
  Local identifier: 7003, Remote identifier: 21254, Local address: 103.103.103.103,
 Remote address: 93.93.93.93,
  Encoding: SDH/SONET, Switching: PSC-1, Minimum bandwidth: 155.52Mbps, Maximum
bandwidth: 155.52Mbps, Total bandwidth: 155.52Mbps,
  Available bandwidth: 0bps
    Name        State Local ID  Remote ID      Bandwidth  Used   LSP-name
    so-0/0/1    Up       21252      21252     155.52Mbps   Yes  gmpls-r1-to-r3

user@R3> show link-management te-link
 TE link name:  te-tester3, State: Up
  Local identifier: 7003, Remote identifier: 21254,  Local address: 93.93.93.93,
Remote address: 103.103.103.103,
  Encoding: SDH/SONET, Switching: PSC-1, Minimum bandwidth: 0bps, Maximum
bandwidth: 0bps, Total bandwidth: 0bps,
  Available bandwidth: 0bps
    Name        State Local ID  Remote ID      Bandwidth Used  LSP-name
    so-0/0/1  Dn       21252      21252     155.52Mbps     No
```

Meaning The sample output for the **show link-management te-link** command issued on the three routers in the network in Figure 18 on page 184 shows the resources allocated to the traffic-engineered links **te-tester2** and **te-tester3**. The resources are the SONET interfaces **so-0/0/0** and **so-0/0/1.** On **R1** and **R2, t**he SONET interfaces are used for the LSP **gmpls-r1-to-r3**, as indicated by **Yes** in the **Used** field.However, the SONET interface **so-0/0/1** on **R3** is down (**Dn**) and is not used for the LSP (**Used No**). Further investigation is required to discover why the SONET interface on **R3** is down.

Sample Outut Use the show log *filename* command to display the contents of the specified log file. In this case, the log file rsvp.log is configured at the [edit protocols rsvp traceoptions] hierarchy level. When the log file is configured, you must issue the monitor start *filename* command to begin logging messages to the file.

```
user@R1> show configuration protocols rsvp
traceoptions {
    file rsvp.log size 3m world-readable;
    flag state detail;
    flag error detail;
    flag packets detail;
}

user@R1> monitor start rsvp.log
```

NOTE: The find Error option entered after the pipe (|) searches the output for an instance of the term *Error*.

Sample Output

```
user@R3>
show log rsvp.log | find Error
Dec 28 17:23:32  Error  Len 20 Session preempted flag 0 by 192.168.4.1 TE-link
103.103.103.103
```

```
[...Output truncated...]
Dec 28 17:23:32 RSVP new resv state,session 192.168.4.1(port/tunnel ID 46115
Ext-ID 192.168.1.1)Proto 0
Dec 28 17:23:32         RSVP-LMP reset LMP request for gmpls-r1-to-r3
Dec 28 17:23:32         RSVP->LMP request -  resource for LSP gmpls-r1-to-r3
Dec 28 17:23:32     LMP->RSVP resource request gmpls-r1-to-r3 failed cannot find resource
encoding type SDH/SONET remote label 21252 bandwidth bw[0
Dec 28 17:23:32     RSVP-LMP reset LMP request for gmpls-r1-to-r3
Dec 28 17:23:32 RSVP originate PathErr 192.168.4.1->192.168.2.1 MPLS label allocation failure LSP
 gmpls-r1-to-r3(2/46115)
Dec 28 17:23:32 RSVP send PathErr 192.168.4.1->192.168.2.1 Len=196 tester3
Dec 28 17:23:32   Session7 Len 16 192.168.4.1(port/tunnel ID 46115 Ext-ID
192.168.1.1) Proto 0
Dec 28 17:23:32   Hop      Len 20 192.168.4.1/0x086e4770 TE-link 103.103.103.103
Dec 28 17:23:32   Error    Len 20 MPLS label allocation failure flag 0 by
192.168.4.1 TE-link 103.103.103.103
Dec 28 17:23:32   Sender7  Len 12 192.168.1.1(port/lsp ID  2)
Dec 28 17:23:32   Tspec    Len 36 rate 0bps size 0bps peak 155.52Mbps m 20 M 1500
Dec 28 17:23:32   ADspec   Len 48 MTU 1500
Dec 28 17:23:32   RecRoute Len 20  103.103.103.103 90.90.90.90
Dec 28 17:23:32   SuggLabel Len  8  21252
Dec 28 17:23:32   UpstrLabel Len  8  21252
```

Meaning The sample output from the egress router **R3** for the **show log rsvp.log** command is a snippet taken from the log file. The snippet shows a Link Management Protocol (LMP) resource request for the LSP **gmpls-r1-to-r3.** The request has problems with the encoding type (SDH/SONET), indicating a possible error with the SONET interface connecting **R2** and **R3**. Further investigation of the configuration of the LMP on **R2** and **R3** is required.

Sample Output Use the show configuration *statement-path* command to display a specific configuration hierarchy; in this instance, link-management.

```
user@R2> show configuration protocols link-management
te-link te-tester2 {
    local-address 100.100.100.100;
    remote-address 90.90.90.90;
    remote-id 22292;
    interface so-0/0/0 {
        local-address 100.100.100.100;
        remote-address 90.90.90.90;
        remote-id 21253;
    }
}
te-link te-tester3  {
    local-address 103.103.103.103;
    remote-address 93.93.93.93;
    remote-id 21254;
     interface so-0/0/1  {
        local-address 103.103.103.103;
        remote-address 93.93.93.93;
        remote-id 21252;
    }
}
peer tester2 {
    address 10.35.1.6;
    control-channel gre.0;
    te-link te-tester2;
}
peer tester3 {
```

```
    address 10.35.1.2;
    control-channel gre.1;
    te-link te-tester3;
}

user@R3> show configuration protocols link-management
te-link te-tester3  {
    local-address 93.93.93.93;
    remote-address 103.103.103.103;
    remote-id 21254;
    }
     interface at-0/3/1 {
        local-address 93.93.93.93;
        remote-address 103.103.103.103;
        remote-id 21252;
    }
}
peer tester3 {
    address 10.35.1.1;
    control-channel gre.0;
    te-link te-tester3;
}
```

Meaning The sample output from transit router **R2** and ingress router **R3** for the **show configuration protocols link-management** command shows that the interface type on the two routers is different. The resource allocated to **te-tester3** on transit router **R2** is a SONET interface, while the resource allocated to **te-tester3** on egress router **R3** is an ATM interface. The interface type on each end of the data or control channels must be of the same type. In this case, both ends should be SONET or ATM.

Solution The solution to the problem of different interface or encapsulation types at either end of the GMPLS LSP is to make sure that the interface type is the same at both ends. In this case, the ATM interface was deleted from the link-management configuration on **R3**, and a SONET interface was configured instead.

The following commands illustrate the correct configuration and commands to verify that the GMPLS LSP is up and using the data channel:

```
user@R3> show configuration protocols link-management
user@R3> show mpls lsp
user@R3> show link-management te-link
```

Sample Output

```
user@R3> show configuration protocols link-management
te-link te-tester3 {
    local-address 93.93.93.93;
    remote-address 103.103.103.103;
    remote-id 21254;
     interface so-0/0/1 { # SONET interface replaces the incorrect ATM interface
        local-address 93.93.93.93;
        remote-address 103.103.103.103;
        remote-id 21252;
    }
}
peer tester3 {
    address 10.35.1.1;
    control-channel gre.0;
    te-link te-tester3;
}
```



```
user@R3> show mpls lsp
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
Egress LSP: 1 sessions
To              From            State   Rt Style Labelin Labelout LSPname
192.168.4.1   192.168.1.1   Up    0 1 FF  21252    -  gmpls-r1-to-r3
Bidir
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

user@R3> show link-management te-link
 TE link name: te-tester3, State: Up
  Local identifier: 7003, Remote identifier: 21254, Local address: 93.93.93.93,
Remote address: 103.103.103.103,
  Encoding: SDH/SONET, Switching: PSC-1, Minimum bandwidth: 155.52Mbps, Maximum
bandwidth: 155.52Mbps, Total bandwidth: 155.52Mbps,
  Available bandwidth: 0bps
    Name        State Local ID  Remote ID      Bandwidth Used  LSP-name
     so-0/0/1  Up    21252   21252   155.52Mbps Yes gmpls-r1-to-r3
```

Meaning The sample output for the **show protocols link-management**, **show mpls lsp**, and **show link-management te-link** commands from ingress router **R3** show that the problem is solved. LMP is correctly configured, and the LSP **gmpls-r1-to-r3** is up and using the data channel **so-0/0/1**.

Conclusion In conclusion, both ends of a GMPLS data channel must be the same encapsulation or interface type. This case illustrates the correct configuration of the data channel. The principles are the same for the control channel.

Router Configurations Output that shows the configurations of the ingress router in the network. The **no-more** option entered after the pipe **(|)** prevents the output from being paginated if the output is longer than the length of the terminal screen.

Sample Output The following sample output is for ingress router R1:

```
user@R1> show configuration | no-more
[...Output truncated...]
interfaces {
    so-0/0/0 {
        unit 0 {
            family inet {
                address 10.0.12.1/32 {
                    destination 10.0.12.2;
                }
            }
            family mpls;
        }
    }
    fe-0/1/0 {
        unit 0 {
            family inet {
                address 10.0.12.13/30;
            }
            family mpls;
```

```
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.70.143/21;
            }
        }
    }
    gre {
        unit 0 {
            tunnel {
                source 10.0.12.13;
                destination 10.0.12.14;
            }
            family inet {
                address 10.35.1.6/30;
            }
            family mpls;
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.1.1/32;
            }
        }
    }
}
routing-options {
    static {
        /* corporate and alpha net */
        route 172.16.0.0/12 {
            next-hop 192.168.71.254;
            retain;
            no-readvertise;
        }
        /* old lab nets */
        route 192.168.0.0/16 {
            next-hop 192.168.71.254;
            retain;
            no-readvertise;
        }
        route 0.0.0.0/0 {
            discard;
            retain;
            no-readvertise;
        }
    }
    router-id 192.168.1.1;
    autonomous-system 65432;
}
protocols {
    rsvp {
        traceoptions {
            file rsvp.log size 3m world-readable;
            flag state detail;
            flag error detail;
            flag packets detail;
        }
        interface fxp0.0 {
```

```
                disable;
            }
            interface all;
            interface lo0.0;
            interface gre.0 {
                disable;
            }
            peer-interface tester2;
        }
        mpls {
            label-switched-path gmpls-r1-to-r3 {
                from 192.168.1.1;
                to 192.168.4.1;
                lsp-attributes {
                    switching-type psc-1;
                    encoding-type sonet-sdh;
                }
                no-cspf;
                primary p1;
            }
            path p1 {
                100.100.100.100 strict;
                93.93.93.93 strict;
            }
            interface all;
        }
        ospf {
            traffic-engineering;
            area 0.0.0.0 {
                interface lo0.0;
                interface fe-0/1/0.0;
                interface fxp0.0 {
                    disable;
                }
                interface gre.0 {
                    disable;
                }
                peer-interface tester2;
            }
        }
        link-management {
            te-link tester2 {
                local-address 90.90.90.90;
                remote-address 100.100.100.100;
                remote-id 21253;
                interface so-0/0/0 {
                    local-address 90.90.90.90;
                    remote-address 100.100.100.100;
                    remote-id 21253;
                }
            }
            peer tester2 {
                address 10.35.1.5;
                control-channel gre.0;
                te-link tester2;
            }
        }
    }
```

Sample Output The following sample output is for transit router R2:

```
user@R2>show configuration | no-more
[...Output truncated...]
interfaces {
    so-0/0/0 {
        unit 0 {
            family inet {
                address 10.0.12.2/32 {
                    destination 10.0.12.1;
                }
            }
            family mpls;
        }
    }
    so-0/0/1 {
        unit 0 {
            family inet {
                address 10.0.24.1/32 {
                    destination 10.0.24.2;
                }
            }
            family mpls;
        }
    }
    fe-0/1/0 {
        unit 0 {
            family inet {
                address 10.0.12.14/30;
            }
            family mpls;
        }
    }
    fe-0/1/2 {
        unit 0 {
            family inet {
                address 10.0.24.13/30;
            }
            family mpls;
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.70.144/21;
            }
        }
    }
    gre {
        unit 0 {
            tunnel {
                source 10.0.12.14;
                destination 10.0.12.13;
            }
            family inet {
                address 10.35.1.5/30;
            }
            family mpls;
        }
        unit 1 {
            tunnel {
                source 10.0.24.13;
                destination 10.0.24.14;
```

```
            }
            family inet {
                address 10.35.1.1/30;
            }
            family mpls;
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.2.1/32;
            }
        }
    }
}
routing-options {
    static {
        route 172.16.0.0/12 {
            next-hop 192.168.71.254;
            retain;
            no-readvertise;
        }
        route 192.168.0.0/16 {
            next-hop 192.168.71.254;
            retain;
            no-readvertise;
        }
        route 0.0.0.0/0 {
            discard;
            retain;
            no-readvertise;
        }
    }
    router-id 192.168.2.1;
    autonomous-system 65432;
}
protocols {
    rsvp {
        traceoptions {
            file rsvp.log size 3m world-readable;
            flag packets detail;
            flag state detail;
            flag error detail;
        }
        interface fxp0.0;
        interface lo0.0;
        interface all;
        interface gre.0 {
            disable;
        }
        peer-interface tester2;
        peer-interface tester3;
    }
    mpls {
        interface all;
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface lo0.0;
            interface fxp0.0 {
```

```
                    disable;
                }
                interface gre.0 {
                    disable;
                }
                interface fe-0/1/0.0;
                interface fe-0/1/2.0;
                interface gre.1 {
                    disable;
                }
                peer-interface tester2;
                peer-interface tester3;
            }
        }
        link-management {
            te-link te-tester2 {
                local-address 100.100.100.100;
                remote-address 90.90.90.90;
                remote-id 22292;
                interface so-0/0/0 {
                    local-address 100.100.100.100;
                    remote-address 90.90.90.90;
                    remote-id 21253;
                }
            }
            te-link te-tester3 {
                local-address 103.103.103.103;
                remote-address 93.93.93.93;
                remote-id 21254;
                interface so-0/0/1 {
                    local-address 103.103.103.103;
                    remote-address 93.93.93.93;
                    remote-id 21252;
                }
            }
            peer tester2 {
                address 10.35.1.6;
                control-channel gre.0;
                te-link te-tester2;
            }
            peer tester3 {
                address 10.35.1.2;
                control-channel gre.1;
                te-link te-tester3;
            }
        }
    }
```

Sample Output The following sample output is for egress router R3:

```
user@R3> show configuration | no-more
[...Output truncated...]
interfaces {
    so-0/0/1 {
        unit 0 {
            family inet {
                address 10.0.24.2/32;
            }
            family mpls;
        }
    }
```



```
    fe-0/1/2 {
        unit 0 {
            family inet {
                address 10.0.24.14/30;
            }
            family mpls;
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.70.146/21;
            }
        }
    }
    gre {
        unit 0 {
            tunnel {
                source 10.0.24.14;
                destination 10.0.24.13;
            }
            family inet {
                address 10.35.1.2/30;
            }
            family mpls;
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.4.1/32;
            }
        }
    }
}
routing-options {
    static {
        route 172.16.0.0/12 {
            next-hop 192.168.71.254;
            retain;
            no-readvertise;
        }
        route 192.168.0.0/16 {
            next-hop 192.168.71.254;
            retain;
            no-readvertise;
        }
        route 0.0.0.0/0 {
            discard;
            retain;
            no-readvertise;
        }
    }
    router-id 192.168.4.1;
    autonomous-system 65432;
}
protocols {
    rsvp {
        traceoptions {
            file rsvp.log size 3m world-readable;
            flag packets detail;
```

```
            flag error;
            flag state;
            flag lmp;
        }
        interface fxp0.0 {
            disable;
        }
        interface all;
        interface lo0.0;
        interface gre.0 {
            disable;
        }
        peer-interface tester3;
    }
    mpls {
        interface all;
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface fxp0.0 {
                disable;
            }
            interface fe-0/1/2.0;
            interface gre.0 {
                disable;
            }
            interface lo0.0;
            peer-interface tester3;
        }
    }
    link-management {
        te-link te-tester3 {
            local-address 93.93.93.93;
            remote-address 103.103.103.103;
            remote-id 21254;
            interface so-0/0/1 {
                local-address 93.93.93.93;
                remote-address 103.103.103.103;
                remote-id 21252;
            }
        }
        peer tester3 {
            address 10.35.1.1;
            control-channel gre.0;
            te-link te-tester3;
        }
    }
}
```

PART 3

Index

Index

Symbols

A

B

C

D

E

F

G

H

N

O

P

R

S

T

U