



Junos[®] OS

MPLS Network Operations Guide



Published: 2013-02-25

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

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

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

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

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

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

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

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

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

Junos® OS MPLS Network Operations Guide
Copyright © 2013, Juniper Networks, Inc.
All rights reserved.

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

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Abbreviated Table of Contents

	About This Guide	xiii
Part 1	Monitoring an MPLS Network	
Chapter 1	Configuring MPLS on a Network	3
Chapter 2	Checking the MPLS and RSVP Configuration	45
Chapter 3	Determining the LSP State	59
Chapter 4	Verifying RSVP Signal Processing	69
Chapter 5	Verifying LSP Use	77
Part 2	Working with Problems on Your Network	
Chapter 6	Working with the Layered MPLS Troubleshooting Model	85
Chapter 7	Verifying the Physical Layer	93
Chapter 8	Checking the Data Link Layer	103
Chapter 9	Verifying the IP and IGP Layers	117
Chapter 10	Checking the RSVP Layer	155
Chapter 11	Checking the MPLS Layer	173
Chapter 12	Checking the BGP Layer	193
Part 3	Index	
	Index	211

Table of Contents

	About This Guide	xiii
	Objectives	xiii
	Audience	xiv
	Supported Routing Platforms	xiv
	Using the Index	xiv
	Using the Examples in This Manual	xv
	Merging a Full Example	xv
	Merging a Snippet	xvi
	Document Conventions	xvi
	Documentation Feedback	xviii
	Requesting Technical Support	xviii
	Self-Help Online Tools and Resources	xviii
	Opening a Case with JTAC	xix
Part 1	Monitoring an MPLS Network	
Chapter 1	Configuring MPLS on a Network	3
	Checklist for Configuring and Verifying an MPLS Network	3
	Configuring MPLS on Your Network	6
	Configure IP Addresses on Router Interfaces	8
	Configure IS-IS as the IGP	9
	Enable IS-IS on Routers in Your Network	10
	Configure ISO Addressing	12
	Enable IS-IS on Router Interfaces	13
	Verify That IS-IS Adjacencies Are Established	14
	Configure OSPF as the IGP	15
	Enable OSPF on Routers in Your Network	17
	Verify That OSPF Neighbors Are Established	19
	Set Up BGP on Routers in Your Network	19
	Define the Local Autonomous System	20
	Configure BGP Neighbor Connections	21
	Configure a Simple Routing Policy	23
	Verify That BGP Sessions Are Up	25
	Enable MPLS and RSVP	25
	Enable MPLS and RSVP on Routers	26
	Enable MPLS on Transit Router Interfaces	27
	Establish an LSP in Your Network	28
	Configure the LSP	29
	Verify the LSP	30
	Example Configurations for an MPLS Topology	32

Chapter 2	Checking the MPLS and RSVP Configuration	45
	Checklist for Checking the MPLS and RSVP Configuration	45
	Verifying the MPLS Configuration	46
	Verify MPLS Interfaces	47
	Verify the RSVP Protocol	49
	Verify RSVP Interfaces	50
	Verify Protocol Families	52
	Verify MPLS Labels	55
	Use the traceroute Command to Verify MPLS Labels	55
	Use the ping Command to Verify MPLS Labels	56
Chapter 3	Determining the LSP State	59
	Checklist for Determining LSP Status	59
	Determining LSP Status	59
	Check the Status of the LSP	60
	Display Extensive Status About the LSP	61
	Determining LSP Statistics	65
Chapter 4	Verifying RSVP Signal Processing	69
	Checklist for Verifying RSVP Signal Processing	69
	Checking That RSVP Path Messages Are Sent and Received	70
	Examining the History Log	71
	Determining the Current RSVP Neighbor State	72
	Enabling RSVP Traceoptions	73
Chapter 5	Verifying LSP Use	77
	Checklist for Verifying LSP Use	77
	Verifying LSP Use in Your Network	77
	Verifying an LSP on the Ingress Router	78
	Verifying an LSP on a Transit Router	79
Part 2	Working with Problems on Your Network	
Chapter 6	Working with the Layered MPLS Troubleshooting Model	85
	Checklist for Working with the Layered MPLS Troubleshooting Model	85
	Understanding the Layered MPLS Troubleshooting Model	85
Chapter 7	Verifying the Physical Layer	93
	Checklist for Verifying the Physical Layer	93
	Verifying the Physical Layer	94
	Verify the LSP	95
	Verify Router Connection	98
	Verify Interfaces	98
	Take Appropriate Action	99
	Verify the LSP Again	100

Chapter 8	Checking the Data Link Layer	103
	Checklist for Checking the Data Link Layer	103
	Checking the Data Link Layer	104
	Verify the LSP	105
	Verify Interfaces	106
	Take Appropriate Action	111
	Verify the LSP Again	111
Chapter 9	Verifying the IP and IGP Layers	117
	Checklist for Verifying the IP and IGP Layers	117
	Verifying the IP and IGP Layers	119
	Verifying the IP Layer	121
	Verify the LSP	122
	Verify IP Addressing	123
	Verify Neighbors or Adjacencies at the IP Layer	125
	Take Appropriate Action	128
	Verify the LSP Again	129
	Verifying the OSPF Protocol	132
	Verify the LSP	133
	Verify OSPF Interfaces	136
	Verify OSPF Neighbors	138
	Verify the OSPF Protocol Configuration	138
	Take Appropriate Action	139
	Verify the LSP Again	141
	Verifying the IS-IS Protocol	144
	Verify the LSP	145
	Verify IS-IS Adjacencies and Interfaces	147
	Verify the IS-IS Configuration	149
	Take Appropriate Action	150
	Verify the LSP Again	151
Chapter 10	Checking the RSVP Layer	155
	Checklist for Checking the RSVP Layer	155
	Checking the RSVP Layer	156
	Verify the LSP	158
	Verify RSVP Sessions	160
	Verify RSVP Neighbors	162
	Verify RSVP Interfaces	163
	Verify the RSVP Protocol Configuration	166
	Take Appropriate Action	167
	Verify the LSP Again	168
Chapter 11	Checking the MPLS Layer	173
	Checklist for Checking the MPLS Layer	173
	Checking the MPLS Layer	174
	Verify the LSP	176
	Verify the LSP Route on the Transit Router	179
	Verify the LSP Route on the Ingress Router	180
	Verify MPLS Labels with the traceroute Command	182
	Verify MPLS Labels with the ping Command	183

	Verify the MPLS Configuration	184
	Take Appropriate Action	187
	Verify the LSP Again	188
Chapter 12	Checking the BGP Layer	193
	Checklist for Checking the BGP Layer	193
	Checking the BGP Layer	194
	Check That BGP Traffic Is Using the LSP	195
	Check BGP Sessions	196
	Verify the BGP Configuration	198
	Examine BGP Routes	204
	Verify Received BGP Routes	205
	Take Appropriate Action	206
	Check That BGP Traffic Is Using the LSP Again	207
Part 3	Index	
	Index	211

List of Figures

Part 1	Monitoring an MPLS Network	
Chapter 1	Configuring MPLS on a Network	3
	Figure 1: MPLS Network Topology	7
	Figure 2: IS-IS Network Topology	9
	Figure 3: OSPF Network Topology	16
	Figure 4: BGP Network Topology	20
Chapter 2	Checking the MPLS and RSVP Configuration	45
	Figure 5: MPLS Network Topology	46
Chapter 3	Determining the LSP State	59
	Figure 6: MPLS Network Topology	60
Chapter 5	Verifying LSP Use	77
	Figure 7: MPLS Topology for Verifying LSP Use	78
Part 2	Working with Problems on Your Network	
Chapter 6	Working with the Layered MPLS Troubleshooting Model	85
	Figure 8: Layered MPLS Network Troubleshooting Model	86
	Figure 9: MPLS Basic Network Topology Example	88
Chapter 7	Verifying the Physical Layer	93
	Figure 10: Verifying the Physical Layer	94
	Figure 11: MPLS Network Broken at the Physical Layer	95
Chapter 8	Checking the Data Link Layer	103
	Figure 12: Checking the Data Link Layer	104
	Figure 13: MPLS Network Broken at the Data Link Layer	105
Chapter 9	Verifying the IP and IGP Layers	117
	Figure 14: IP and IGP Layers	120
	Figure 15: MPLS Network Broken at the IP and IGP Layers	121
	Figure 16: MPLS Network Broken at the IP Layer	122
	Figure 17: MPLS Network Broken at the OSPF Protocol Layer	133
	Figure 18: MPLS Network Broken at the IS-IS Protocol Layer	144
Chapter 10	Checking the RSVP Layer	155
	Figure 19: Checking the RSVP Layer	156
	Figure 20: MPLS Network Broken at the RSVP Layer	157
Chapter 11	Checking the MPLS Layer	173
	Figure 21: Checking the MPLS Layer	175

Chapter 12

Figure 22: MPLS Network Broken at the MPLS Layer	175
Checking the BGP Layer	193
Figure 23: Checking the BGP Layer	194
Figure 24: MPLS Network Broken at the BGP Layer	195

List of Tables

	About This Guide	xiii
	Table 1: Notice Icons	xvi
	Table 2: Text and Syntax Conventions	xvii
Part 1	Monitoring an MPLS Network	
Chapter 1	Configuring MPLS on a Network	3
	Table 3: Checklist for Configuring and Verifying an MPLS Network	3
Chapter 2	Checking the MPLS and RSVP Configuration	45
	Table 4: Checklist for Checking the MPLS and RSVP Configuration	45
Chapter 3	Determining the LSP State	59
	Table 5: Checklist for Determining the LSP State	59
Chapter 4	Verifying RSVP Signal Processing	69
	Table 6: Checklist for Verifying RSVP Signal Processing	69
	Table 7: RSVP Tracing Flags	76
Chapter 5	Verifying LSP Use	77
	Table 8: Checklist for Verifying LSP Use	77
	Table 9: MPLS Label Range Allocations	80
Part 2	Working with Problems on Your Network	
Chapter 6	Working with the Layered MPLS Troubleshooting Model	85
	Table 10: Checklist for Working with the Layered MPLS Troubleshooting Model	85
Chapter 7	Verifying the Physical Layer	93
	Table 11: Checklist for Verifying the Physical Layer	93
Chapter 8	Checking the Data Link Layer	103
	Table 12: Checklist for Checking the Data Link Layer	103
Chapter 9	Verifying the IP and IGP Layers	117
	Table 13: Checklist for Verifying the IP and IGP Layers	117
Chapter 10	Checking the RSVP Layer	155
	Table 14: Checklist for Checking the RSVP Layer	155
Chapter 11	Checking the MPLS Layer	173
	Table 15: Checklist for Checking the MPLS Layer	173
Chapter 12	Checking the BGP Layer	193

Table 16: Checklist for Checking the BGP Layer	193
--	-----

About This Guide

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

- [Objectives on page xiii](#)
- [Audience on page xiv](#)
- [Supported Routing Platforms on page xiv](#)
- [Using the Index on page xiv](#)
- [Using the Examples in This Manual on page xv](#)
- [Document Conventions on page xvi](#)
- [Documentation Feedback on page xviii](#)
- [Requesting Technical Support on page xviii](#)

Objectives

This guide provides descriptions of the tasks you need to perform to configure, monitor, and troubleshoot an example MPLS network. 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 Guides*—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
- MX 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. The example does not become active until you commit the 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.xml;
    }
  }
}
interfaces {
  fxp0 {
    disable;
    unit 0 {
      family inet {
        address 10.0.0.1/24;
      }
    }
  }
}
```

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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

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

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

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

For more information about the **load** command, see the 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 xvii defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the configure command: user@host> configure
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> show chassis alarms No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none"> Introduces or emphasizes important new terms. Identifies book names. Identifies RFC and Internet draft titles. 	<ul style="list-style-type: none"> A policy <i>term</i> is a named structure that defines match conditions and actions. <i>Junos OS System Basics Configuration Guide</i> RFC 1997, <i>BGP Communities Attribute</i>
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# set system domain-name <i>domain-name</i>
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none"> To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level. The console port is labeled CONSOLE.
< > (angle brackets)	Enclose optional keywords or variables.	stub <default-metric <i>metric</i> >;
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	broadcast multicast (<i>string1</i> <i>string2</i> <i>string3</i>)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only
[] (square brackets)	Enclose a variable for which you can substitute one or more values.	community name members [<i>community-ids</i>]
Indentation and braces ({ })	Identify a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } }
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	

J-Web GUI Conventions

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

Convention	Description	Examples
Bold text like this	Represents J-Web graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"> In the Logical Interfaces box, select All Interfaces. To cancel the configuration, click Cancel.
> (bold right angle bracket)	Separates levels in a hierarchy of J-Web selections.	In the configuration editor hierarchy, select Protocols>Ospf .

Documentation Feedback

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

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

Requesting Technical Support

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

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

Self-Help Online Tools and Resources

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

- Find CSC offerings: <http://www.juniper.net/customers/support/>
- Find product documentation: <http://www.juniper.net/techpubs/>

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

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

Opening a Case with JTAC

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

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

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

PART 1

Monitoring an MPLS Network

- [Configuring MPLS on a Network on page 3](#)
- [Checking the MPLS and RSVP Configuration on page 45](#)
- [Determining the LSP State on page 59](#)
- [Verifying RSVP Signal Processing on page 69](#)
- [Verifying LSP Use on page 77](#)

CHAPTER 1

Configuring MPLS on a Network

This chapter describes how to configure a network to run Multiprotocol Label Switching (MPLS), including the components and supporting protocols.

- [Checklist for Configuring and Verifying an MPLS Network on page 3](#)
- [Configuring MPLS on Your Network on page 6](#)
- [Configure IP Addresses on Router Interfaces on page 8](#)
- [Configure IS-IS as the IGP on page 9](#)
- [Configure OSPF as the IGP on page 15](#)
- [Set Up BGP on Routers in Your Network on page 19](#)
- [Enable MPLS and RSVP on page 25](#)
- [Establish an LSP in Your Network on page 28](#)
- [Example Configurations for an MPLS Topology on page 32](#)

Checklist for Configuring and Verifying an MPLS Network

This checklist provides the steps and commands to configure and verify an MPLS network. The checklist includes links to more detailed information about the commands to configure and verify MPLS and supporting protocols, and example configurations for an example MPLS topology. (See [Table 3 on page 3.](#))

Table 3: Checklist for Configuring and Verifying an MPLS Network

Tasks	Command or Action
“Configure IP Addresses on Router Interfaces” on page 8	
1. Configure IP Addresses on Router Interfaces on page 8	[edit] edit interfaces <i>type-fpc/pic/port</i> unit <i>logical-unit-number</i> set family inet address <i>address</i> show commit
2. Configure IS-IS as the IGP on page 9	

Table 3: Checklist for Configuring and Verifying an MPLS Network (*continued*)

Tasks	Command or Action
a. Enable IS-IS on Routers in Your Network on page 10	<pre>[edit] edit protocols isis set level 1 disable set interface <i>type-fpc/pic/port</i> level <i>level-number</i> metric <i>metric</i> set interface fxp0.0 disable set interface lo0.0 set interface lo0 passive show commit</pre>
b. Configure ISO Addressing on page 12	<pre>[edit] edit interfaces set lo0 unit <i>number</i> family iso address <i>address</i> show commit</pre>
c. Enable IS-IS on Router Interfaces on page 13	<pre>[edit] edit interfaces set <i>type-fpc/pic/port</i> unit <i>number</i> family iso show commit</pre>
d. Verify That IS-IS Adjacencies Are Established on page 14	<pre>show isis adjacency</pre>
3. Configure OSPF as the IGP on page 15	
a. Enable OSPF on Routers in Your Network on page 17	<pre>[edit] edit protocols ospf [edit protocols ospf] set area <i>area-id</i> interface <i>type-fpc/pic/port</i> set interface fxp0.0 disable set area 0.0.0.0 interface lo0 set area 0.0.0.0 interface lo0 passive set traffic engineering [edit routing-options] set router-id <i>router-id</i> show commit</pre>
b. Verify That OSPF Neighbors Are Established on page 19	<pre>show ospf neighbor</pre>
4. Set Up BGP on Routers in Your Network on page 19	
a. Define the Local Autonomous System on page 20	<pre>[edit] edit routing-options set autonomous-system <i>as-number</i> show commit</pre>

Table 3: Checklist for Configuring and Verifying an MPLS Network (*continued*)

Tasks	Command or Action
b. Configure BGP Neighbor Connections on page 21	<pre>[edit] edit protocols bgp set group <i>group-name</i> type <i>type</i> neighbor <i>neighbor-address</i> set group <i>group-name</i> local-address <i>local-address</i> show commit</pre>
c. Configure a Simple Routing Policy on page 23	<pre>[edit] edit routing-options set static route <i>destination</i> /24 reject [edit policy-options] set policy-statement <i>policy-name</i> term <i>term-name</i> from route-filter <i>address</i> exact set policy-statement <i>policy-name</i> term <i>term-name</i> then accept [edit protocols bgp] set export <i>policy-name</i> show commit</pre>
d. Verify That BGP Sessions Are Up on page 25	<pre>show bgp summary</pre>
5. Enable MPLS and RSVP on page 25	
a. Enable MPLS and RSVP on Routers on page 26	<pre>[edit] edit protocols set mpls interface all set rsvp interface all [edit protocols mpls] set interface fxp0.0 disable [edit protocols rsvp] set interface fxp0.0 disable show commit</pre>
b. Enable MPLS on Transit Router Interfaces on page 27	<pre>[edit] edit interfaces set <i>type-fpc/pic/port</i> unit <i>number</i> family mpls show commit</pre>
6. Establish an LSP in Your Network on page 28	
a. Configure the LSP on page 29	<pre>[edit] edit protocols mpls set label-switched-path <i>lsp-path-name</i> to <i>address</i> show commit</pre>

Table 3: Checklist for Configuring and Verifying an MPLS Network (*continued*)

Tasks	Command or Action
b. Verify the LSP on page 122	<code>show mpls lsp extensive</code>
“Example Configurations for an MPLS Topology” on page 32	<code>show configuration no-more</code>

Configuring MPLS on Your Network

Purpose For MPLS to run on the routers in your network, you must enable MPLS and the Resource Reservation Protocol (RSVP), configure an interior gateway protocol (IGP) and Border Gateway Protocol (BGP) to run over the relevant interfaces, and configure each interface with the following:

- Basic IP information
- MPLS support


In addition, you must configure a label-switched path (LSP) from the ingress router to the egress router. For more information on ingress and egress routers, see the *Junos MPLS Applications Configuration Guide*.

You can configure your MPLS network with either Intermediate System-to-Intermediate System (IS-IS) or Open Shortest Path First (OSPF) as the IGP. The example network in [Figure 1 on page 7](#) is configured with IS-IS. To configure interfaces with OSPF, see the *Junos Routing Protocols Configuration Guide*.

An IGP is required for the Constrained Shortest Path First (CSPF) LSP, which is the default with the Junos OS. The example network in [Figure 1 on page 7](#) focuses on CSPF LSPs.

[Figure 1 on page 7](#) illustrates the example MPLS network topology used in this section and throughout this book. The example network uses IS-IS Level 2 and a policy to create traffic. However, IS-IS Level 1 or an OSPF area can be used and the policy omitted if the network has existing BGP traffic.

Figure 1: MPLS Network Topology



The MPLS network in [Figure 1 on page 7](#) illustrates a router-only network with SONET interfaces that consist 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 R6 that allow a new route to be advertised into the network
- Two unidirectional LSPs between routers R1 and R6, which allow for bidirectional traffic

The network shown in [Figure 1 on page 7](#) 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.

See “[Example Configurations for an MPLS Topology](#)” on [page 32](#) for complete configurations for all routers in this example MPLS network. The following sections outline the steps for configuring MPLS on a network based on the topology shown in [Figure 1 on page 7](#).

You can enable MPLS throughout the rest of the network by repeating Step 1, “[Configure IP Addresses on Router Interfaces](#)” on [page 8](#) through Step 5, “[Enable MPLS and RSVP](#)” on [page 25](#) as appropriate on other routers until all routers and interfaces are enabled for MPLS.

To configure the MPLS network, follow these steps:

1. [Configure IP Addresses on Router Interfaces on page 8](#)
2. [Configure IS-IS as the IGP on page 9](#)
3. [Configure OSPF as the IGP on page 15](#)

4. [Set Up BGP on Routers in Your Network on page 19](#)
5. [Enable MPLS and RSVP on page 25](#)
6. [Establish an LSP in Your Network on page 28](#)

Related Documentation • For more information on configuring MPLS, see the *Junos MPLS Applications Configuration Guide*.

Configure IP Addresses on Router Interfaces

Purpose Before you can run MPLS on your network, you must have an IP address configured on all interfaces. Repeat this procedure as appropriate on other router interfaces in your network until all interfaces have an IP address.

Action To configure an IP address, follow these steps:

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

```
[edit]
user@host# edit interfaces type-fpc/pic/port unit logical-unit-number
```
2. Configure the IP address:

```
[edit interfaces type-fpc/pic/port unit number]
user@host# set family inet address address
```
3. Verify the configuration:

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

Sample Output

```
user@R1> edit
Entering configuration mode

[edit]
user@R1# edit interfaces so-0/0/2 unit 0

[edit interfaces so-0/0/2 unit 0]
user@R1# set family inet address 10.1.13.1/30

[edit interfaces so-0/0/2 unit 0]
user@R1# show
family inet {
    address 10.1.13.1/30 ;
}

[edit interfaces so-0/0/2 unit 0]
user@R1# commit
commit complete
```

Meaning The sample output shows an interface configured with an IP address. The IP address is assigned when you configure the protocol family. In this instance, the IP address is included in the **inet** family. The **family** statement identifies which protocol packets are accepted

into the interfaces. For example, valid IP packets are dropped if the interface is not configured with the **family inet** statement.


Related Documentation

- For more information on interface addressing, see the *Junos Network Interfaces Configuration Guide*.
- [Checklist for Configuring and Verifying an MPLS Network on page 3](#)

Configure IS-IS as the IGP

Before you can run MPLS on your network, you should have an IGP running on all specified routers and interfaces. The IGP can be either IS-IS or OSPF. For the steps to configure OSPF, see “[Configure OSPF as the IGP](#)” on page 15.

Figure 2: IS-IS Network Topology



The IS-IS IGP in the MPLS network in [Figure 2 on page 9](#) consists of the following:

- All routers are configured for Level 2, therefore default CSPF LSPs can occur.
- All routers are in IS-IS area 49.0004. However, the routers in this network could be in any area because Level 2 adjacencies occur between all directly connected Level 2 routers regardless of which area they are in.
- Level 2 adjacencies between all directly connected Level 2 routers as follows:
 - R1 is adjacent to R2, R3, and R5
 - R2 is adjacent to R1, R3, R4, and R6
 - R3 is adjacent to R1, R2, R4, and R6
 - R4 is adjacent to R2, R3, R5, and R6
 - R5 is adjacent to R1, R4, and R6
 - R6 is adjacent to R2, R3, R4, and R5

When you configure IS-IS as the IGP, you must enable IS-IS on the router, configure International Organization for Standardization (ISO) addressing, and enable IS-IS on all router interfaces.

You can enable IS-IS throughout the rest of the network by repeating Step 1, “[Enable IS-IS on Routers in Your Network](#)” on page 10 through Step 3, “[Enable IS-IS on Router Interfaces](#)” on page 13 as appropriate on other routers until all routers and interfaces establish IS-IS adjacencies.

To configure IS-IS and establish IS-IS adjacencies, follow these steps:

1. [Enable IS-IS on Routers in Your Network](#) on page 10
2. [Configure ISO Addressing](#) on page 12
3. [Enable IS-IS on Router Interfaces](#) on page 13
4. [Verify That IS-IS Adjacencies Are Established](#) on page 14

Enable IS-IS on Routers in Your Network

Action To enable IS-IS on routers in your network, follow these steps:

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

```
[edit]  
user@host# edit protocols isis
```
2. Disable Level 1 if appropriate for your network:

```
[edit protocols isis]  
user@host# set level 1 disable
```
3. Configure the interface:

```
[edit protocols isis]  
user@host# edit interface type-fpc/pic/port level level-number metric metric
```
4. Disable the management interface if you have included the **interface all** statement, as shown in [Sample Output 2 on page 11](#):

```
[edit protocols isis]  
user@host# set interface fxp0.0 disable
```
5. Include the loopback interface (lo0) if you have listed all interfaces separately, as shown in [Sample Output 1 on page 11](#):

```
[edit protocols isis]  
user@host# set interface lo0.0
```
6. Set the loopback interface (lo0) to passive:

```
[edit protocols isis]  
user@R1# set interface lo0 passive
```
7. Verify and commit the configuration:

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

Sample Output 1

```

user@R1> edit
Entering configuration mode

[edit]
user@R1# edit protocols isis

[edit protocols isis]
user@R1# set level 1 disable

[edit protocols isis]
user@host# edit interface all level 2 metric 10

[edit protocols isis]
user@host# set interface lo0.0

[edit protocols isis]
user@host# set interface lo0 passive

[edit protocols isis]
user@R1# show
level 1 disable;
interface so-0/0/0.0;
interface so-0/0/1.0;
interface so-0/0/2.0;
interface lo0.0;
    passive;
}

[edit protocols isis]
user@R1# commit
commit complete

```

Sample Output 2

```

[edit protocols isis]
user@R6# show
level 1 disable;
interface all {
    level 2 metric 15;
}
interface fxp0.0 {
    disable;
}
interface lo0.0 {
    passive;
}

```

Meaning Sample Output 1 shows that IS-IS Level 1 is disabled, making this a Level 2 router. All routers in the network shown in [Figure 1 on page 7](#) are running at one IS-IS level (Level 2), therefore default CSPF LSPs can occur.

Because **R1** in Sample Output 1 has all IS-IS enabled interfaces listed, including the loopback interface (**lo0**), you do not need to include the **disable** statement for the management interface (**fxp0**). All interfaces have unit number **0**, the default if a unit number is not specified. When you configure an interface at the **[edit protocols isis]** hierarchy level, and you do not include the logical unit, the default **0** is appended to the interface name, for example, **so-0/0/1.0**.

Sample Output 2 does not list the interfaces configured with IS-IS; instead, all interfaces are configured, including the loopback interface (**lo0**) and the management interface (**fxp0**). Therefore, you do not need to include a separate statement for the loopback (**lo0**) interface. However, in this instance, it is best practice to disable the management interface (**fxp0**) so that IS-IS packets are not sent over it. If you do not disable the management interface (**fxp0**) when you include the **interface-all** statement, the IS-IS protocol can form adjacencies over the management backbone, but traffic does not flow because transit traffic does not go out of the management interface.

Sample Output 2 also shows that all interfaces on **R6** are configured with a metric of **15**. A metric is not required to configure IS-IS on your interfaces. The default metric value is **10** (with the exception of the loopback [**lo0**] interface, which has a default metric of **0**). A metric is included to demonstrate that you can configure a metric for IS-IS if the default (**10**) is not appropriate for your network.

Both sample outputs show the **passive** statement included in the configuration of the loopback (**lo0**) interface. Including the **passive** statement is considered best practice and ensures the following:

- Protocols are not run over the loopback (**lo0**) interface
- When the router ID (RID) is configured manually, ensures that the loopback (**lo0**) interface is advertised to other networks.



.....
NOTE: It is considered best practice to configure the RID manually to avoid duplicate RID problems.
.....

Configure ISO Addressing

Purpose For a router to support IS-IS, you must configure an ISO network entity title (NET) address on one of the router's interfaces, preferably the loopback interface (**lo0**).

Action To configure ISO addressing, follow these steps:

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

```
[edit]  
user@host# edit interfaces
```
2. Include a NET address for the loopback interface:

```
[edit interfaces]  
user@host# set lo0 unit number family iso address address
```
3. Verify and commit the configuration:

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


Sample Output

```

user@R1> edit
Entering configuration mode

[edit]
user@R1# edit interfaces

[edit interfaces]
user@R1# set lo0 unit 0 family iso address 49.0004.1000.0000.0001.00

[edit interfaces]
user@R1# show
[...Output truncated...]
lo0 {
    unit 0 {
        family inet {
            address 10.0.0.1/32;
        }
        family iso {
            address 49.0004.1000.0000.0001.00;
        }
    }
}

[edit interfaces]
user@R1# commit
commit complete

```

Meaning The sample output shows that the loopback (**lo0**) interface is configured with the NET address **49.0004.1000.0000.0001.00**. The loopback interface (**lo0**) becomes a point of connection from the router to the IS-IS network. Every router in an IS-IS network must have at least one ISO NET address that identifies a point of connection to the IS-IS network. The NET address is generally configured on the loopback (**lo0**) interface. Routers that participate in multiple areas can have multiple NET addresses.

All the routers in the network shown in [Figure 1 on page 7](#) share a Level 2 database containing identical information. A common Level 2 database occurs in this case because all adjacencies are Level 2, and all routers are within the same IS-IS area (**49.0004**). Level 2 LSP flooding reaches all routers in the network due to the presence of a single level. For more information on determining the NET address, see the *Junos Routing Protocols Configuration Guide*.

Enable IS-IS on Router Interfaces

Purpose Enable reception and transmission of ISO protocol data units (PDUs) on each router interface in the network with the **family** statement, which identifies which protocol packets are accepted into the interfaces. For example, valid IS-IS packets are dropped if the interface is not configured with the **family iso** statement.

Action To configure support for IS-IS on router interfaces in your network, follow these steps:

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

```

[edit]
user@host# edit interfaces

```

2. Configure IS-IS:

```
[edit interfaces]
user@host# set type-fpc/pic/port unit number family iso
```

3. Verify and commit the configuration:

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

Sample Output

```
user@R1> edit
Entering configuration mode

[edit]
user@R1# edit interfaces

[edit interfaces]
user@R1# set so-0/0/2 unit 0 family iso

[edit interfaces]
userR1# show
[...Output truncated...]
so-0/0/2 {
    unit 0 {
        family inet {
            address 10.1.13.1/30;
        }
        family iso;
    }
}

[edit interfaces]
user@R1# commit
commit complete
```

Meaning The sample output shows that the interface **so-0/0/2** is configured with IS-IS.

Verify That IS-IS Adjacencies Are Established

Purpose After configuring IS-IS, you must verify that neighboring routers have formed adjacencies with each other.

Action To verify IS-IS adjacencies, enter the following Junos OS command-line interface (CLI) operational mode command:

```
user@host> show isis adjacency
```

Sample Output

```
user@R1> show isis adjacency
Interface      System      L State      Hold (secs) SNPA
so-0/0/0.0     R2           2 Up         25
so-0/0/1.0     R5           2 Up         23
so-0/0/2.0     R3           2 Up         20
```

```
user@R3> show isis adjacency
Interface      System      L State      Hold (secs) SNPA
so-0/0/0.0     R4           2 Up         25
so-0/0/1.0     R2           2 Up         25
so-0/0/2.0     R1           2 Up         26
so-0/0/3.0     R6           2 Up         25
```

```
user@R6> show isis adjacency
Interface      System      L State      Hold (secs) SNPA
so-0/0/0.0     R5           2 Up         19
so-0/0/1.0     R4           2 Up         22
so-0/0/2.0     R2           2 Up         22
so-0/0/3.0     R3           2 Up         19
```

Sample Output

```
user@R1> show isis adjacency
Interface      System      L State      Hold (secs) SNPA
so-0/0/0.0     R2           2 Up         25
so-0/0/1.0     R5           2 Up         23
so-0/0/2.0     R3           2 Up         20
```

```
user@R3> show isis adjacency
Interface      System      L State      Hold (secs) SNPA
so-0/0/0.0     R4           2 Up         25
so-0/0/1.0     R2           2 Up         25
so-0/0/2.0     R1           2 Up         26
so-0/0/3.0     R6           2 Up         25
```

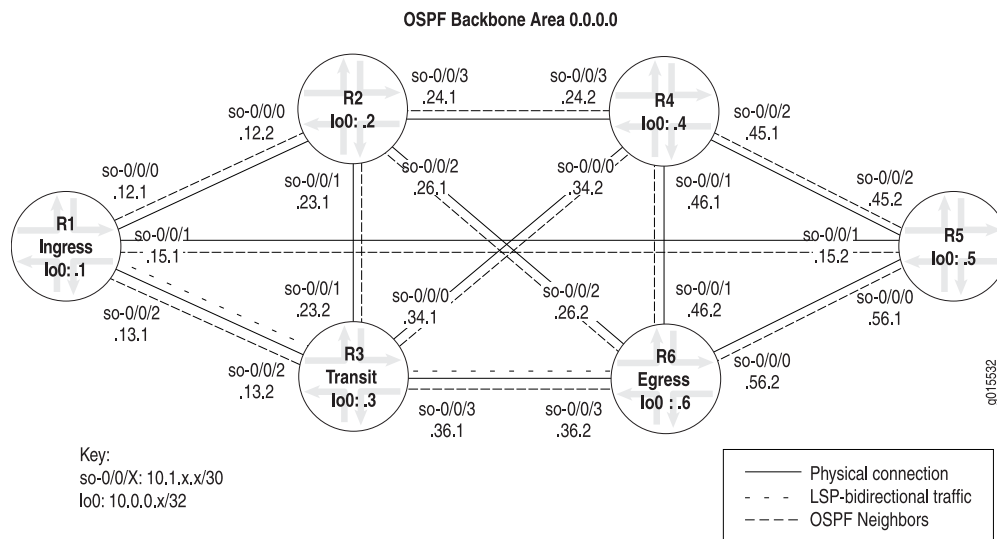
```
user@R6> show isis adjacency
Interface      System      L State      Hold (secs) SNPA
so-0/0/0.0     R5           2 Up         19
so-0/0/1.0     R4           2 Up         22
so-0/0/2.0     R2           2 Up         22
so-0/0/3.0     R3           2 Up         19
```

Meaning The sample output from the ingress, transit, and egress routers shows that all routers in the network shown in [Figure 1 on page 7](#) have formed IS-IS adjacencies.

Configure OSPF as the IGP

Before you can run MPLS on your network, you must have an IGP running on all specified routers and interfaces. The IGP can be either OSPF or IS-IS. For the steps to configure IS-IS, see [“Configure IS-IS as the IGP” on page 9](#).

Figure 3: OSPF Network Topology



The OSPF IGP in the MPLS network in [Figure 2 on page 9](#) consists of the following:

- All routers are configured for the backbone OSPF area 0.0.0.0.
- All routers have the RID manually configured to avoid possible problems when the OSPF RID changes; for example, when multiple loopback addresses are configured.
- All routers have traffic engineering enabled. When traffic engineering is enabled for OSPF, the SPF algorithm takes into account the various LSPs configured under MPLS and configures OSPF to generate link-state advertisements (LSAs) that carry traffic engineering parameters. These routes are installed into the primary routing table **inet.0**, but the LSPs are installed by default into the **inet.3** routing table.
- Adjacencies between all OSPF neighbors are as follows:
 - R1 is adjacent to R2, R3, and R5
 - R2 is adjacent to R1, R3, R4, and R6
 - R3 is adjacent to R1, R2, R4, and R6
 - R4 is adjacent to R2, R3, R5, and R6
 - R5 is adjacent to R1, R4, and R6
 - R6 is adjacent to R2, R3, R4, and R5

When you configure OSPF as the IGP, you must enable OSPF and traffic engineering on the router. We also recommend that you manually configure the RID and include the loopback interface (**lo0**) at the **[edit protocols ospf]** hierarchy level.

You can enable OSPF throughout the rest of the network by repeating this step as appropriate on other routers until all routers and interfaces establish OSPF neighbors.

To configure OSPF and establish OSPF neighbors, follow these steps:

1. [Enable OSPF on Routers in Your Network on page 17](#)
2. [Verify That OSPF Neighbors Are Established on page 19](#)

Enable OSPF on Routers in Your Network

Action To enable OSPF on routers in your MPLS network, follow these steps:

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


```
[edit]
user@host# edit protocols ospf
```
2. Configure the area and the interface:


```
[edit protocols ospf]
user@host# set area area-id interface type-fpc/pic/port
```
3. Disable the management interface if you have included the **interface all** statement in the previous step:


```
[edit protocols ospf]
user@host# set interface fxp0.0 disable
```
4. Include the loopback (lo0) interface if you intend to manually configure the RID:


```
[edit protocols ospf]
user@host# set area 0.0.0.0 interface lo0
```
5. Set the loopback interface (lo0) to passive:


```
[edit protocols ospf]
user@host# set area 0.0.0.0 interface lo0 passive
```
6. Configure traffic engineering:


```
[edit protocols ospf]
user@host# set traffic-engineering
```
7. Manually configure the RID at the [**routing-options**] hierarchy level:


```
[edit]
user@host# edit routing-options
[edit routing-options]
user@host# set router-id router-id
```
8. Verify and commit the entire configuration:


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

Sample Output

```
user@R6> edit
Entering configuration mode

[edit]
user@R6# edit protocols ospf

[edit protocols ospf]
user@R6# set area 0.0.0.0 interface so-0/0/0.0

[edit protocols ospf]
user@R6# set area 0.0.0.0 interface lo0

[edit protocols ospf]
user@R6# set area 0.0.0.0 interface lo0 passive

[edit protocols ospf]
user@R6# set traffic-engineering

[edit protocols ospf]
user@R6# show
traffic-engineering;
area 0.0.0.0 {
    interface so-0/0/0.0;
    interface so-0/0/1.0;
    interface so-0/0/2.0;
    interface so-0/0/3.0;
    interface lo0.0 {
        passive;
    }
}

[edit protocols ospf]
user@R6# commit
commit complete

[edit]
user@R6# edit routing-options

[edit routing-options]
user@R6# set router-id 10.0.0.6

[edit routing-options]
user@R6# show
[...Output truncated...]
router-id 10.0.0.6;
autonomous-system 65432;

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

Meaning The sample output shows that OSPF, with traffic engineering, is enabled on the interfaces on egress router **R6**. In addition, the RID is configured manually to avoid possible problems when the OSPF RID changes; for example, when multiple loopback addresses are configured. The RID uniquely identifies the router within the OSPF network. It is transmitted within the LSAs used to populate the link-state database and calculate the shortest-path tree. In a link-state network, it is important that two routers do not share the same RID value, otherwise IP routing problems may occur.

The sample outputs also shows the **passive** statement included in the configuration of the loopback (**lo0**) interface. Including the **passive** statement is considered best practice and ensures the following:

- Protocols are not run over the loopback (**lo0**) interface
- When the router ID (RID) is configured manually, ensures that the loopback (**lo0**) interface is advertised to other networks.

Verify That OSPF Neighbors Are Established

Purpose After configuring OSPF, you must verify that neighboring routers have formed adjacencies with each other.

Action To verify OSPF neighbors, enter the following Junos OS CLI operational mode command:

```
user@host> show ospf neighbor
```

Sample Output

```
user@R1> show ospf neighbor
  Address      Interface      State      ID           Pri    Dead
  10.1.12.2     so-0/0/0.0    Full      10.0.0.2     128    37
  10.1.15.2     so-0/0/1.0    Full      10.0.0.5     128    35
  10.1.13.2     so-0/0/2.0    Full      10.0.0.3     128    38
```

```
user@R3> show ospf neighbor
  Address      Interface      State      ID           Pri    Dead
  10.1.34.2     so-0/0/0.0    Full      10.0.0.4     128    38
  10.1.23.1     so-0/0/1.0    Full      10.0.0.2     128    35
  10.1.13.1     so-0/0/2.0    Full      10.0.0.1     128    37
  10.1.36.2     so-0/0/3.0    Full      10.0.0.6     128    36
```

```
user@R6> show ospf neighbor
  Address      Interface      State      ID           Pri    Dead
  10.1.56.1     so-0/0/0.0    Full      10.0.0.5     128    39
  10.1.46.1     so-0/0/1.0    Full      10.0.0.4     128    37
  10.1.26.1     so-0/0/2.0    Full      10.0.0.2     128    36
  10.1.36.1     so-0/0/3.0    Full      10.0.0.3     128    37
```

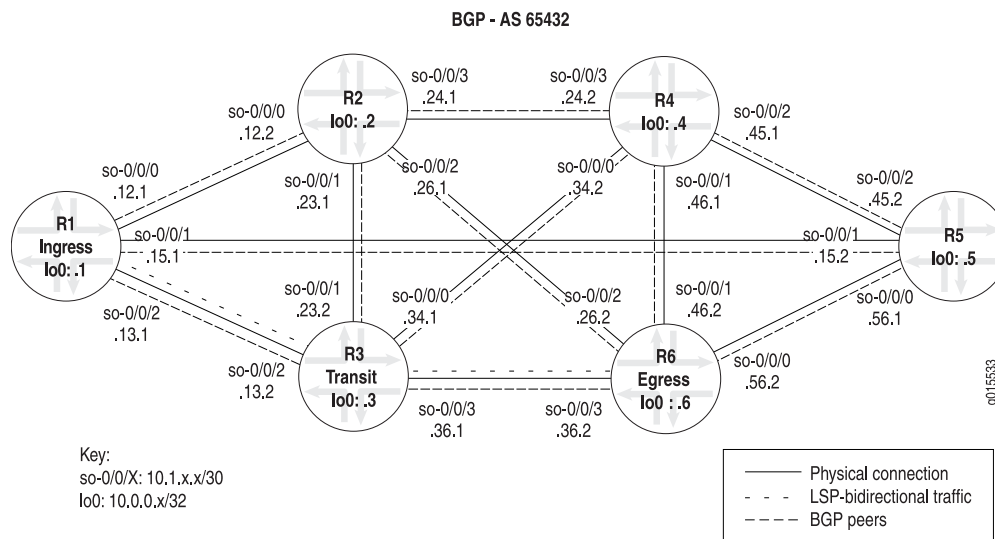
Meaning The sample output from the ingress, transit, and egress routers shows that all routers in the network shown in [Figure 1 on page 7](#) have formed OSPF neighbor adjacencies.

Set Up BGP on Routers in Your Network

Before BGP can function in your MPLS network, you must define the autonomous system (AS) number on the routers in your network, and configure at least one group that includes at least one peer.

Optionally, you can configure a routing policy. The routing policy allows you to control the information shared with BGP neighbors and provides the opportunity to filter and modify the information you receive.

Figure 4: BGP Network Topology



The BGP configuration in the MPLS network in [Figure 4 on page 20](#) consists of the following:

- A full-mesh IBGP topology, using AS 65432.
- All IBGP sessions peer between loopback addresses because significant stability advantages are gained.
- All routers are configured with one group, **group internal**.
- A **send-statics** policy on routers R1 and R6 allows a new route to be advertised into the network.

The example network uses IS-IS Level 2 and a policy to create routes that are reachable through the LSP. However, IS-IS Level 1 or an OSPF area can be used and the policy omitted if the network has existing BGP traffic.

You can set up BGP throughout the rest of the network by repeating Step 1, “[Define the Local Autonomous System](#)” on page 20 through Step 3, “[Configure a Simple Routing Policy](#)” on page 23 as appropriate on other routers until all routers are set up with BGP.

To set up BGP on routers in your network, follow these steps:

1. [Define the Local Autonomous System](#) on page 20
2. [Configure BGP Neighbor Connections](#) on page 21
3. [Configure a Simple Routing Policy](#) on page 23
4. [Verify That BGP Sessions Are Up](#) on page 25

Define the Local Autonomous System

Purpose Before BGP can function, you need to define a local AS number on the routers in your network. In the example network in [Figure 4 on page 20](#), all routers are in AS 65432.

Action To define an AS number on routers in your network, follow these steps:

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

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

2. Configure all interfaces to a specific AS:

```
[edit routing-options]
user@host# set autonomous-system as-number
```

3. Verify the configuration:

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

Sample Output

```
user@R1> edit
Entering configuration mode

[edit]
user@R1# edit routing-options

[edit routing-options]
user@R1# set autonomous-system 65432

[edit routing-options]
user@R1# show
[...Output truncated...]
autonomous-system 65432;

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

Meaning The output shows that router **R1** resides in **AS 65432**. All other routers in the example network shown in [Figure 4 on page 20](#) also reside in **AS 65432**.

Configure BGP Neighbor Connections

Purpose You must configure at least one group that includes at least one peer for BGP to run in your network. First determine which neighbors are internal or external to your local AS boundary. Internal neighbors are inside your local AS boundary. In the example network shown in [Figure 4 on page 20](#), all the routers are in one AS and are therefore internal. In this example, all IBGP sessions peer between loopback addresses because significant stability advantages are gained. For more information about configuring BGP neighbor connections, see the *Junos Routing Protocols Configuration Guide*.

Action To configure BGP neighbor connections, follow these steps:

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

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

2. Configure the group and peer's IP address:

```
[edit protocols bgp]
user@host# set group group-name type type neighbor neighbor-address
```



NOTE: For external neighbors, use the following form of the command that includes the peer's AS number:

```
user@host# set group group-name neighbor neighbor-address peer-as
peer-as-number
```

3. Configure the local address:

```
[edit protocols bgp]
user@host# set group group-name local-address local-address
```

4. Verify and commit the configuration:

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

Sample Output

```
user@R1> edit
Entering configuration mode

[edit]
user@R1# edit protocols bgp

[edit protocols bgp]
user@R1# set group internal type internal neighbor 10.0.0.2

[edit protocols bgp]
users@R1# set group internal local-address 10.0.0.1

[edit protocols bgp]
user@R1# show
group internal {
    type internal;
    local-address 10.0.0.1;
    neighbor 10.0.0.2;
    neighbor 10.0.0.3;
    neighbor 10.0.0.5;
    neighbor 10.0.0.4;
    neighbor 10.0.0.6;
}

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

Meaning The sample output shows that router **R1** is in an internal group with five BGP neighbors. The **local-address** statement is included in this example configuration because IBGP is used. It is considered best practice to configure a local address when you use an IBGP. BGP messages are sourced from the loopback address because the **local-address** statement is included in the configuration. Generally, you would not configure a local address when external BGP is configured.

Configure a Simple Routing Policy

Purpose Routing policy allows you to control the information shared with BGP neighbors and provides the opportunity to filter and modify the information you receive. Typically, a network is injected into BGP using a policy. This may also be done through a static route. In the network in [Figure 4 on page 20](#), a static route export policy is used to inject routes into BGP.

Action To configure a simple routing policy, follow these steps:

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

```
[edit]
user@host# edit routing-options
```
2. Configure a static route for redistribution to other autonomous systems:

```
[edit routing-options]
user@host# set static route destination/24 reject
```
3. Configure a routing policy that matches and accepts the configured static routes into BGP updates:

```
[edit]
user@host# edit policy-options
[edit policy-options]
user@host# set policy-statement policy-name term term-name from route-filter
address exact
user@host# set policy-statement policy-name term term-name then accept
```
4. Apply the policy created in Step 3 to all BGP neighbors:

```
[edit]
user@host# edit protocols bgp
[edit protocols bgp]
user@host# set export policy-name
```
5. Verify and commit the configuration:

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

Sample Output

```
user@R1> edit
Entering configuration mode

[edit]
user@R1# edit routing-options

[edit routing-options]
user@R1# set static route 100.100.1.0/24 reject

[edit routing-options]
user@R1# show
[...Output truncated...]
    route 100.100.1.0/24 reject;
}
router-id 10.0.0.1;
autonomous-system 65432;

[edit routing-options]
user@R1# top

[edit]
user@R1# edit policy-options

[edit policy-options]
user@R1# set policy-statement send-statics term statics from route-filter 100.100.1.0/24 exact

[edit policy-options]
user@R1# set policy-statement send-statics term statics then accept

[edit policy-options]
user@R1# top

[edit]
user@R1# edit protocols bgp

[edit protocols bgp]
user@R1# set export send-statics

[edit protocols bgp]
user@R1# show
export send-statics;
group internal {
    type internal;
    local-address 10.0.0.1;
    neighbor 10.0.0.2;
    neighbor 10.0.0.3;
    neighbor 10.0.0.5;
    neighbor 10.0.0.4;
    neighbor 10.0.0.6;
}

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

Meaning The sample output shows that routing policy **send-statics** is configured on the router. The routing policy matches and accepts the configured static routes into the routing table and injects the routes into BGP updates. Typically, a routing policy is applied at the group level, although it can be applied at the global level, as shown in this example.

Verify That BGP Sessions Are Up

Purpose After configuring BGP, you must verify that BGP peers are established and the sessions are up.

Action To verify BGP peers and sessions, enter the following Junos OS CLI operational mode command:

```
user@host> show bgp summary
```

Sample Output

```
user@R1> show bgp summary
```

```
Groups: 1 Peers: 5 Down peers: 0
```

Table	Tot Paths	Act Paths	Suppressed	History	Damp	State	Pending
inet.0	1	1	0	0	0	0	0
Peer	AS	InPkt	OutPkt	OutQ	Flaps	Last Up/Dwn	State #Active/Received/Damped...
10.0.0.2	65432	1369	1373	0	0	11:25:11	0/0/0 0/0/0
10.0.0.3	65432	1369	1372	0	0	11:24:55	0/0/0 0/0/0
10.0.0.4	65432	1369	1372	0	0	11:25:03	0/0/0 0/0/0
10.0.0.5	65432	1369	1372	0	0	11:25:07	0/0/0 0/0/0
10.0.0.6	65432	1343	1344	0	1	11:10:55	1/1/0 0/0/0

```
user@R3> show bgp summary
```

```
Groups: 1 Peers: 4 Down peers: 0
```

Table	Tot Paths	Act Paths	Suppressed	History	Damp	State	Pending
inet.0	2	2	0	0	0	0	0
Peer	AS	InPkt	OutPkt	OutQ	Flaps	Last Up/Dwn	State #Active/Received/Damped...
10.0.0.1	65432	1375	1375	0	6	11:26:57	1/1/0 0/0/0
10.0.0.2	65432	43016	43016	0	0	2w0d22h	0/0/0 0/0/0
10.0.0.4	65432	74460	74461	0	0	3w4d20h	0/0/0 0/0/0
10.0.0.6	65432	1347	1347	0	6	11:13:10	1/1/0 0/0/0

```
user@R6> show bgp summary
```

```
Groups: 1 Peers: 5 Down peers: 0
```

Table	Tot Paths	Act Paths	Suppressed	History	Damp	State	Pending
inet.0	1	1	0	0	0	0	0
Peer	AS	InPkt	OutPkt	OutQ	Flaps	Last Up/Dwn	State #Active/Received/Damped...
10.0.0.1	65432	1348	1350	0	0	11:13:46	1/1/0 0/0/0
10.0.0.2	65432	1347	1351	0	0	11:14:02	0/0/0 0/0/0
10.0.0.3	65432	1347	1350	0	0	11:13:58	0/0/0 0/0/0
10.0.0.4	65432	1347	1350	0	0	11:13:54	0/0/0 0/0/0
10.0.0.5	65432	1347	1350	0	0	11:13:50	0/0/0 0/0/0

Meaning The sample output from the ingress, transit, and egress routers shows that all routers in the network shown in [Figure 4 on page 20](#) have BGP peers established and sessions up.

Enable MPLS and RSVP

You can enable MPLS and RSVP throughout the rest of the network by repeating Step 1, “[Enable MPLS and RSVP on Routers](#)” on page 26 and Step 2, “[Enable MPLS on Transit Router Interfaces](#)” on page 27 as appropriate on other routers until all routers are enabled with MPLS and RSVP.



NOTE: Even though the MPLS and RSVP protocols are enabled, you must complete all five steps in [“Configuring MPLS on Your Network” on page 6](#) to have the MPLS protocol running on your network.

1. [Enable MPLS and RSVP on Routers on page 26](#)
2. [Enable MPLS on Transit Router Interfaces on page 27](#)

Enable MPLS and RSVP on Routers

Action To enable MPLS and RSVP on routers in your network, follow these steps:

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

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

2. Configure MPLS and RSVP:

```
[edit protocols]
user@host# set mpls interface all
user@host# set rsvp interface all
```

3. Disable the management interface for MPLS and RSVP:

```
[edit protocols mpls]
user@host# set interface fxp0.0 disable
[edit protocols rsvp]
user@host# set interface fxp0.0 disable
```

4. Verify and commit the configuration:

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

Sample Output

```

user@R1> edit
Entering configuration mode

[edit]
user@R1# edit protocols

[edit protocols]
user@R1# set mpls interface all

[edit protocols]
user@R1# set rsdp interface all

[edit protocols]
user@R1# show
rsdp {
    interface all;
    interface fxp0.0 {
        disable;
    }
}
mpls {
    interface all;
    interface fxp0.0 {
        disable;
    }
}

[edit protocols]
user@R1# commit
commit complete

```

Meaning The sample output shows that router **R1** has MPLS and RSVP enabled on all interfaces, except for the management interface (**fxp0.0**), which is disabled. It is considered best practice to disable the management interlace (**fxp0.0**) for MPLS and RSVP to preempt any problems. The sample network shown in [Figure 1 on page 7](#) has all interfaces (with the management interface [**fxp0.0**]) disabled on all routers configured with the MPLS and RSVP protocols.

Typically every interface that you want to use is listed. For an example of a router configured with specific interfaces, see [“Enable IS-IS on Routers in Your Network” on page 10](#).

Enable MPLS on Transit Router Interfaces

Purpose Even though transit interfaces are enabled with MPLS when you include the **family mpls** statement in the configuration, MPLS as a whole is not configured on your router or in your network. You must complete all five steps in this topic to have the MPLS protocol running on your network.



NOTE: The management interface (**fxp0**) and the loopback interface (**lo0**) are not transit interfaces.

Action To configure transit interfaces to support MPLS, follow these steps:

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

```
[edit]
user@host# edit interfaces
```
2. Configure MPLS:

```
[edit interfaces]
user@host# set type-fpc/pic/port unit number family mpls
```
3. Verify and commit the configuration:

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

Sample Output

```
user@R1> edit
Entering configuration mode

[edit]
user@R1# edit interfaces

[edit interfaces]
user@R1# set so-0/0/2 unit 0 family mpls

[edit interfaces]
user@R1# show
so-0/0/2 {
  unit 0 {
    family inet {
      address 10.1.13.1/30;
    }
    family iso;
    family mpls;
  }
}

[edit interfaces]
user@R1# commit
commit complete
```

Meaning The sample output shows that the interface **so-0/0/2** is configured to support MPLS. The **family** statement identifies which protocol packets are accepted into the interfaces. For example, valid MPLS packets are dropped if the interface is not configured with the MPLS protocol.

Establish an LSP in Your Network

Create a label-switched path on specified routers in your network using the loopback address of the ingress and egress routers.

To establish an LSP in your network, follow these steps:

1. [Configure the LSP on page 29](#)
2. [Verify the LSP on page 30](#)

Configure the LSP

Action To configure an LSP in your network, follow these steps:

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

```
[edit]
user@host# edit protocols mpls
```
2. Configure the LSP on the ingress and egress routers:

```
[edit protocols mpls]
user@host# set label-switched-path lsp-path-name to address
```
3. Verify and commit the configuration:

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

Sample Output 1

```
user@R1> edit
Entering configuration mode

[edit]
user@R1# edit protocols mpls

[edit protocols mpls]
user@R1# set label-switched-path R1-to-R6 to 10.0.0.6

[edit protocols mpls]
user@R1# show
label-switched-path R1-to-R6 {
    to 10.0.0.6;
}
interface all;
interface fxp0.0 {
    disable;
}

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

Sample Output 2

```
[edit protocols mpls]
user@R6# show
label-switched-path R6-to-R1 {
    to 10.0.0.1;
}
interface all;
interface fxp0.0 {
    disable;
}
```

Meaning The sample output shows that two CSPF LSPs (**R1-to-R6** and **R6-to-R1**) are configured between routers **R1** and **R6**. CSPF is enabled by default with the Junos OS. The example network shown in [Figure 1 on page 7](#) focuses on CSPF LSPs.

The CSPF algorithm is an advanced form of the SPF algorithm used in OSPF and IS-IS route computations. CSPF is used in computing paths for LSPs that are subject to multiple

constraints. When computing paths for LSPs, CSPF considers not only the topology of the network, but also the attributes of the LSP and the links, and attempts to minimize congestion by intelligently balancing the network load.

Typically in a network, LSPs are configured to every other egress router, resulting in a full mesh of LSPs that correspond to the BGP full mesh. In the example network shown in [Figure 1 on page 7](#), two LSPs are configured between **R1** and **R6** to allow for bidirectional traffic. The first LSP is from **R1** to **R6** (**R1-to-R6**) and the second is from **R6** to **R1** (**R6-to-R1**). If only one LSP was configured, for example, from **R1** to **R6**, only unidirectional traffic would be allowed.

Verify the LSP

Purpose After configuring the LSP, you must verify that the LSP is up. LSPs can be ingress, transit, or egress. Use the **show mpls lsp** command for quick verification of the LSP state, with the **extensive** option (**show mpls lsp extensive**) as a follow-up if the LSP is down. If your network has numerous LSPs, you might consider specifying the name of the LSP, using the **name** option (**show mpls lsp name *name*** or **show mpls lsp name *name* extensive**).

Action To verify that the LSP is up, enter the following command from the ingress router:

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

Sample Output

```

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

10.0.0.6
  From: 10.0.0.1, State: Up, ActiveRoute: 1, LSPname: R1-to-R6
  ActivePath: (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
10.1.13.2 S 10.1.36.2 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

      10.1.13.2 10.1.36.2
    6 Dec 13 11:50:15 Selected as active path
    5 Dec 13 11:50:15 Record Route: 10.1.13.2 10.1.36.2
    4 Dec 13 11:50:15 Up
    3 Dec 13 11:50:15 Originate Call
    2 Dec 13 11:50:15 CSPF: computation result accepted
    1 Dec 13 11:49:45 CSPF failed: no route toward 10.0.0.6[6 times]
  Created: Mon Dec 13 11:47:19 2004
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, LSPstate: Up, ActiveRoute: 0
  LSPname: R6-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: 127, Since: Mon Dec 13 11:50:10 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 39136 protocol 0
  PATH rcvfrom: 10.1.13.2 (so-0/0/2.0) 28709 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.1.36.2 10.1.13.2 <self>
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** show two LSPs in which this router participates: ingress LSP **R1-to-R6** and egress LSP **R6-to-R1** (the reverse LSP which allows bidirectional traffic). Both LSPs have active routes to the destination: **10.0.0.6** for the ingress LSP and **10.0.0.1** for the egress LSP. The state for both LSPs is up.

For more information on verifying the LSP, see [“Checklist for Determining LSP Status” on page 59](#).

Example Configurations for an MPLS Topology

Purpose The configurations in this topic are for the six routers in the example network illustrated in [Figure 1 on page 7](#).

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@R1> show configuration | no-more
system {
  host-name R1;
  [...Output truncated...]
}
interfaces {
  so-0/0/0 {
    unit 0 {
      family inet {
        address 10.1.12.1/30;
      }
      family iso;
      family mpls;
    }
  }
  so-0/0/1 {
    unit 0 {
      family inet {
        address 10.1.15.1/30;
      }
      family iso;
      family mpls;
    }
  }
  so-0/0/2 {
    unit 0 {
      family inet {
        address 10.1.13.1/30;
      }
      family iso;
      family mpls;
    }
  }
  fxp0 {
    unit 0 {
      family inet {
        address 192.168.70.143/21;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.0.0.1/32;
      }
      family iso {
        address 49.0004.1000.0000.0001.00;
      }
    }
  }
}
is
}
routing-options {
  static {
    [...Output truncated...]
    route 100.100.1.0/24 reject;
  }
  router-id 10.0.0.1;
  autonomous-system 65432;
}
protocols {
  rsvp {

```

#family mpls is not
#configured because the
#loopback (lo0) interface

#not a transit interface

```

        inactive: traceoptions {
            file rsvp.log;
            flag packets;
        }
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
    mpls {
        label-switched-path R1-to-R6 {
            to 10.0.0.6;
        }
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
    bgp {
        export send-statics;
        group internal {
            type internal;
            local-address 10.0.0.1;
            neighbor 10.0.0.2;
            neighbor 10.0.0.3;
            neighbor 10.0.0.5;
            neighbor 10.0.0.4;
            neighbor 10.0.0.6;
        }
    }
    isis {
        level 1 disable;
        interface all {
            level 2 metric 10;
        }
        interface fxp0.0 {
            disable;
        }
        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

```

user@R2> show configuration | no-more
system {
    host-name R2;
    [...Output truncated...]
    interfaces {
        so-0/0/0 {
            unit 0 {
                family inet {

```

```

        address 10.1.12.2/30;
    }
    family iso;
    family mpls;
}
}
so-0/0/1 {
    unit 0 {
        family inet {
            address 10.1.23.1/30;
        }
        family iso;
        family mpls;
    }
}
so-0/0/2 {
    unit 0 {
        family inet {
            address 10.1.26.1/30;
        }
        family iso;
        family mpls;
    }
}
so-0/0/3 {
    unit 0 {
        family inet {
            address 10.1.24.1/30;
        }
        family iso;
        family mpls;
    }
}
fxp0 {
    unit 0 {
        family inet {
            address 192.168.70.144/21;
        }
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.0.0.2/32;
        }
        family iso {
            address 49.0004.1000.0000.0002.00;
        }
    }
}
}
routing-options {
    [...Output truncated...]
    router-id 10.0.0.2;
    autonomous-system 65432;
}
protocols {
    rsvp {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
}

```

#family mpls is not
#configured because the
#loopback (lo0) interface is
#not a transit interface

```

    }
  }
  mpls {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  bgp {
    group internal {
      type internal;
      local-address 10.0.0.2;
      neighbor 10.0.0.1;
      neighbor 10.0.0.3;
      neighbor 10.0.0.4;
      neighbor 10.0.0.6;
    }
  }
  isis {
    level 1 disable;
    interface all {
      level 2 metric 10;
    }
    interface fxp0.0 {
      disable;
    }
    interface lo0.0;
    passive
  }
}

```

Sample Output 3 user@R3> show configuration | no-more

```

system {
  host-name R3;
  [...Output truncated...]
}
interfaces {
  so-0/0/0 {
    unit 0 {
      family inet {
        address 10.1.34.1/30;
      }
      family iso;
      family mpls;
    }
  }
  so-0/0/1 {
    unit 0 {
      family inet {
        address 10.1.23.2/30;
      }
      family iso;
      family mpls;
    }
  }
  so-0/0/2 {
    unit 0 {
      family inet {
        address 10.1.13.2/30;
      }
      family iso;
      family mpls;
    }
  }
}

```



```

    }
  }
  so-0/0/3 {
    unit 0 {
      family inet {
        address 10.1.36.1/30;
      }
      family iso;
      family mpls;
    }
  }
  fxp0 {
    unit 0 {
      family inet {
        address 192.168.70.145/21;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.0.0.3/32;
      }
      family iso {
        address 49.0004.1000.0000.0003.00;
      }
    }
  }
}
#family mpls is not
#configured because the
#loopback (lo0) interface is
#not a transit interface

routing-options {
  static {
    [...Output truncated...]
    router-id 10.0.0.3;
    autonomous-system 65432;
  }
  protocols {
    rsvp {
      interface all;
      interface fxp0.0 {
        disable;
      }
    }
    mpls {
      interface all;
      interface fxp0.0 {
        disable;
      }
    }
    bgp {
      group internal {
        type internal;
        local-address 10.0.0.3;
        neighbor 10.0.0.1;
        neighbor 10.0.0.2;
        neighbor 10.0.0.4;
        neighbor 10.0.0.6;
      }
    }
    isis {
      level 1 disable;
    }
  }
}

```

```
        interface all {
            level 2 metric 10;
        }
        interface fxp0.0 {
            disable;
        }
        interface lo0.0;
        passive
    }
}
```

Sample Output 4

```
user@R4> show configuration | no-more
system {
    host-name R4;
    [...Output truncated...]
}
interfaces {
    so-0/0/0 {
        unit 0 {
            family inet {
                address 10.1.34.2/30;
            }
            family iso;
            family mpls;
        }
    }
    so-0/0/1 {
        unit 0 {
            family inet {
                address 10.1.46.1/30;
            }
            family iso;
            family mpls;
        }
    }
    so-0/0/2 {
        unit 0 {
            family inet {
                address 10.1.45.1/30;
            }
            family iso;
            family mpls;
        }
    }
    so-0/0/3 {
        unit 0 {
            family inet {
                address 10.1.24.2/30;
            }
            family iso;
            family mpls;
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.70.146/21;
            }
        }
    }
    lo0 {
        unit 0 {
```

```

        family inet {
            address 10.0.0.4/32;
        }
        family iso {
            address 49.0004.1000.0000.0004.00;
        }
    }
}
#family mpls is not
#configured because the
#loopback (lo0) interface is
#not a transit interface
routing-options {
    static {
        [...Output truncated...]
        router-id 10.0.0.4;
        autonomous-system 65432;
    }
}
protocols {
    rsvp {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
    mpls {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
    bgp {
        group internal {
            type internal;
            local-address 10.0.0.4;
            neighbor 10.0.0.2;
            neighbor 10.0.0.3;
            neighbor 10.0.0.5;
            neighbor 10.0.0.6;
        }
    }
    isis {
        level 1 disable;
        interface all {
            level 2 metric 10;
        }
        interface fxp0.0 {
            disable;
        }
        interface lo0.0;
        passive
    }
}
}

```

Sample Output 5

```

user@R5> show configuration | no-more
system {
    host-name R5;
    [...Output truncated...]
}
interfaces {
    so-0/0/0 {
        unit 0 {
            family inet {
                address 10.1.56.1/30;
            }
        }
    }
}

```

```

        family iso;
        family mpls;
    }
}
so-0/0/1 {
    unit 0 {
        family inet {
            address 10.1.15.2/30;
        }
        family iso;
        family mpls;
    }
}
so-0/0/2 {
    unit 0 {
        family inet {
            address 10.1.45.2/30;
        }
        family iso;
        family mpls;
    }
}
fxp0 {
    unit 0 {
        family inet {
            address 192.168.70.147/21;
        }
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.0.0.5/32;
        }
        family iso {
            address 49.0004.1000.0000.0005.00;
        }
    }
}
#family mpls is not
#configured because the
#loopback (lo0) interface is

}
#not a transit interface
routing-options {
    static {
        [...Output truncated...]
        router-id 10.0.0.5;
        autonomous-system 65432;
    }
}
protocols {
    rsvp {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
    mpls {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
    bgp {

```

```

        group internal {
            type internal;
            local-address 10.0.0.5;
            neighbor 10.0.0.1;
            neighbor 10.0.0.4;
            neighbor 10.0.0.6;
        }
    }
    isis {
        level 1 disable;
        interface all {
            level 2 metric 10;
        }
        interface fxp0.0 {
            disable;
        }
        interface lo0.0;
        passive
    }
}

```

Sample Output 6 user@R6> show configuration | no-more

```

system {
    host-name R6;
    [...Output truncated...]
}
interfaces {
    so-0/0/0 {
        unit 0 {
            family inet {
                address 10.1.56.2/30;
            }
            family iso;
            family mpls;
        }
    }
    so-0/0/1 {
        unit 0 {
            family inet {
                address 10.1.46.2/30;
            }
            family iso;
            family mpls;
        }
    }
    so-0/0/2 {
        unit 0 {
            family inet {
                address 10.1.26.2/30;
            }
            family iso;
            family mpls;
        }
    }
    so-0/0/3 {
        unit 0 {
            family inet {
                address 10.1.36.2/30;
            }
            family iso;
            family mpls;
        }
    }
}

```

```

    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.70.148/21;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 10.0.0.6/32;
            }
            family iso {
                address 49.0004.1000.0000.0006.00;
            }
        }
    }
}
routing-options {
    static {
        [...Output truncated...]
        route 100.100.6.0/24 reject;
    }
    router-id 10.0.0.6;
    autonomous-system 65432;
}
protocols {
    rsvp {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
    mpls {
        label-switched-path R6-to-R1 {
            to 10.0.0.1;
        }
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
    bgp {
        group internal {
            type internal;
            local-address 10.0.0.6;
            export send-statics;
            neighbor 10.0.0.2;
            neighbor 10.0.0.3;
            neighbor 10.0.0.4;
            neighbor 10.0.0.5;
            neighbor 10.0.0.1;
        }
    }
    isis {
        level 1 disable;
        interface all {
            level 2 metric 10;
        }
    }
}

```

#family mpls is not
#configured because the
#loopback (lo0) interface is

#not a transit interface

```

        interface fxp0.0 {
            disable;
        }
        interface lo0.0;
            passive
    }
}
policy-options {
    policy-statement send-statics {
        term statics {
            from {
                route-filter 100.100.6.0/24 exact;
            }
            then accept;
        }
    }
}

```

Meaning Sample Outputs 1 through 6 show the configurations of all six routers in the example network illustrated in [Figure 1 on page 7](#). LSPs **R1-to-R6** and **R6-to-R1** are configured on **R1** and **R6**, respectively.

Two static routes, **100.100.1/24** on **R1** and **100.100.6/24** on **R6**, are configured at the [**edit routing-options static route**] hierarchy level. Both prefixes are included in the send-statics policy at the [**edit policy-options send statics**] hierarchy level so the routes can become BGP routes.

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

CHAPTER 2

Checking the MPLS and RSVP Configuration

This chapter describes how to verify the correct configuration of both the Multiprotocol Label Switching (MPLS) protocol and Resource Reservation Protocol (RSVP). Incorrect configuration of either protocol prevents successful label-switched path (LSP) creation.

- [Checklist for Checking the MPLS and RSVP Configuration on page 45](#)
- [Verifying the MPLS Configuration on page 46](#)
- [Verify MPLS Labels on page 55](#)

Checklist for Checking the MPLS and RSVP Configuration

Purpose This checklist provides the steps and commands for checking the MPLS and RSVP configuration of the Multiprotocol Label Switching (MPLS) network. The checklist provides links to an overview of the MPLS configuration and more detailed information about the commands used to investigate the problem.

[Table 4 on page 45](#) provides commands for checking the MPLS and RSVP configuration.

Table 4: Checklist for Checking the MPLS and RSVP Configuration

Tasks	Command or Action
“Verifying the MPLS Configuration” on page 46	
1. Verify MPLS Interfaces on page 47	<code>show mpls interface</code>
2. Verify the RSVP Protocol on page 49	<code>show rsvp version</code>
3. Verify RSVP Interfaces on page 50	<code>show rsvp interface</code>
4. Verify Protocol Families on page 52	<code>show interfaces terse</code>
5. Verify MPLS Labels on page 55	
a. Use the traceroute Command to Verify MPLS Labels on page 55	<code>traceroute host-name or ip-address-of-remote-host</code>

Table 4: Checklist for Checking the MPLS and RSVP Configuration (*continued*)

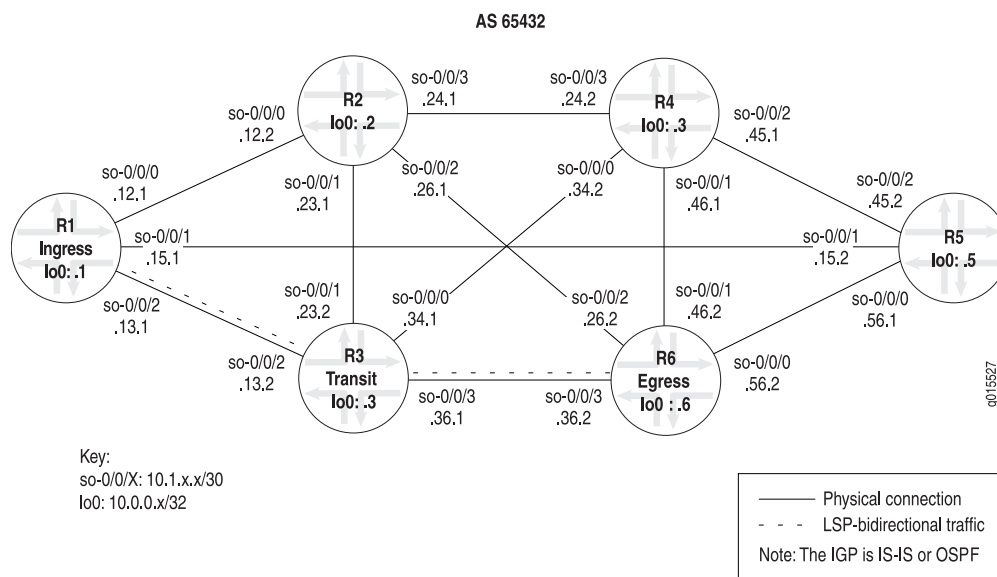
Tasks	Command or Action
b. Use the ping Command to Verify MPLS Labels on page 56	On the egress router, enter the following command: ping mpls rsvp <i>lsp-name</i> detail

Verifying the MPLS Configuration

After configuring MPLS on your network, you must verify the correct configuration of both the MPLS and RSVP protocols. Incorrect configuration of either protocol prevents successful LSP creation.

Figure 5 on page 46 illustrates the network with the example configurations used in this topic. For more details about the router configurations in this network, see “Configuring MPLS on Your Network” on page 6.

Figure 5: MPLS Network Topology



To verify the MPLS configuration, follow these steps:

1. [Verify MPLS Interfaces on page 47](#)
2. [Verify the RSVP Protocol on page 49](#)
3. [Verify RSVP Interfaces on page 50](#)
4. [Verify Protocol Families on page 52](#)

Verify MPLS Interfaces

Purpose If the MPLS protocol is not configured correctly on the routers in your network, the interfaces are not able to perform MPLS switching.

Action To verify MPLS interfaces, enter the following Junos OS command-line interface (CLI) operational mode command:

```
user@host> show mpls interface
```

Sample Output 1

The following sample output is for all routers in the network shown in [Figure 5 on page 46](#).

```
user@R1> show mpls interface
Interface      State      Administrative groups
so-0/0/0.0     Up         <none>
so-0/0/1.0     Up         <none>
so-0/0/2.0     Up         <none>
```

```
user@R2> show mpls interface
Interface      State      Administrative groups
so-0/0/0.0     Up         <none>
so-0/0/1.0     Up         <none>
so-0/0/2.0     Up         <none>
so-0/0/3.0     Up         <none>
```

```
user@R3> show mpls interface
Interface      State      Administrative groups
so-0/0/0.0     Up         <none>
so-0/0/1.0     Up         <none>
so-0/0/2.0     Up         <none>
so-0/0/3.0     Up         <none>
```

```
user@R4> show mpls interface
Interface      State      Administrative groups
so-0/0/0.0     Up         <none>
so-0/0/1.0     Up         <none>
so-0/0/2.0     Up         <none>
so-0/0/3.0     Up         <none>
```

```
user@R5> show mpls interface
Interface      State      Administrative groups
so-0/0/0.0     Up         <none>
so-0/0/1.0     Up         <none>
so-0/0/2.0     Up         <none>
```

```
user@R6> show mpls interface
Interface      State      Administrative groups
so-0/0/0.0     Up         <none>
so-0/0/1.0     Up         <none>
so-0/0/2.0     Up         <none>
so-0/0/3.0     Up         <none>
```

Sample Output 2

```
user@R6> show mpls interface
Interface      State      Administrative groups
so-0/0/0.0     Up         <none>
so-0/0/1.0     Up         <none>
so-0/0/3.0     Up         <none>          # so-0/0/2.0 is missing
```

Sample Output 3

```
user@host> show mpls interface
MPLS not configured
```

Meaning Sample Output 1 shows that all MPLS interfaces on all routers in the network are enabled

(Up) and can perform MPLS switching. If you fail to configure the correct interface at the [edit protocols mpls] hierarchy level or include the **family mpls** statement at the [edit interfaces *type-fpc/pic/port unit number*] hierarchy level, the interface cannot perform MPLS switching, and does not appear in the output for the **show mpls interface** command.

Administrative groups are not configured on any of the interfaces shown in the example network in [Figure 5 on page 46](#). However, if they were, the output would indicate which affinity class bits are enabled on the router.

Sample Output 2 shows that interface **so-0/0/2.0** is missing and therefore might be incorrectly configured. For example, the interface might not be included at the [edit protocols mpls] hierarchy level, or the **family mpls** statement might not be included at the [edit interfaces *type-fpc/pic/port unit number*] hierarchy level. If the interface is configured correctly, RSVP might not have signaled over this interface yet. For more information on determining which interface is incorrectly configured, see [“Verify Protocol Families” on page 52](#).

Sample Output 3 shows that the MPLS protocol is not configured at the [edit protocols mpls] hierarchy level.

For more information on configuring MPLS on routers in your network, see [“Configuring MPLS on Your Network” on page 6](#)

Verify the RSVP Protocol

Purpose If the RSVP protocol is not enabled on the routers in your network, the interface cannot signal LSPs.

Action To verify that the RSVP protocol is enabled, enter the following Junos OS CLI command:

```
user@host> show rsvp version
```

Sample Output

```
user@R1> show rsvp version
Resource ReSerVation Protocol, version 1. rfc2205
  RSVP protocol           = Enabled
  R(refresh timer)        = 30 seconds
  K(keep multiplier)       = 3
  Preemption              = Normal
  Soft-preemption cleanup = 30 seconds
  Graceful restart         = Disabled
  Restart helper mode      = Enabled
  Restart time             = 0 msec
```

Meaning The sample output shows that the RSVP protocol is enabled on **R1**. The supported RSVP protocol is version 1, as defined in RFC 2205.

The RSVP refresh timer is set to 30 seconds, indicating that every 30 seconds, plus or minus 50 percent, the router will refresh the RSVP state with its directly connected neighbors by sending either a **Path** or a **Resv** message. The variable refresh time helps prevent harmonic oscillations in network traffic caused by periodic protocol updates.

The keepalive multiplier, **K(keep multiplier)**, is input to a formula that helps determine the lifetime of an RSVP session. The session lifetime is reset each time the state is updated. The lifetime represents the duration of an RSVP session that does not receive any state updates (**Path** or **Resv** messages). The formula is:

$$\text{RSVP session lifetime} = (\text{keep-multiplier} + 0.5) * 1.5 * \text{refresh-time}$$

The RSVP **preemption** state is currently configured for normal preemption, indicating that only an LSP with a stronger priority can preempt an existing session; that is, the setup value of the new LSP is lower than the hold value of the existing LSP. Other options include **aggressive** preemption, which always preempts when there is insufficient bandwidth, and **disabled**, which prevents any preemption, regardless of LSP priority values.

Graceful restart is currently disabled and **Restart helper mode** is enabled. There are four combinations for **Graceful restart** and **restart helper mode**:

1. Both **Graceful restart** and **Restart helper mode** are enabled.
2. **Graceful restart** is enabled but **Restart helper mode** is disabled. An LSR with this configuration can restart gracefully but cannot help a neighbor with its restart and recovery procedures.
3. **Graceful restart** is disabled but **Restart helper mode** is enabled. An LSR with this configuration can only help a restarting neighbor. It cannot restart gracefully itself.
4. **Graceful restart** and **Restart helper mode** are both disabled. This configuration completely disables RSVP graceful restart (including restart and recovery procedures and helper mode). It is the same as an LSR that is not supported by RSVP graceful restart.

Restart time is the estimated time (in milliseconds) for an LSR to restart the RSVP traffic engineering component. In the example output, the restart time is 0 milliseconds, indicating that it is disabled.

The output is identical for all routers in the network shown in [Figure 5 on page 46](#).

Verify RSVP Interfaces

Purpose If the RSVP protocol is not configured correctly on the routers in your network, the interfaces cannot signal LSPs.

Action To verify RSVP interfaces, enter the following Junos OS CLI operational mode command:

```
user@host> show rsvp interface
```

Sample Output 1

```
user@R1> show rsvp interface
```

```
RSVP interface: 4 active
```

Interface	State	Active resv	Subscr- ption	Static BW	Available BW	Reserved BW	Highwater mark
so-0/0/0.0	Up	2	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/1.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/2.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps

```
user@R2> show rsvp interface
```

```
RSVP interface: 5 active
```

Interface	State	Active resv	Subscr- ption	Static BW	Available BW	Reserved BW	Highwater mark
so-0/0/0.0	Up	1	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/1.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/2.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/3.0	Up	1	100%	155.52Mbps	155.52Mbps	0bps	0bps

```
user@R3> show rsvp interface
```

```
RSVP interface: 5 active
```

Interface	State	Active resv	Subscr- ption	Static BW	Available BW	Reserved BW	Highwater mark
so-0/0/0.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/1.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/2.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/3.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps

```
user@R4> show rsvp interface
```

```
RSVP interface: 5 active
```

Interface	State	Active resv	Subscr- ption	Static BW	Available BW	Reserved BW	Highwater mark
so-0/0/0.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/1.0	Up	1	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/2.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/3.0	Up	1	100%	155.52Mbps	155.52Mbps	0bps	0bps

```
user@R5> show rsvp interface
```

```
RSVP interface: 4 active
```

Interface	State	Active resv	Subscr- ption	Static BW	Available BW	Reserved BW	Highwater mark
so-0/0/0.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/1.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/2.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps

```
user@R6> show rsvp interface
```

```
RSVP interface: 5 active
```

Interface	State	Active resv	Subscr- ption	Static BW	Available BW	Reserved BW	Highwater mark
so-0/0/0.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/1.0	Up	1	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/2.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/3.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps

Sample Output 2

```
user@R6> show rsvp interface
```

```
RSVP interface: 3 active
```

Active	Subscr-	Static	Available	Reserved	Highwater
--------	---------	--------	-----------	----------	-----------

Interface	State	resv	ption	BW	BW	BW	mark
so-0/0/0.0	Up	1	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/1.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/2.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps

#so-0/0/3.0 is missing

Sample Output 3

```
user@host# show rsvp interface
RSVP not configured
```

Meaning Sample Output 1 shows that all interfaces on all routers in the network are enabled with RSVP, including the management interface (**fxp0**). The output for all routers in the network includes similar information, so we will examine R6 in detail.

R6 has five interfaces enabled with RSVP (**Up**). Interface **so-0/1/1.0** has a single active RSVP reservation (**Active resv**) that did not change the default subscription percentage of 100 percent (**Subscription**). Interface **so-0/1/1.0** did not assign a static bandwidth (**Static BW**) to the logical unit and therefore inherited 100 percent of the physical interface rate as the bandwidth available (**Available BW**) for RSVP sessions. Interface **so-0/1/1.0** has no bandwidth assigned (**Reserved BW**), and no RSVP bandwidth allocation at any single instant in time (**Highwater mark**).

Sample Output 2 shows that interface **so-0/0/3.0** is missing. If you do not configure the correct interface at the **[edit protocols rsvp]** hierarchy level, the interface cannot signal LSPs, and does not appear in the output for the **show rsvp interface** command. For more information on configuring MPLS on routers in your network, see [“Configuring MPLS on Your Network” on page 6](#).

Sample Output 3 shows that the RSVP protocol is not configured at the **[edit protocols rsvp]** hierarchy level.

Verify Protocol Families

Purpose If a logical interface does not have MPLS enabled, it cannot perform MPLS switching. This step allows you to quickly determine which interfaces are configured with MPLS and other protocol families.

Action To verify the protocol families configured on the routers in your network, enter the following Junos OS CLI operational mode command:

```
user@host> show interfaces terse
```


Sample Output 1

```

user@R1> show interfaces terse
Interface      Admin Link Proto Local Remote
so-0/0/0       up   up   up   10.1.12.1/30
so-0/0/0.0     up   up   inet 10.1.12.1/30
               up   up   iso
               up   up   mpls

so-0/0/1       up   up
so-0/0/1.0     up   up   inet 10.1.15.1/30
               up   up   iso
               up   up   mpls

so-0/0/2       up   up
so-0/0/2.0     up   up   inet 10.1.13.1/30
               up   up   iso
               up   up   mpls

so-0/0/3       up   down

user@R2> show interfaces terse
Interface      Admin Link Proto Local Remote
so-0/0/0       up   up
so-0/0/0.0     up   up   inet 10.1.12.2/30
               up   up   iso
               up   up   mpls

so-0/0/1       up   up
so-0/0/1.0     up   up   inet 10.1.23.1/30
               up   up   iso
               up   up   mpls

so-0/0/2       up   up
so-0/0/2.0     up   up   inet 10.1.26.1/30
               up   up   iso
               up   up   mpls

so-0/0/3       up   up
so-0/0/3.0     up   up   inet 10.1.24.1/30
               up   up   iso
               up   up   mpls

user@R3> show interfaces terse
Interface      Admin Link Proto Local Remote
so-0/0/0       up   up
so-0/0/0.0     up   up   inet 10.1.34.1/30
               up   up   iso
               up   up   mpls

so-0/0/1       up   up
so-0/0/1.0     up   up   inet 10.1.23.2/30
               up   up   iso
               up   up   mpls

so-0/0/2       up   up
so-0/0/2.0     up   up   inet 10.1.13.2/30
               up   up   iso
               up   up   mpls

so-0/0/3       up   up
so-0/0/3.0     up   up   inet 10.1.36.1/30
               up   up   iso
               up   up   mpls

user@R4> show interfaces terse
Interface      Admin Link Proto Local Remote
so-0/0/0       up   up
so-0/0/0.0     up   up   inet 10.1.34.2/30

```

```

iso
mpls
so-0/0/1          up    up
so-0/0/1.0        up    up    inet  10.1.46.1/30
iso
mpls
so-0/0/2          up    up
so-0/0/2.0        up    up    inet  10.1.45.1/30
iso
mpls
so-0/0/3          up    up
so-0/0/3.0        up    up    inet  10.1.24.2/30
iso
mpls

user@R5> show interfaces terse
Interface      Admin Link Proto Local Remote
so-0/0/0       up    up
so-0/0/0.0     up    up    inet  10.1.56.1/30
iso
mpls
so-0/0/1       up    up
so-0/0/1.0     up    up    inet  10.1.15.2/30
iso
mpls
so-0/0/2       up    up
so-0/0/2.0     up    up    inet  10.1.45.2/30
iso
mpls
so-0/0/3       up    down

user@R6> show interfaces terse
Interface      Admin Link Proto Local Remote
so-0/0/0       up    up
so-0/0/0.0     up    up    inet  10.1.56.2/30
iso
mpls
so-0/0/1       up    up
so-0/0/1.0     up    up    inet  10.1.46.2/30
iso
mpls
so-0/0/2       up    up
so-0/0/2.0     up    up    inet  10.1.26.2/30
iso
mpls
so-0/0/3       up    up
so-0/0/3.0     up    up    inet  10.1.36.2/30
iso
mpls

```

Sample Output 2

```

user@R6> show interfaces terse
Interface      Admin Link Proto Local Remote
so-0/0/0       up    up
so-0/0/0.0     up    up    inet  10.1.56.2/30
iso
mpls
so-0/0/1       up    up
so-0/0/1.0     up    up    inet  10.1.46.2/30
iso

```

```

so-0/0/2          up    up    mpls
so-0/0/2.0        up  up  inet 10.1.26.2/30
                   iso  #The mpls statement is missing.

so-0/0/3          up    up
so-0/0/3.0        up  up  inet 10.1.36.2/30
                   iso
                   mpls

```

Meaning Sample Output 1 shows the interface, the administrative status of the link (**Admin**), the data link layer status of the link (**Link**), the protocol families configured on the interface (**Proto**), and the local and remote addresses on the interface.

All interfaces on all routes in the network shown in [Figure 5 on page 46](#) are administratively enabled and functioning at the data link layer with MPLS and IS-IS, and have an **inet** address. All are configured with an IPv4 protocol family (**inet**), and have the IS-IS (**iso**) and MPLS (**mpls**) protocol families configured at the `[edit interfaces type-fpc/pic/port unit number]` hierarchy level.

Sample Output 2 shows that interface **so-0/0/2.0** on **R6** does not have the **mpls** statement included at the `[edit interfaces type-fpc/pic/port unit number]` hierarchy level. For information on how to configure MPLS on an interface, see [“Configuring MPLS on Your Network” on page 6](#).

Verify MPLS Labels

You can use the **traceroute** command or the **ping mpls** command to verify that packets are being sent over the LSP.

To verify MPLS labels and that packets are sent over the LSP, follow these steps:

1. [Use the traceroute Command to Verify MPLS Labels on page 55](#)
2. [Use the ping Command to Verify MPLS Labels on page 56](#)

Use the traceroute Command to Verify MPLS Labels

Purpose You can use the **traceroute** command to verify that packets are being sent over the LSP.

Action To verify MPLS labels, enter the following Junos OS CLI operational mode command, where **host-name** is the IP address or the name of the remote host:

```
user@host> traceroute host-name
```

Sample Output 1

```
user@R1> traceroute 100.100.6.1
traceroute to 100.100.6.1 (100.100.6.1), 30 hops max, 40 byte packets
 1  10.1.12.2 (10.1.12.2)  0.861 ms  0.718 ms  0.679 ms
    MPLS Label=100048 CoS=0 TTL=1 S=1
 2  10.1.24.2 (10.1.24.2)  0.822 ms  0.731 ms  0.708 ms
    MPLS Label=100016 CoS=0 TTL=1 S=1
 3  10.1.46.2 (10.1.46.2)  0.571 ms !N  0.547 ms !N  0.532 ms !N
```

Sample Output 2

```
user@R1> traceroute 10.0.0.6
traceroute to 10.0.0.6 (10.0.0.6), 30 hops max, 40 byte packets
 1  10.1.13.2 (10.1.13.2)  0.605 ms  0.548 ms  0.503 ms
 2  10.0.0.6 (10.0.0.6)  0.761 ms  0.676 ms  0.675 ms
```

Meaning Sample Output 1 shows that MPLS labels are used to forward packets through the network. Included in the output is a label value (**MPLS Label=100048**), the time-to-live value (**TTL=1**), and the stack bit value (**S=1**).

The **MPLS Label** field is used to identify the packet to a particular LSP. It is a 20-bit field, with a maximum value of ($2^{20}-1$), or approximately 1,000,000.

The TTL value contains a limit on the number of hops that this MPLS packet can travel through the network (1). It is decremented at each hop, and if the TTL value drops below one, the packet is discarded.

The bottom of the stack bit value (**S=1**) indicates that is the last label in the stack and that this MPLS packet has one label associated with it. The MPLS implementation in the Junos OS supports a stacking depth of 3 on the M-series routers and up to 5 on the T-series platforms. For more information on MPLS label stacking, see RFC 3032, *MPLS Label Stack Encoding*.

MPLS labels appear in Sample Output 1 because the **traceroute** command is issued to a BGP destination where the BGP next hop for that route is the LSP egress address. The Junos OS default behavior uses LSPs for BGP traffic when the BGP next hop equals the LSP egress address.

Sample Output 2 shows that MPLS labels do not appear in the output for the **traceroute** command. If the BGP next hop does not equal the LSP egress address or the destination is an IGP route, the BGP traffic does not use the LSP. Instead of using the LSP, the BGP traffic is using the IGP (IS-IS, in this case) to reach the egress address (**R6**).

Use the ping Command to Verify MPLS Labels

Purpose On the egress router (the router receiving the MPLS echo packets), the IP address 127.0.0.1/32 is configured by default on its loopback (**lo0**) interface, resulting in echo requests being sent as MPLS packets destined for the address 127.0.0.1 and the well-known port 3503. When the echo request arrives at the egress router, the receiver checks the contents of the packet and sends a reply containing the correct return value.

The sender of the echo request waits 2 seconds for the echo reply, then times out. In the example network shown in [Figure 5 on page 46](#), the egress router is **R6**.

Action To verify MPLS labels, in operational mode, enter the following command to ping the egress router:

```
user@ingress-router> ping mpls rsvp lsp-name detail
```

Sample Output 1

```
user@R1> ping mpls rsvp R1-to-R6 detail
Request for seq 1, to interface 69, label 100064
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 69, label 100064
Reply for seq 2, return code: Egress-ok
Request for seq 3, to interface 69, label 100064
Reply for seq 3, return code: Egress-ok
Request for seq 4, to interface 69, label 100064
Reply for seq 4, return code: Egress-ok
Request for seq 5, to interface 69, label 100064
Reply for seq 5, return code: Egress-ok

--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

Meaning Sample Output 1 from ingress router **R1** shows that an echo request is sent with a label (100064), indicating that the echo requests were sent over the LSP **R1-to-R6**.

CHAPTER 3

Determining the LSP State

This chapter describes how to display the status and statistics of the Multiprotocol Label Switching (MPLS) protocol running on all routers in a network. You can use a variety of operational mode commands to determine status and statistics information useful in diagnosing problem situations.

- [Checklist for Determining LSP Status on page 59](#)
- [Determining LSP Status on page 59](#)
- [Determining LSP Statistics on page 65](#)

Checklist for Determining LSP Status

Purpose This checklist provides the steps and commands to verify the state of a label-switched path (LSP) in an MPLS network. The checklist includes links to more detailed information about the commands to verify the LSP and supporting protocols. [Table 5 on page 59](#) provides commands for determining the LSP state.

Table 5: Checklist for Determining the LSP State

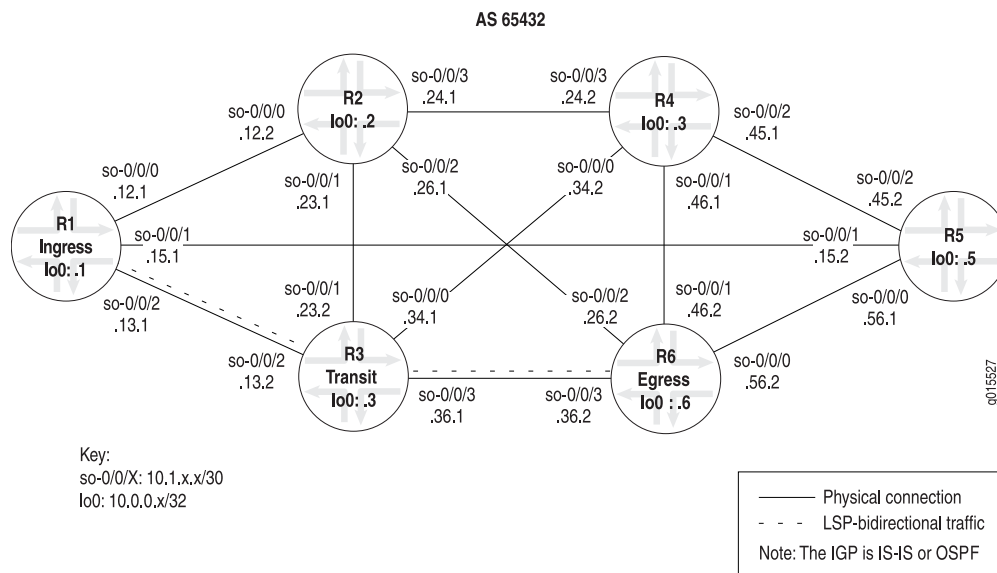
Tasks	Command or Action
“Determining LSP Status” on page 59	
1. Check the Status of the LSP on page 60	<code>show mpls lsp</code>
2. Display Extensive Status About the LSP on page 61	<code>show mpls lsp extensive</code>
“Determining LSP Statistics” on page 65	
	<code>show rsvp session detail</code>

Determining LSP Status

Display detailed information about Resource Reservation Protocol (RSVP) objects and the label-switched path (LSP) history to pinpoint a problem with the LSP.

[Figure 6 on page 60](#) illustrates the network topology used in this topic. For more details about the router configurations in this network, see [“Checklist for Configuring and Verifying an MPLS Network” on page 3](#).

Figure 6: MPLS Network Topology



To determine the LSP state, follow these steps:

1. [Check the Status of the LSP on page 60](#)
2. [Display Extensive Status About the LSP on page 61](#)

Check the Status of the LSP

Purpose Display the status of the label-switched pathe (LSP).

Action To determine the LSP status, on the ingress router, enter the following Junos OS command-line interface (CLI) operational mode command:

```
user@host> show mpls lsp
```

Sample Output

```
user@R1> show mpls lsp
Ingress LSP: 1 sessions
To          From          State Rt ActivePath      P      LSPname
10.0.0.6    10.0.0.1    Up 1      * R1-to-R6
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions
To          From          State Rt Style  Labelin Labelout LSPname
10.0.0.1    10.0.0.6    Up 01 FF  3    - R6-to-R1
Total 1 displayed, Up 1, Down 0

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

Meaning The sample output is from the ingress router (R1), and shows ingress, egress, and transit LSP information. Ingress information is for the sessions that originate from this router, egress information is for sessions that terminate on this router, and transit information is for sessions that transit through this router.

There is one ingress route from **R1 (10.0.0.1)** to **R6 (10.0.0.6)**. This route is currently up, and is an active route installed in the routing table (**Rt**). The LSP **R1-to-R6** is the primary path (**P**) as opposed to the secondary path, and is indicated by an asterisk (*****). The route to **R6** does not contain a named path (**ActivePath**).

There is one egress LSP from **R6** to **R1**. The **State** is up, with no routes installed in the routing table. RSVP reservation style (**Style**) consists of two parts. The first is the number of active reservations (**1**). The second is the reservation style, which is **FF** (fixed filter). The reservation style can be **FF**, **SE** (shared explicit), or **WF** (wildcard filter). There are three incoming labels (**Labelin**) and no labels going out (**Labelout**) for this LSP.

There are no transit LSPs.

For more information on checking the LSP state, see [“Checklist for Working with the Layered MPLS Troubleshooting Model”](#) on page 85.

Display Extensive Status About the LSP

Purpose Display extensive information about LSPs, including all past state history and the reasons why an LSP might have failed.

Action To display extensive information about LSPs, on the ingress router, enter the following Junos OS CLI operational mode command:

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

Sample Output

```

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

10.0.0.6
  From: 10.0.0.1, State: Up , ActiveRoute: 1 , LSPName: R1-to-R6
  ActivePath: (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
10.1.13.2 S 10.1.36.2 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

    10.1.13.2 10.1.36.2
91 Aug 17 12:22:52 Selected as active path
90 Aug 17 12:22:52 Record Route: 10.1.13.2 10.1.36.2
89 Aug 17 12:22:52 Up
88 Aug 17 12:22:52 Originate Call
87 Aug 17 12:22:52 CSPF: computation result accepted
86 Aug 17 12:22:23 CSPF failed: no route toward 10.0.0.6[13920 times]
85 Aug 12 19:12:51 Clear Call
84 Aug 12 19:12:50 10.1.56.2: MPLS label allocation failure
83 Aug 12 19:12:47 Deselected as active
82 Aug 12 19:12:47 10.1.56.2: MPLS label allocation failure
81 Aug 12 19:12:47 ResvTear received
80 Aug 12 19:12:47 Down
79 Aug 12 19:12:31 10.1.56.2: MPLS label allocation failure[4 times]
78 Aug 12 19:09:58 Selected as active path
77 Aug 12 19:09:58 Record Route: 10.1.15.2 10.1.56.2
76 Aug 12 19:09:58 Up
75 Aug 12 19:09:57 Originate Call
74 Aug 12 19:09:57 CSPF: computation result accepted
73 Aug 12 19:09:29 CSPF failed: no route toward 10.0.0.6[11 times]
72 Aug 12 19:04:36 Clear Call
71 Aug 12 19:04:23 Deselected as active
70 Aug 12 19:04:23 ResvTear received
69 Aug 12 19:04:23 Down
68 Aug 12 19:04:23 CSPF failed: no route toward 10.0.0.6
67 Aug 12 19:04:23 10.1.15.2: Session preempted
66 Aug 12 16:45:35 Record Route: 10.1.15.2 10.1.56.2
65 Aug 12 16:45:35 Up
64 Aug 12 16:45:35 Clear Call
63 Aug 12 16:45:35 CSPF: computation result accepted
62 Aug 12 16:45:35 ResvTear received
61 Aug 12 16:45:35 Down
60 Aug 12 16:45:35 10.1.13.2: Session preempted
59 Aug 12 14:50:52 Selected as active path
58 Aug 12 14:50:52 Record Route: 10.1.13.2 10.1.36.2
57 Aug 12 14:50:52 Up
56 Aug 12 14:50:52 Originate Call
55 Aug 12 14:50:52 CSPF: computation result accepted
54 Aug 12 14:50:23 CSPF failed: no route toward 10.0.0.6[7 times]
53 Aug 12 14:47:22 Deselected as active
52 Aug 12 14:47:22 CSPF failed: no route toward 10.0.0.6
51 Aug 12 14:47:22 Clear Call
50 Aug 12 14:47:22 CSPF: link down/deleted
10.1.12.1(R1.00/10.0.0.1)->10.1.12.2(R2.00/10.0.0.2)
49 Aug 12 14:47:22 CSPF: link down/deleted

```

```

10.1.15.1(R1.00/10.0.0.1)->10.1.15.2(R5.00/10.0.0.5)
 48 Aug 12 14:47:22 10.1.15.1: MPLS label allocation failure
 47 Aug 12 14:47:22 Clear Call
 46 Aug 12 14:47:22 CSPF: computation result accepted
 45 Aug 12 14:47:22 10.1.12.1: MPLS label allocation failure
 44 Aug 12 14:47:22 MPLS label allocation failure
 43 Aug 12 14:47:22 Down
 42 Jul 23 11:27:21 Selected as active path
    Created: Sat Jul 10 18:18:44 2004
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, LSPstate: Up , ActiveRoute: 0
  LSPname: R6-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: 141, Since: Tue Aug 17 12:23:14 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 39024 protocol 0
  PATH rcvfrom: 10.1.15.2 (so-0/0/1.0) 130 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.1.56.2 10.1.15.2 <self>
Total 1 displayed, Up 1, Down 0

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

```

Meaning The sample output is from the ingress router (**R1**), and shows ingress, egress, and transit LSP information in detail, including all past state history and the reasons why an LSP failed. Ingress information is for sessions that originate from this router, egress information is for sessions that terminate on this router, and transit information is for sessions that transit through this router.

There is one ingress route from **R1 (10.0.0.1)** to **R6 (10.0.0.6)**. This route is currently up (**State**), with one route actively using the LSP, **R1-to-R6**. The LSP active path is the primary path. Even if the LSP does not contain a **primary** or **secondary** keyword, the router still treats the LSP as a primary LSP, indicating that if the LSP fails, the router will attempt to signal inactive LSPs at 30-second intervals, by default.

Load balancing is **Random**, which is the default, indicating that when selecting the physical path for an LSP, the router randomly selects among equal-cost paths that have an equal hop count. Other options that you can configure are **Least-fill** and **Most-fill**. **Least-fill** places the LSP over the least utilized link of the equal-cost paths with equal hop count. **Most-fill** places the LSP over the most utilized link of the equal-cost paths sharing an equal hop count. Utilization is based on the percentage of available bandwidth.

The **Encoding type** field shows Generalized MPLS (GMPLS) signaling parameters (**Packet**), indicating IPv4. The **Switching type** is **Packet**, and the Generalized Payload Identifier (**GPID**) is IPv4.

The primary path is the active path, as indicated by an asterisk (*). The state of the LSP is **Up**.

The Explicit Route Object (**ERO**) includes the Constrained Shortest Path First (CSPF) cost (**20**) for the physical path that the LSP follows. The presence of the CSPF metric indicates that this is a CSPF LSP. The absence of the CSPF metric indicates a no-CSPF LSP.

The field **10.1.13.2 S** indicates the actual ERO. The RSVP signaling messages went to **10.1.13.2** strictly (as a next hop) and finished at **10.1.36.2** strictly. All ERO addresses are strict hops when the LSP is a CSPF LSP. Loose hops can only display in a no-CSPF LSP.

The received Record Route Object (**RRO**) has the following protection flags:

- **0x01**—Local protection available. The link downstream of this node is protected by a local repair mechanism. This flag can only be set if the Local protection flag was set in the SESSION_ATTRIBUTE object of the corresponding path message.
- **0x02**—Local protection in use. A local repair mechanism is in use to maintain this tunnel (usually because of an outage of the link it was routed over previously).
- **0x04**—Bandwidth protection. The downstream router has a backup path providing the same bandwidth guarantee as the protected LSP for the protected section.
- **0x08**—Node protection. The downstream router has a backup path providing protection against link and node failure on the corresponding path section. If the downstream router can set up only a link-protection backup path, the “Local protection available” bit is set but the “Node protection” bit is cleared.
- **0x10**—Preemption pending. The preempting node sets this flag if a pending preemption is in progress for the traffic engineered LSP. This indicates to the ingress label edge router (LER) of this LSP that it should be rerouted.

For more information on protection flags, see the *Junos Routing Protocols and Policies Command Reference*.

The field **10.1.13.2.10.1.36.2** is the actual received record route (**RRO**). Note that the addresses in the **RRO** field match those in the **ERO** field. This is the normal case for CSPF LSPs. If the RRO and ERO addresses do not match for a CSPF LSP, the LSP has to reroute or detour.

The lines numbered 91 through 42 contain the 49 most recent entries to the history log. Each line is time stamped. The most recent entries have the largest log history number and are at the top of the log, indicating that line 91 is the most recent history log entry. When you read the log, start with the oldest entry (**42**) to the most recent (**91**).

The history log was started on July 10, and displays the following sequence of activities: an LSP was selected as active, was found to be down, MPLS label allocation failed several times, was deleted several times, was preempted because of an ResvTear, was deselected as active, and was cleared. In the end, the router computed a CSPF ERO, signaled the call, the LSP came up with the listed RRO (line 90), and was listed as active.

For more information on error messages, see the *Junos MPLS Network Operations Guide Log Reference*.

The total number of ingress LSPs displayed is 1, with 1 up and 0 down. The number in the **Up** field plus the number in the **Down** field should equal the total.

There is one egress LSP session from **R6** to **R1**. The **State** is up, with no routes installed in the routing table. RSVP reservation style (**Style**) consists of two parts. The first is the number of active reservations (1). The second is the reservation style, which is **FF** (fixed filter). The reservation style can be **FF**, **SE** (shared explicit), or **WF** (wildcard filter). There are three incoming labels (**Labelin**) and no labels going out (**Labelout**) for this LSP.

There are no transit LSPs.

For more information on checking the LSP state, see [“Checklist for Working with the Layered MPLS Troubleshooting Model” on page 85](#).

Determining LSP Statistics

- | | |
|----------------|--|
| Purpose | Display detailed information about RSVP objects to assist the diagnosis of an LSP problem. |
| Action | To verify RSVP objects, enter the following Junos OS CLI operational mode command:

user@host> show rsvp session detail |

Sample Output

```

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

10.0.0.6
  From: 10.0.0.1, LSPstate: Up, ActiveRoute: 1
  LSPname: R1-to-R6 , LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 100064
  Resv style: 1 FF, Label in: -, Label out: 100064
  Time left: -, Since: Tue Aug 17 12:22:52 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 12 receiver 44251 protocol 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  PATH sentto: 10.1.13.2 (so-0/0/2.0) 182 pkts
  RESV rcvfrom: 10.1.13.2 (so-0/0/2.0) 159 pkts
  Explct route: 10.1.13.2 10.1.36.2
  Record route: <self> 10.1.13.2 10.1.36.2
Total 1 displayed, Up 1, Down 0

Egress RSVP: 1 sessions

10.0.0.1
  From: 10.0.0.6 , LSPstate: Up, ActiveRoute: 0
  LSPname: R6-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: 135, Since: Tue Aug 17 12:23:14 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 39024 protocol 0
  PATH rcvfrom: 10.1.15.2 (so-0/0/1.0) 158 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.1.56.2 10.1.15.2 <self>
Total 1 displayed, Up 1, Down 0

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

```

Meaning The sample output shows that there is one ingress and one egress RSVP session. The ingress session has a source address of **10.0.0.1 (R1)**, and the session is up, with one active route. The LSP name is **R1-to-R6** and it is the primary path for the LSP.

The recovery label (**100064**) is sent by a graceful restart router to its neighbor to recover a forwarding state. It is probably the old label that the router advertised before it went down.

This session is using the fixed filter (**FF**) reservation style (**Resv style**). Since this is an ingress router, there is no inbound label. The outbound label (provided by the next downstream router) is **100064**.

The **Time Left** field provides the number of seconds remaining in the RSVP session, and the **Tspec** object provides information about the controlled load rate (**rate**) and maximum

burst size (**peak**), an infinite value (**Infbps**) for the guaranteed delivery option, and the indication that packets smaller than 20 bytes are treated as 20 bytes, while packets larger than 1500 bytes are treated as 1500 bytes.

The port number is the IPv4 tunnel ID, while the sender/receiver port number is the LSP ID. The IPv4 tunnel ID is unique for the life of the LSP, while the sender/receiver LSP ID can change, for example, with an SE style reservation.

The **PATH rcvfrom** field includes the source of the path message. Since this is the ingress router, the local client originated the path message.

The **PATH sentto** field includes the path message destination (**10.1.13.2**) and outgoing interface (**so-0/0/2.0**). The **RESV rcvfrom** field includes both the source of the Resv message received (**10.1.13.2**) and the incoming interface (**so-0/0/2.0**).

The RSVP explicit route and the route record values are identical: **10.1.13.2** and **10.1.36.2**. In most cases, the explicit route and the record route values are identical. Differences indicate that some path rerouting has occurred, typically during Fast-Reroute.

The **Total** fields indicate the total number of ingress, egress, and transit RSVP sessions, with the total being equal to the sum of the up and down sessions. In this example, there is one ingress session, one egress session, and no transit RSVP sessions.

Verifying RSVP Signal Processing

This chapter describes how to determine that the Resource Reservation Protocol (RSVP) path messages are sent and received.

- [Checklist for Verifying RSVP Signal Processing on page 69](#)
- [Checking That RSVP Path Messages Are Sent and Received on page 70](#)
- [Examining the History Log on page 71](#)
- [Determining the Current RSVP Neighbor State on page 72](#)
- [Enabling RSVP Traceoptions on page 73](#)

Checklist for Verifying RSVP Signal Processing

Purpose This checklist provides the steps and commands to verify Resource Reservation Protocol (RSVP) signal processing in an MPLS network. The checklist includes links to more detailed information about checking RSVP messages, examining the history log, determining RSVP neighbors, and commands to configure RSVP traceoptions.

[Table 6 on page 69](#) provides commands for verifying RSVP signal processing.

Table 6: Checklist for Verifying RSVP Signal Processing

Tasks	Command or Action
“Checking That RSVP Path Messages Are Sent and Received” on page 70	<code>show rsvp statistics</code>
“Examining the History Log” on page 71	<code>show mpls lsp extensive</code>
“Determining the Current RSVP Neighbor State” on page 72	<code>show rsvp neighbor</code>
“Enabling RSVP Traceoptions” on page 73	<code>[edit]</code> <code>edit protocols rsvp traceoptions</code> <code>set file filename.log</code> <code>set flag packets</code> <code>show</code> <code>commit</code> <code>run show log rsvp.log</code> <code>deactivate traceoptions</code> <code>show</code> <code>commit</code>

Checking That RSVP Path Messages Are Sent and Received

Purpose The presence or absence of various RSVP messages can help determine if there is a problem with Multiprotocol Label Switching (MPLS) in your network. For example, if path messages occur in the output without Resv messages, it might indicate that label-switched paths (LSPs) are not being created.

Action To check that RSVP Path messages are sent and received, enter the following Junos OS command-line interface (CLI) operational mode command:

```
user@host> show rsvp statistics
```

Sample Output

```
user@R1> show rsvp statistics
PacketType          Total          Last 5 seconds
      Sent      Received      Sent      Received
Path      114523      80185      1      0
PathErr      5      10      0      0
PathTear      12      6      0      0
Resv FF      80515      111476      0      0
Resv WF              0      0      0      0
Resv SE              0      0      0      0
ResvErr              0      0      0      0
ResvTear      0      5      0      0
ResvConf              0      0      0      0
Ack              0      0      0      0
SRefresh              0      0      0      0
Hello      915851      915881      0      0
EndtoEnd RSVP              0      0      0      0

Errors              Total          Last 5 seconds
Rcv pkt bad length      0      0
Rcv pkt unknown type    0      0
Rcv pkt bad version      0      0
Rcv pkt auth fail        0      0
Rcv pkt bad checksum      0      0
Rcv pkt bad format        0      0
Memory allocation fail    0      0
No path information        0      0
Resv style conflict        0      0
Port conflict              0      0
Resv no interface          0      0
PathErr to client      15      0
ResvErr to client        0      0
Path timeout              0      0
Resv timeout              0      0
Message out-of-order      0      0
Unknown ack msg            0      0
Recv nack                  0      0
Recv duplicated msg-id      0      0
No TE-link to rcv Hop      0      0
```

Meaning The sample output shows RSVP messages sent and received. The total number of RSVP Path messages is 11,4532 sent and 80,185 received. Within the last 5 seconds, no messages have been sent or received.

A total of 5 **PathErr** messages were sent and 10 received. When path errors occur (usually because of parameter problems in a path message), the router sends a unicast PathErr message to the sender that issued the path message. In this case, **R1** sent at least 10 path messages with an error, as indicated by the 10 PathErr messages that **R1** has received. The downstream router sent **R1** five path messages with an error, as indicated by the five PathErr messages that **R1** has sent. PathErr messages transmit in the opposite direction to path messages.

A total of 12 **PathTear** messages were sent and 6 received, none in the last 5 seconds. In contrast to PathErr messages, PathTear messages travel in the same direction as path messages. Since path messages are both sent and received, PathTear messages are also sent and received. However, if only path messages are sent, then only the PathTear messages that are sent appear in the output.

A total of 80,515 reservation (**Resv**) messages with the fixed filter (**FF**) reservation style were sent and 111,476 received, none in the last 5 seconds. An **FF** reservation style indicates that within each session, each receiver establishes its own reservation with each upstream sender, and that all selected senders are listed. No messages for the wildcard filter (**WF**) or shared explicit (**SE**) reservation styles are sent or received. For more information on RSVP reservation styles, see the *Junos MPLS Applications Configuration Guide*.

Other RSVP message types are not sent or received. For information on the ResvErr, ResvTear, and Resvconf message types, see the *Junos MPLS Applications Configuration Guide*.

Ack and summary refresh (SRefresh) messages do not appear in the output. Ack and summary refresh messages are defined in RFC 2961 and are part of the RSVP extensions. Ack messages are used to reduce the amount of RSVP control traffic in the network.

A total of 915,851 hello messages were sent and 915,881 received, with none transmitted or received in the last 5 seconds. The RSVP hello interval is 9 seconds. If more than one hello message is sent or received in the last 5 seconds, it implies that more than one interface supports RSVP.

EndtoEnd RSVP messages are legacy RSVP messages that are not used for RSVP traffic engineering. These counters increment only when RSVP forwards legacy RSVP messages issued by a virtual private network (VPN) customer for transit across the backbone to the other site(s) in the VPN. They are called end-to-end messages because they are intended for the opposite side of the network and only have meaning at the two ends of the provider network.

The **Errors** section of the output shows statistics about RSVP packets with errors. A total of 15 **PathErr to client** packets were sent to the Routing Engine. The total combines the sent and received **PathErr** packets. For more information about error statistics and packets, see the *Junos System Basics and Services Command Reference*.

Examining the History Log

Purpose The history log for the **show mpls lsp** extensive command contains information that is useful in determining a possible reason for any errors in MPLS functioning in your network.

Action To examine the history log, enter the following Junos OS CLI operational mode command:

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

Sample Output

```
user@R6> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.1
From: 10.0.0.6, State: Up, ActiveRoute: 1, LSPname: R6-to-R1
ActivePath: (primary)
LoadBalance: Random
Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary State: Up
Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
10.1.56.1 S 10.1.15.1 S
Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

    10.1.56.1 10.1.15.1
    6 Aug 17 12:19:04 Selected as active path
    5 Aug 17 12:19:03 Record Route: 10.1.56.1 10.1.15.1
    4 Aug 17 12:19:03 Up
    3 Aug 17 12:19:03 Originate Call
    2 Aug 17 12:19:03 CSPF: computation result accepted
    1 Aug 17 12:18:34 CSPF failed: no route toward 10.0.0.1
Created: Tue Aug 17 12:18:33 2004
Total 1 displayed, Up 1, Down 0
[...Output truncated...]
```

Meaning Lines 1 through 6 contain the six most recent entries to the history log. Each line is time stamped. The most recent entries have the largest log history number and are at the top of the log, indicating that line 6 is the most recent entry in the history log.

The history log was started on August 17, and displays the following sequence of activities: a call failed because the address could not be reached (line 1); 31 seconds later, probably because the addressing problem was resolved, the call was signaled (line 2); the call was completed (line 3); the LSP came up with a route (lines 4 and 5); and the LSP was selected as active (line 6).

For more details about the messages that can appear in the history log, see *Junos MPLS Operations Guide: Log Files*.

Determining the Current RSVP Neighbor State

Purpose Display a list of RSVP neighbors that were learned dynamically when exchanging RSVP packets. Once a neighbor is learned, it is never removed from the list of RSVP neighbors.

Action To determine the current RSVP neighbor state, enter the following Junos OS CLI operational mode command:

```
user@host> show rsvp neighbor
```

Sample Output

```
user@R6> show rsvp neighbor
RSVP neighbor: 2 learned
Address  Idle Up/Dn LastChange  HelloInt  HelloTx/Rx  MsgRcvd
10.1.36.1   5  1/0  1w5d 6:30:50    9    116734/116734  23558
10.1.56.1  10  1/0  2w2d 23:44:15   9    161600/161600  23570
```

Meaning The sample output shows that **R6** has learned about two different RSVP neighbors. Each neighbor has one line of output that includes the neighbor RSVP address, the length of time the interface was idle, the current interface up/down counter, the time of the last interface state change, the current RSVP hello interval, the total number of RSVP hello messages transmitted and received, and the total number of RSVP messages received on the interface.

The **show rsvp neighbor** command only indicates a neighbor after a session is established. Once an interface is displayed in this command output, it always appears, even if the RSVP neighbor state is down.

The RSVP neighbor **10.1.36.1** was idle for 5 seconds, came up once and has not gone down, indicating that the interface is currently in an **Up** state. As long as the up counter is one greater than the down counter, the RSVP interface is up. If the up/down counters are equal, the interface is down.

The last state change occurred 6 hours and 30 minutes ago. The current hello interval is 9 seconds. A total of 116,734 hello messages were transmitted and received on this interface, and a total of 23,558 RSVP Path/Resv messages were processed.

The RSVP neighbor **10.1.56.1** was idle for 10 seconds, came up once and has not gone down, indicating that the interface is currently in an **Up** state. The last state change occurred 23 hours and 44 minutes ago. The current Hello interval is 9 seconds. A total of 161,600 hello messages were transmitted and received on this interface, and a total of 23,570 RSVP Path/Resv messages were processed.

Enabling RSVP Traceoptions

Purpose Global routing protocol tracing operations track all general routing operations and record them in a log file. Any global tracing operations that you configure are inherited by the individual routing protocols. To modify the global tracing operations for an individual protocol, enable tracing when configuring that protocol.

The error descriptions logged by the remote operations daemon can often provide more detailed information to help you solve the problem faster.

Action To enable traceoptions for RSVP packets in your network, follow these steps:

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

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

2. Configure the RSVP log file:

```
[edit protocols rsvp traceoptions]  
user@host# set file filename.log
```

3. Configure the tracing operations:

```
[edit protocols rsvp traceoptions]  
user@host# set flag packets
```

4. Verify and commit the configuration:

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

5. View the contents of the log file:

```
user@host# run show log rsvp.log
```

6. Stop monitoring the **rsvp** log file:

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

7. Verify and commit the new configuration:

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

```

Sample Output user@R1> edit
                  Entering configuration mode

                  [edit]
user@R1# edit protocols rsvp traceoptions

[edit protocols rsvp traceoptions]
user@R1# set file rsvp.log

[edit protocols rsvp traceoptions]
user@R1# set flag packets

[edit protocols rsvp traceoptions]
user@R1# show
file rsvp.log;
flag packets;

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

[edit protocols rsvp traceoptions]
user@R1# run show log rsvp.log
Aug 26 10:05:54 trace_on: Tracing to "/var/log/rsvp.log" started
Aug 26 10:05:54 RSVP send Hello New 10.1.13.1->10.1.13.2 Len=32 so-0/0/2.0
Aug 26 10:05:55 RSVP recv Resv 10.1.13.2->10.1.13.1 Len=128 so-0/0/2.0
Aug 26 10:05:55 RSVP send Hello New 10.1.12.1->10.1.12.2 Len=32 so-0/0/0.0
Aug 26 10:05:55 RSVP send Hello New 10.1.15.1->10.1.15.2 Len=32 so-0/0/1.0
Aug 26 10:05:55 RSVP recv Hello New 10.1.12.2->10.1.12.1 Len=32 so-0/0/0.0
Aug 26 10:05:55 RSVP recv Hello New 10.1.15.2->10.1.15.1 Len=32 so-0/0/1.0
Aug 26 10:05:57 RSVP recv Path 10.0.0.6->10.0.0.1 Len=208 so-0/0/1.0
Aug 26 10:05:57 RSVP send Resv 10.1.15.1->10.1.15.2 Len=120 so-0/0/1.0
---(more)---[abort]

[edit protocols rsvp traceoptions]
user@R1# up

[edit protocols rsvp]
user@R1# deactivate traceoptions

[edit protocols rsvp]
user@R1# show
inactive: traceoptions {
    file rsvp.log;
    flag packets;
}
interface so-0/0/0.0;
interface so-0/0/1.0;
interface so-0/0/2.0;
interface fxp0.0 {
    disable;
}

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

```

Meaning The sample output shows the configuration of RSVP traceoptions, the output for the log file, and the deactivation of the traceoptions configuration.

To specify more than one tracing operation, include multiple flag statements in the configuration, at the following hierarchy level:

```
[edit protocols rsvp traceoptions]  
user@R1# set flag flag
```

Table 7: RSVP Tracing Flags

Flag	Description
all	All tracing operations
error	All detected error conditions
event	RSVP-related events
lmp	RSVP-LMP interactions
packets	All RSVP packets
path	All path messages
pathtear	PathTear messages
resv	Resv messages
resvtear	ResvTear messages
route	Routing information
state	Session state transitions

For more information on configuring traceoptions, see the *Junos MPLS Applications Configuration Guide* and the *Junos Routing Protocols Configuration Guide*.

Verifying LSP Use

- This chapter describes how to verify the availability and valid use of a label-switched path (LSP) in your network.
- [Checklist for Verifying LSP Use on page 77](#)
 - [Verifying LSP Use in Your Network on page 77](#)

Checklist for Verifying LSP Use

Purpose This checklist provides the steps and commands to verify the use of the LSP in an MPLS network. The checklist includes links to more detailed information about verifying the LSP on the ingress and transit routers in the network.

This checklist describes how to verify the availability and valid use of a label-switched path (LSP) in your network.

Table 8: Checklist for Verifying LSP Use


Tasks	Command or Action
“Verifying LSP Use in Your Network” on page 77	
“Verifying an LSP on the Ingress Router” on page 78	<code>show route table inet.3</code>
“Verifying an LSP on a Transit Router” on page 79	<code>show route table mpls.0</code>

[Table 8 on page 77](#) provides commands for verifying LSP use.

Verifying LSP Use in Your Network

Purpose When you verify the valid use of an LSP on the ingress and transit routers in your network, you can determine if there is a problem with Multiprotocol Label Switching (MPLS) in your network. [Figure 7 on page 78](#) describes the example network used in this topic.

Figure 7: MPLS Topology for Verifying LSP Use



The MPLS network in [Figure 7 on page 78](#) illustrates a router-only network with SONET interfaces that consist of the following components:

- A full-mesh interior Border Gateway Protocol (IBGP) topology, using AS 65432
- MPLS and Resource Reservation Protocol (RSVP) enabled on all routers
- A **send-statics** policy on routers R1 and R6 that allows a new route to be advertised into the network
- An LSP between routers R1 and R6

The network shown in [Figure 7 on page 78](#) is a Border Gateway Protocol (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 the full configuration for each router in the example network, see [“Checklist for Configuring and Verifying an MPLS Network” on page 3](#).

To verify LSP use in your network, follow these steps:

1. [Verifying an LSP on the Ingress Router on page 78](#)
2. [Verifying an LSP on a Transit Router on page 79](#)

Verifying an LSP on the Ingress Router

Purpose You can verify the availability of an LSP when it is up by examining the **inet.3** routing table on the ingress router. The **inet.3** routing table contains the host address of each LSP's egress router. This routing table is used on ingress routers to route BGP packets to the destination egress router. BGP uses the **inet.3** routing table on the ingress router to help resolve next-hop addresses.

Action To verify an LSP on an ingress router, enter the following Junos OS command-line interface (CLI) operational mode command:

```
user@host> show route table inet.3
```

Sample Output

```
user@R1> show route table inet.3
inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.6/32          *[RSVP/7] 4w0d 22:40:57, metric 20
                    > via so-0/0/2.0, label-switched-path R1-to-R6
```

Meaning The sample output shows the **inet.3** routing table. By default, only BGP and MPLS virtual private networks (VPNs) can use the **inet.3** route table to resolve next-hop information. One destination is listed in the route table, **10.0.0.6**. This destination (**10.0.0.6**) is signaled by RSVP, and is the current active path, as indicated by the asterisk (*). The protocol preference for this route is **7**, and the metric associated with it is **20**. The label-switched path is **R1-to-R6**, through interface **so-0/0/2.0**, which is the physical next-hop transit interface.

Typically, the penultimate router in the LSP either pops the packet's label or changes the label to a value of 0. If the penultimate router pops the top label and an IPv4 packet is underneath, the egress router routes the IPv4 packet, consulting the IP routing table **inet.0** to determine how to forward the packet. If another type of label (such as one created by Label Distribution Protocol (LDP) tunneling or VPNs, but not IPv4) is underneath the top label, the egress router does not examine the **inet.0** routing table. Instead, it examines the **mpls.0** routing table for forwarding decisions.

If the penultimate router changes the packet's label to a value of 0, the egress router strips off the 0 label, indicating that an IPv4 packet follows. The packet is examined by the **inet.0** routing table for forwarding decisions.

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

When BGP resolves a next-hop prefix, it examines both the **inet.0** and **inet.3** routing tables, seeking the next hop with the lowest preference; for example, RSVP preference 7 is preferred over OSPF preference 10. The RSVP signaled LSP is used to reach the BGP next hop. This is the default when the BGP next hop equals the LSP egress address. Once the BGP next hop is resolved through an LSP, the BGP traffic uses the LSP to forward BGP transit traffic.

Verifying an LSP on a Transit Router

Purpose You can verify the availability of an LSP when it is up by examining the **mpls.0** routing table on a transit router. MPLS maintains the **mpls.0** routing table, which contains a list of the next label-switched router in each LSP. This routing table is used on transit routers to route packets to the next router along an LSP.

Action To verify an LSP on a transit router, enter the following Junos OS CLI operational mode command:

```
user@host> show route table mpls.0
```

Sample Output

```
user@R3> show route table mpls.0
mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                * [MPLS/0] 7w3d 22:20:56, metric 1
                  Receive
1                * [MPLS/0] 7w3d 22:20:56, metric 1
                  Receive
2                * [MPLS/0] 7w3d 22:20:56, metric 1
                  Receive
100064            * [RSVP/7] 2w1d 04:17:36, metric 1
                  > via so-0/0/3.0, label-switched-path R1-to-R6
100064 (S=0)      * [RSVP/7] 2w1d 04:17:36, metric 1
                  > via so-0/0/3.0, label-switched-path R1-to-R6
```

Meaning The sample output from transit router **R3** shows route entries in the form of MPLS label entries, indicating that there is only one active route, even though there are five active entries.

The first three MPLS labels are reserved MPLS labels defined in RFC 3032. Packets received with these label values are sent to the Routing Engine for processing. 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 two entries with the **100064** label are for the same LSP, **R1-to-R6**. There are two entries because the stack values in the MPLS header may be different. The second entry, **100064 (S=0)**, indicates that the stack depth is not 1 and additional label values are included in the packet. In contrast, the first entry of **100064** has an inferred S=1 which indicates a stack depth of 1 and makes it the last label in the packet. The dual entry indicates that this is the penultimate router. For more information on MPLS label stacking, see RFC 3032, *MPLS Label Stack Encoding*.

The incoming label is the MPLS header of the MPLS packet, and is assigned by RSVP to the upstream neighbor. Juniper Networks routers dynamically assign labels for RSVP traffic-engineered LSPs in the range from 100,000 through 1,048,575.

The router assigns labels starting at label 100,000, in increments of 16. The sequence of label assignments is 100,000, 100,016, 100,032, 100,048, and so on. At the end of the assigned labels, the label numbers start over at 100001, incrementing in units of 16. Juniper Networks reserves labels for various purposes. [Table 9 on page 80](#) lists the various label range allocations for incoming labels.

Table 9: MPLS Label Range Allocations

Incoming Label	Status
0 through 15	Reserved by IETF
16 through 1023	Reserved for static LSP assignment

Table 9: MPLS Label Range Allocations (*continued*)

Incoming Label	Status
1024 through 9999	Reserved for internal use (for example, CCC labels)
10,000 through 99,999	Reserved for static LSP assignment
100,000 through 1,048,575	Reserved for dynamic label assignment

PART 2

Working with Problems on Your Network

- [Working with the Layered MPLS Troubleshooting Model on page 85](#)
- [Verifying the Physical Layer on page 93](#)
- [Checking the Data Link Layer on page 103](#)
- [Verifying the IP and IGP Layers on page 117](#)
- [Checking the RSVP Layer on page 155](#)
- [Checking the MPLS Layer on page 173](#)
- [Checking the BGP Layer on page 193](#)

CHAPTER 6

Working with the Layered MPLS Troubleshooting Model

This chapter describes the different layers that you must verify when troubleshooting a Multiprotocol Label Switching (MPLS) network. The chapter also includes the example network used throughout the book to illustrate various problems that can occur in an MPLS network.

- [Checklist for Working with the Layered MPLS Troubleshooting Model on page 85](#)
- [Understanding the Layered MPLS Troubleshooting Model on page 85](#)

Checklist for Working with the Layered MPLS Troubleshooting Model

Problem This checklist provides a link to more detailed information about the layered Multiprotocol Label Switching network.

Solution [Table 10 on page 85](#) provides commands for working with the layered MPLS troubleshooting model.

Table 10: Checklist for Working with the Layered MPLS Troubleshooting Model

Tasks	Command or Action
“Understanding the Layered MPLS Troubleshooting Model” on page 85	<code>show mpls lsp</code> <code>show mpls lsp extensive</code> <code>show mpls lsp name <i>name</i></code> <code>show mpls lsp name <i>name</i> extensive</code>

Understanding the Layered MPLS Troubleshooting Model

Problem The layered MPLS troubleshooting model is a disciplined approach to investigating problems with an MPLS network. [Figure 8 on page 86](#) illustrates the layers in the model, and the commands you can use to structure your investigation. Because of the complexity of the MPLS network, you can obtain much better results from your investigations if you progress through the layers and verify the functioning of each layer on the ingress, egress, and transit routers before moving on to the next layer.

Solution Figure 8 on page 86 shows the layered MPLS troubleshooting model that you can use to troubleshoot problems with your MPLS network.

Figure 8: Layered MPLS Network Troubleshooting Model

BGP Layer	tracroute <i>host-name</i> show bgp summary show configuration protocols bgp show route <i>destination-prefix</i> detail show route receive protocol bgp <i>neighbor-address</i>
MPLS Layer	show mpls lsp show mpls lsp extensive show route table mpls,0 show route <i>address</i> traceroute <i>address</i> ping mpls rsvp <i>Lsp-name</i> detail
RSVP Layer	show rsvp session show rsvp neighbor show rsvp interface
<div>↩</div> IGP and IP Layers Functioning <div>↪</div>	
OSPF Layer show ospf neighbor show configuration protocols ospf show ospf interface	IS-IS Layer show isis adjacency show configuration protocols isis show isis interface
IP Layer show ospf neighbor extensive show interfaces terse	IP Layer show isis adjacency extensive show interfaces terse
Data Link Layer	show interfaces extensive <i>JUNOS Interfaces Network Operations Guide</i>
Physical Layer	show interfaces show interfaces terse ping <i>host</i>

g015528

As you move from one layer of the model to the next, you verify the correct functioning of a different component of the MPLS network and eliminate that layer as the source of the problem.

Physical Layer When you investigate the physical layer, you check that the routers are connected, and the interfaces are up and configured correctly. To check the physical layer, enter the **show interfaces**, **show interfaces terse**, and **ping** commands. If there is a problem in the physical layer, take appropriate action to fix it; then check that the LSP is operating as expected using the **show mpls lsp extensive** command. For more information on checking the physical layer, see [“Checklist for Verifying the Physical Layer” on page 93](#).

Data Link Layer When you investigate the data link layer, you check the encapsulation mode, for example, Point-to-Point Protocol (PPP) or Cisco High-Level Data Link Control (HDLC); PPP options, for example, header encapsulation; frame check sequence (FCS) size; and whether keepalive frames are enabled or disabled. To check the data link layer, enter the **show interfaces extensive** command. If there is a problem in the data link layer, take appropriate action to fix it; then check that the LSP is operating as expected using the **show mpls lsp extensive** command. For more information on checking the data link layer, see [“Checking the Data Link Layer” on page 104](#) and the *Junos Interfaces Operations Guide*.

IP Layer When you investigate the IP layer, you verify that interfaces have correct IP addressing, and that the interior gateway protocol (IGP) neighbor adjacencies are established. To check the IP layer, enter the **show interfaces terse**, **show ospf neighbor extensive**, and **show isis adjacency extensive** commands. If there is a problem in the IP layer, take appropriate action to fix it; then check that the LSP is operating as expected using the **show mpls lsp extensive** command.

IGP Layer When you investigate the IGP layer, you verify that the the Open Shortest Path First (OSPF) or Intermediate System-to-Intermediate System (IS-IS) protocols are configured and running correctly. For more information about configuring OSPF and IS-IS, see [“Configuring MPLS on Your Network” on page 6](#).

- If you have the OSPF protocol configured, you must check the IP layer first, and then the OSPF configuration. When you investigate the OSPF layer, you check that the protocol, interfaces, and traffic engineering are configured correctly. To check the OSPF layer, enter the **show configuration protocols ospf** and **show ospf interface** commands. If the problem exists in the OSPF layer, take appropriate action to fix it; then check that the LSP is operating as expected using the **show mpls lsp extensive** command. For more information about checking the OSPF layer, see [“Verifying the OSPF Protocol” on page 132](#).
- If you have the IS-IS protocol configured, because IS-IS and IP are independent of each other, it doesn't matter which one you check first. When you check the IS-IS configuration, you verify that IS-IS adjacencies are up, and the interfaces and IS-IS protocol are configured correctly. To check the IS-IS layer, enter the **show isis adjacency**, **show configuration protocols isis**, and **show isis interfaces** commands. If the problem exists in the IS-IS layer, take appropriate action to fix it; then check that the LSP is operating as expected using the **show mpls lsp extensive** command. For more information about checking the IS-IS layer, see [“Verifying the IS-IS Protocol” on page 144](#).



NOTE: The IS-IS protocol has traffic engineering enabled by default.

RSVP and MPLS Layers After you have both the IP and IGP layers functioning and the problem is still not solved, you can begin to check the Resource Reservation Protocol (RSVP) and MPLS layers to determine if the problem is in one of these layers.

- When you investigate the RSVP layer, you are checking that dynamic RSVP signaling is occurring as expected, neighbors are connected, and interfaces are configured correctly for RSVP. To check the RSVP layer, enter the **show rsvp session**, **show rsvp neighbor**, and **show rsvp interface** commands. If there is a problem in the RSVP layer, take appropriate action to fix it; then check that the LSP is operating as expected using the **show mpls lsp extensive** command.
- When you investigate the MPLS layer, you are checking whether the LSP is up and functioning correctly. To check the MPLS layer, enter the **show mpls lsp**, **show mpls lsp extensive**, **show route table mpls.0**, **show route address**, **traceroute address**, and **ping mpls rsvp lsp-name detail** commands. If there is a problem in the MPLS layer, take

appropriate action to fix it; then check that the LSP is operating as expected using the **show mpls lsp extensive** command.

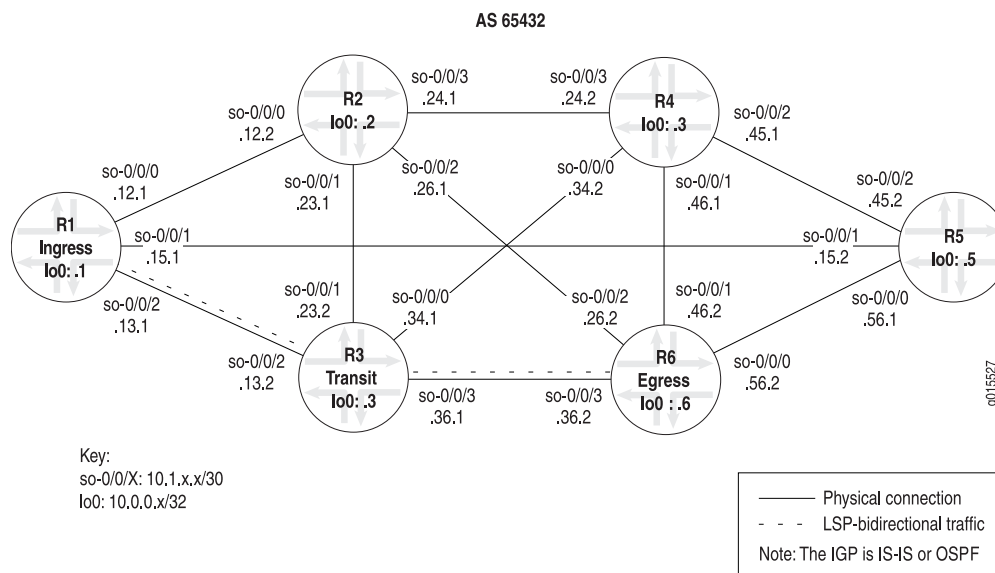
BGP Layer If the problem persists after you have checked the RSVP and MPLS layers, you must verify that the Border Gateway Protocol (BGP) is working correctly. There is no point in checking the BGP layer unless the LSP is established because BGP uses the MPLS LSP to forward traffic. When you check the BGP layer, you verify that the route is present and active, and more importantly, you ensure that the next hop is the LSP. To check the BGP layer, enter the **traceroute host-name**, **show bgp summary**, **show configuration protocols bgp**, **show route destination-prefix detail**, and **show route receive protocol bgp neighbor-address** commands. For more information on checking the BGP layer, see “Checking the BGP Layer” on page 194.

In reality, you could start at any level of the MPLS model to investigate a problem with your MPLS network. However, a disciplined approach, as the one described here, produces more consistent and reliable results.

Figure 9 on page 88 illustrates the basic network topology used in the following topics that demonstrate how to troubleshoot an MPLS network:

- Checklist for Verifying the Physical Layer on page 93
- Checklist for Checking the Data Link Layer on page 103
- Checklist for Verifying the IP and IGP Layers on page 117
- Checklist for Checking the RSVP Layer on page 155
- Checklist for Checking the MPLS Layer on page 173
- Checklist for Checking the BGP Layer on page 193

Figure 9: MPLS Basic Network Topology Example



The MPLS network consists of the following components:

- Router-only network with SONET interfaces
- MPLS protocol enabled on all routers, with interfaces selectively deactivated to illustrate a particular problem scenario
- All interfaces configured with MPLS
- A full-mesh IBGP topology, using AS 65432
- IS-IS or OSPF as the underlying IGP, using one level (IS-IS Level 2) or one area (OSPF area 0.0.0.0)
- A **send-statics** policy on routers R1 and R6, allowing a new route to be advertised into the network
- Two LSPs between routers R1 and R6, allowing for bidirectional traffic.

After you have configured an LSP, it is considered best practice to issue the **show mpls lsp** command to verify that the LSP is up, and to investigate further if you find an error message in the output. The error message can indicate a problem at any layer of the MPLS network.

The LSPs can be ingress, transit, or egress. Use the **show mpls lsp** command for quick verification of the LSP state, with the **extensive** option (**show mpls lsp extensive**) as a follow-up if the LSP is down. If your network has numerous LSPs, you might consider specifying the name of the LSP, using the **name** option (**show mpls lsp name *name*** or **show mpls lsp name *name* extensive**).

Action To begin the investigation of an error in your MPLS network, from the ingress router, enter some or all of the following Junos OS command-line interface (CLI) operational mode commands:

```
user@host> show mpls lsp
user@host> show mpls lsp extensive
user@host> show mpls lsp name name
user@host> show mpls lsp name name extensive
```

Sample Output 1 user@R1> show mpls lsp

Ingress LSP: 1 sessions

To	From	State	Rt	ActivePath	P	LSPName
10.0.0.6	10.0.0.1	Up	1		*	R1-to-R6

Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

To	From	State	Rt	Style	Labelin	Labelout	LSPName
10.0.0.1	10.0.0.6	Up	0	1 FF	3	-	R6-to-R1

Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions

Total 0 displayed, Up 0, Down 0

Sample Output 2 user@R1> show mpls lsp extensive

Ingress LSP: 1 sessions

10.0.0.6

From: 10.0.0.1, State: Up, ActiveRoute: 1, LSPName: R1-to-R6

ActivePath: (primary)

LoadBalance: Random

Encoding type: Packet, Switching type: Packet, GPID: IPv4

*Primary State: Up

Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)

10.1.13.2 S 10.1.36.2 S

Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

10.1.13.2 10.1.36.2

30 Dec 28 13:47:29 Selected as active path

29 Dec 28 13:47:29 Record Route: 10.1.13.2 10.1.36.2

28 Dec 28 13:47:29 Up

27 Dec 28 13:47:29 Originate Call

26 Dec 28 13:47:29 CSPF: computation result accepted

25 Dec 28 13:46:59 CSPF failed: no route toward 10.0.0.6

24 Dec 28 13:46:39 Deselected as active

23 Dec 28 13:46:39 CSPF failed: no route toward 10.0.0.6

22 Dec 28 13:46:39 Clear Call

21 Dec 28 13:46:39 ResvTear received

20 Dec 28 13:46:39 Down

19 Dec 28 13:46:39 10.1.13.2: Session preempted

18 Dec 28 13:42:07 Selected as active path

17 Dec 28 13:42:07 Record Route: 10.1.13.2 10.1.36.2

16 Dec 28 13:42:07 Up

15 Dec 28 13:42:07 Originate Call

14 Dec 28 13:42:07 CSPF: computation result accepted

13 Dec 28 13:41:37 CSPF failed: no route toward 10.0.0.6

12 Dec 28 13:41:16 Deselected as active

11 Dec 28 13:41:16 CSPF failed: no route toward 10.0.0.6

10 Dec 28 13:41:16 Clear Call

9 Dec 28 13:41:16 ResvTear received

8 Dec 28 13:41:16 Down

7 Dec 28 13:41:16 10.1.13.2: Session preempted

6 Dec 13 11:50:15 Selected as active path

5 Dec 13 11:50:15 Record Route: 10.1.13.2 10.1.36.2

4 Dec 13 11:50:15 Up

3 Dec 13 11:50:15 Originate Call

2 Dec 13 11:50:15 CSPF: computation result accepted

1 Dec 13 11:49:45 CSPF failed: no route toward 10.0.0.6[6 times]

---(more)---[abort]

Sample Output 3

```

user@R1> show mpls lsp name R1-to-R6
Ingress LSP: 1 sessions
To          From          State Rt ActivePath      P      LSPname
10.0.0.6    10.0.0.1    Up    1
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions
Total 0 displayed, Up 0, Down 0

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

```

Sample Output 4

```

user@R1> show mpls lsp name R1-to-R6 extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, State: Up, ActiveRoute: 1, LSPname: R1-to-R6
  ActivePath: (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary          State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
    10.1.13.2 S 10.1.36.2 S
      Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

          10.1.13.2 10.1.36.2
30 Dec 28 13:47:29 Selected as active path
29 Dec 28 13:47:29 Record Route: 10.1.13.2 10.1.36.2
28 Dec 28 13:47:29 Up
27 Dec 28 13:47:29 Originate Call
26 Dec 28 13:47:29 CSPF: computation result accepted
25 Dec 28 13:46:59 CSPF failed: no route toward 10.0.0.6
24 Dec 28 13:46:39 Deselected as active
23 Dec 28 13:46:39 CSPF failed: no route toward 10.0.0.6
22 Dec 28 13:46:39 Clear Call
21 Dec 28 13:46:39 ResvTear received
20 Dec 28 13:46:39 Down
19 Dec 28 13:46:39 10.1.13.2: Session preempted
18 Dec 28 13:42:07 Selected as active path
17 Dec 28 13:42:07 Record Route: 10.1.13.2 10.1.36.2
16 Dec 28 13:42:07 Up
15 Dec 28 13:42:07 Originate Call
14 Dec 28 13:42:07 CSPF: computation result accepted
13 Dec 28 13:41:37 CSPF failed: no route toward 10.0.0.6
12 Dec 28 13:41:16 Deselected as active
11 Dec 28 13:41:16 CSPF failed: no route toward 10.0.0.6
10 Dec 28 13:41:16 Clear Call
9 Dec 28 13:41:16 ResvTear received
8 Dec 28 13:41:16 Down
7 Dec 28 13:41:16 10.1.13.2: Session preempted
6 Dec 13 11:50:15 Selected as active path
5 Dec 13 11:50:15 Record Route: 10.1.13.2 10.1.36.2
4 Dec 13 11:50:15 Up
3 Dec 13 11:50:15 Originate Call
2 Dec 13 11:50:15 CSPF: computation result accepted
1 Dec 13 11:49:45 CSPF failed: no route toward 10.0.0.6[6 times]
Created: Mon Dec 13 11:47:19 2004
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions
Total 0 displayed, Up 0, Down 0

```

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

Meaning The sample output from the ingress router **R1** shows that the label-switched path is traversing the network as intended, from **R1** through **R3** to **R6**, and another LSP in the reverse direction, from **R6** through **R3** to **R1**.

If your network has numerous LSPs, you might consider using the **show mpls lsp** command for quick verification of the LSP state, and the **show mpls lsp name *name* extensive** command to continue your investigation if you find that the LSP is down.

For more information about the status and statistics of the **show mpls lsp** command, see [“Checklist for Determining LSP Status” on page 59](#). For more information on the availability and valid use of an LSP, see [“Checklist for Verifying LSP Use” on page 77](#).

In all the following topics, the network topology is broken at different layers of the network so that you can investigate various MPLS network problems. The problems presented are not inclusive. Instead, the problems serve to illustrate one possible process of investigation into the different layers of the troubleshooting model.

- Related Documentation**
- [Verifying the Physical Layer on page 94](#)
 - [Checking the Data Link Layer on page 104](#)
 - [Verifying the IP and IGP Layers on page 119](#)
 - [Checking the RSVP Layer on page 156](#)
 - [Checking the MPLS Layer on page 174](#)
 - [Checking the BGP Layer on page 194](#)

CHAPTER 7

Verifying the Physical Layer

This chapter describes how to investigate a problem at the physical layer of a Multiprotocol Label Switching (MPLS) network.

- [Checklist for Verifying the Physical Layer on page 93](#)
- [Verifying the Physical Layer on page 94](#)

Checklist for Verifying the Physical Layer

Problem This checklist provides the steps and commands for investigating a problem at the physical layer of a Multiprotocol Label Switching (MPLS) network. The checklist provides links to an overview of verifying the physical layer and more detailed information about the commands used to investigate the problem.

Solution [Table 11 on page 93](#) provides commands for verifying the physical layer.

Table 11: Checklist for Verifying the Physical Layer

Tasks	Command or Action
“Verifying the Physical Layer” on page 94	
1. Verify the LSP on page 95	<code>show mpls lsp extensive</code>
2. Verify Router Connection on page 98	<code>ping host</code>
3. Verify Interfaces on page 98	<code>show interfaces terse</code> <code>show configuration interfaces <i>type-fpc/pic/port</i></code>
4. Take Appropriate Action on page 99	The following sequence of commands addresses the specific problem described in this topic: <code>[edit interfaces <i>type-fpc/pic/port</i>]</code> <code>set family mpls</code> <code>show</code> <code>commit</code>
5. Verify the LSP Again on page 100	<code>show mpls lsp extensive</code>

Verifying the Physical Layer

Purpose After you have configured the LSP, issued the **show mpls lsp extensive** command, and determined that there is an error, you can start investigating the problem at the physical layer of the network.

Figure 10 on page 94 illustrates the physical layer of the layered MPLS model.

Figure 10: Verifying the Physical Layer


BGP Layer	traceroute <i>host-name</i> show bgp summary show configuration protocols bgp show route <i>destination-prefix</i> detail show route receive protocol bgp <i>neighbor-address</i>
MPLS Layer	show mpls lsp show mpls lsp extensive show route table mpls.0 show route <i>address</i> traceroute <i>address</i> ping mpls rsvp <i>Lsp-name</i> detail
RSVP Layer	show rsvp session show rsvp neighbor show rsvp interface
<div>↩</div> IGP and IP Layers Functioning <div>↪</div>	
OSPF Layer show ospf neighbor show configuration protocols ospf show ospf interface	IS-IS Layer show isis adjacency show configuration protocols isis show isis interface
IP Layer show ospf neighbor extensive show interfaces terse	IP Layer show isis adjacency extensive show interfaces terse
Data Link Layer	show interfaces extensive <i>"JUNOS Interfaces Operations Guide"</i>
Physical Layer	show interfaces show interfaces terse ping <i>host</i>

With this layer, you must ensure that the routers are connected, and that the interfaces are up and configured correctly on the ingress, egress, and transit routers.

If the network is not functioning at this layer, the label-switched path (LSP) does not work as configured.

Figure 11 on page 95 illustrates the MPLS network and the problem described in this topic.

Figure 11: MPLS Network Broken at the Physical Layer



The network shown in [Figure 11 on page 95](#) is a fully meshed configuration where every directly connected interface can receive and send packets to every other similar interface. The LSP in this network is configured to run from ingress router **R1**, through transit router **R3**, to egress router **R6**. In addition, a reverse LSP is configured to run from **R6** through **R3** to **R1**, creating bidirectional traffic.

However, in this example, traffic does not use the configured LSP. Instead traffic uses the alternate route from **R1** through **R2** to **R6**, and in the reverse direction, from **R6** through **R5** to **R1**.

When you become aware of a situation where an alternate route is used rather than the configured LSP, verify that the physical layer is functioning correctly. You might find that routers are not connected, or that interfaces are not up and configured correctly on the ingress, egress, or transit routers.

The cross shown in [Figure 11 on page 95](#) indicates where the LSP is broken because of a configuration error on ingress router **R1**.

To check the physical layer, follow these steps:

1. [Verify the LSP on page 95](#)
2. [Verify Router Connection on page 98](#)
3. [Verify Interfaces on page 98](#)
4. [Take Appropriate Action on page 99](#)
5. [Verify the LSP Again on page 100](#)

Verify the LSP

Purpose Typically, you use the `show mpls lsp extensive` command to verify the LSP. However, for quick verification of the LSP state, use the `show mpls lsp` command. If the LSP is down, use the `extensive` option (`show mpls lsp extensive`) as a follow-up. If your network has

numerous LSPs, you might consider specifying the name of the LSP, using the **name** option (**show mpls lsp name *name*** or **show mpls lsp name *name* extensive**).

Action To determine whether the LSP is up, enter the following command from the ingress router:

```
user@ingress-router> show mpls lsp extensive
```

Sample Output

```

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

10.0.0.6
  From: 10.0.0.1, State: Up, ActiveRoute: 1, LSPname: R1-to-R6
  ActivePath: (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
10.1.12.2 S 10.1.26.2 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

    10.1.12.2 10.1.26.2
99 Sep 18 14:19:04 CSPF: computation result accepted
98 Sep 18 14:19:04 CSPF: link down/deleted
10.1.13.1(R1.00/10.0.0.1)->10.1.13.2(R3.00/10.0.0.3)
97 Sep 18 14:19:01 Record Route: 10.1.12.2 10.1.26.2
96 Sep 18 14:19:01 Up
95 Sep 18 14:19:01 Clear Call
94 Sep 18 14:19:01 CSPF: computation result accepted
93 Sep 18 14:19:01 MPLS label allocation failure
92 Sep 18 14:19:01 Down
91 Aug 17 12:22:52 Selected as active path
90 Aug 17 12:22:52 Record Route: 10.1.13.2 10.1.36.2
89 Aug 17 12:22:52 Up
[...Output truncated...]
Created: Sat Jul 10 18:18:44 2004
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, LSPstate: Up, ActiveRoute: 0
  LSPname: R6-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: 144, Since: Tue Aug 17 12:23:14 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 39024 protocol 0
  PATH rcvfrom: 10.1.15.2 (so-0/0/1.0) 67333 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.1.56.2 10.1.15.2 <self>
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 the LSP is using an alternate path rather than the configured path. The configured path for the LSP is **R1** through **R3** to **R6**, and for the reverse LSP, **R6** through **R3** to **R1**. The alternate path used by the LSP is **R1** through **R2** to **R6**, and for the reverse LSP, **R6** through **R5** to **R1**.

Verify Router Connection

Purpose Confirm that the appropriate ingress, transit, and egress routers are functioning by examining whether the packets have been received and transmitted with 0% packet loss.

Action To determine that the routers are connected, enter the following command from the ingress and transit routers:

```
user@host> ping host
```

Sample Output

```
user@R1> ping 10.0.0.3 count 3
PING 10.0.0.3 (10.0.0.3): 56 data bytes
64 bytes from 10.0.0.3: icmp_seq=0 ttl=254 time=0.859 ms
64 bytes from 10.0.0.3: icmp_seq=1 ttl=254 time=0.746 ms
64 bytes from 10.0.0.3: icmp_seq=2 ttl=254 time=0.776 ms

--- 10.0.0.3 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.746/0.794/0.859/0.048 ms

user@R3> ping 10.0.0.6 count 3
PING 10.0.0.6 (10.0.0.6): 56 data bytes
64 bytes from 10.0.0.6: icmp_seq=0 ttl=255 time=0.968 ms
64 bytes from 10.0.0.6: icmp_seq=1 ttl=255 time=3.221 ms
64 bytes from 10.0.0.6: icmp_seq=2 ttl=255 time=0.749 ms

--- 10.0.0.6 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.749/1.646/3.221/1.117 ms
```

Meaning The sample output shows that ingress router **R1** is receiving packets from transit router **R3**, and that the transit router is receiving packets from the egress router. Therefore, the routers in the LSP are connected.

Verify Interfaces

Purpose Confirm that the interfaces are configured correctly with the **family mpls** statement.

Action To determine that the relevant interfaces are up and configured correctly, enter the following commands from the ingress, transit, and egress routers:

```
user@host> show interfaces terse
user@host> show configuration interfaces type-fpc/pic/port
```

Sample Output

```

user@R1> show interfaces so* terse
Interface      Admin Link Proto Local Remote
so-0/0/0       up   up   up   10.1.12.1/30
so-0/0/0.0      up   up   inet 10.1.12.1/30
                up   up   iso
                up   up   mpls

so-0/0/1       up   up   up   10.1.15.1/30
so-0/0/1.0      up   up   inet 10.1.15.1/30
                up   up   iso
                up   up   mpls

so-0/0/2       up   up   up   10.1.13.1/30
so-0/0/2.0      up   up   inet 10.1.13.1/30
                up   up   iso   <<< family mpls is missing

so-0/0/3       up   down

```

```

user@R1> show configuration interfaces so-0/0/2
unit 0 {
    family inet {
        address 10.1.13.1/30;
    }
    family iso; <<< family mpls is missing
}

```

Meaning The sample output shows that interface **so-0/0/2.0** on the ingress router does not have the **family mpls** statement configured at the **[edit interfaces type-fpc/pic/port]** hierarchy level, indicating that the interface is incorrectly configured to support the LSP. The LSP is configured correctly at the **[edit protocols mpls]** hierarchy level.

The output from the transit and egress routers (not shown) shows that the interfaces on those routers are configured correctly.

Take Appropriate Action

Problem Depending on the error you encountered in your investigation, you must take the appropriate action to correct the problem. In the example below, the **family mpls** statement, which was missing, is included in the configuration of ingress router **R1**.

Solution To correct the error in this example, enter the following commands:

```

[edit interfaces type-fpc/pic/port]
user@R1# set family mpls
user@R1# show
user@R1# commit

```

Sample Output

```
[edit interfaces so-0/0/2 unit 0]
user@R1# set family mpls

[edit interfaces so-0/0/2 unit 0]
user@R1# show
family inet {
    address 10.1.13.1/30;
}
family iso;
family mpls;

[edit interfaces so-0/0/2 unit 0]
user@R1# commit
commit complete
```

Meaning The sample output from ingress router **R1** shows that the **family mpls** statement is configured correctly for interface **so-0/0/2.0**, and that the LSP is now functioning as originally configured.

Verify the LSP Again

Purpose After taking the appropriate action to correct the error, the LSP needs to be checked again to confirm that the problem in the physical layer has been resolved.

Action To verify that the LSP is up and traversing the network as expected, enter the following command:

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


Sample Output 1

```

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

10.0.0.6
  From: 10.0.0.1, State: Up, ActiveRoute: 1, LSPname: R1-to-R6
  ActivePath: (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
  10.1.13.2 S 10.1.36.2 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

      10.1.13.2 10.1.36.2
    112 Sep 21 16:27:33 Record Route: 10.1.13.2 10.1.36.2
    111 Sep 21 16:27:33 Up
    110 Sep 21 16:27:33 CSPF: computation result accepted
    109 Sep 21 16:27:33 CSPF: link down/deleted
  10.1.12.1(R1.00/10.0.0.1)->10.1.12.2(R2.00/10.0.0.2)
    108 Sep 21 16:27:33 CSPF: link down/deleted
  10.1.15.1(R1.00/10.0.0.1)->10.1.15.2(R5.00/10.0.0.5)
    [Output truncated...]
    Created: Sat Jul 10 18:18:44 2004
  Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, LSPstate: Up, ActiveRoute: 0
  LSPname: R6-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: 149, Since: Tue Sep 21 16:29:43 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 2 receiver 39024 protocol 0
  PATH rcvfrom: 10.1.13.2 (so-0/0/2.0) 7 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.1.36.2 10.1.13.2 <self>
  Total 1 displayed, Up 1, Down 0

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

```

Sample Output 2

```

[edit protocols mpls]
user@R1# show
label-switched-path R1-to-R6 {
  to 10.0.0.6;
}
interface fxp0.0 {
  disable;
}
inactive: interface so-0/0/0.0;
inactive: interface so-0/0/1.0;

```

```
interface so-0/0/2.0;
```

Meaning Sample Output 1 from ingress router **R1** shows that the LSP is now traversing the network along the expected path, from **R1** through **R3** to **R6**, and the reverse LSP, from **R6** through **R3** to **R1**.

Sample Output 2 from ingress router **R1** shows that the LSP is forced to take the intended path because MPLS is deactivated on **R1** interfaces **so-0/0/0.0** and **so-0/0/1.0**. If these interfaces were not deactivated, even though the configuration is now correct, the LSP would still traverse the network through the alternate path.

CHAPTER 8

Checking the Data Link Layer

This chapter describes how to investigate a problem at the data link layer of the Multiprotocol Label Switching (MPLS) network.

- [Checklist for Checking the Data Link Layer on page 103](#)
- [Checking the Data Link Layer on page 104](#)

Checklist for Checking the Data Link Layer

Problem This checklist provides the steps and commands for investigating a problem at the data link layer of the Multiprotocol Label Switching (MPLS) network. The checklist provides links to an overview of the data link layer and more detailed information about the commands used to investigate the problem.

Solution [Table 12 on page 103](#) provides commands for checking the data link layer.

Table 12: Checklist for Checking the Data Link Layer

Tasks	Command or Action
“Checking the Data Link Layer” on page 104	
1. Verify the LSP on page 105	<code>show mpls lsp extensive</code>
2. Verify Interfaces on page 106	<code>show interfaces type-fpc/pic/port extensive</code> <code>show interfaces type-fpc/pic/port</code>
3. Take Appropriate Action on page 111	The following sequence of commands addresses the specific problem described in this topic: <code>[edit interfaces type-fpc/pic/port]</code> <code>show</code> <code>delete encapsulation</code> <code>show</code> <code>commit</code>
4. Verify the LSP Again on page 111	<code>show mpls lsp extensive</code>

Checking the Data Link Layer

Purpose After you have configured the label-switched path (LSP), issued the **show mpls lsp extensive** command, and determined that there is an error, you might find that the error is not in the physical layer. Continue investigating the problem at the data link layer of the network.

Figure 12 on page 104 illustrates the data link layer of the layered MPLS model.


Figure 12: Checking the Data Link Layer

BGP Layer	<tr> <td>MPLS Layer</td><td> <tr> <td>RSVP Layer</td><td> <tr> <td colspan="2"> <div>↙ IGP and IP Layers Functioning ↘</div> <tr> <td>OSPF Layer</td><td>IS-IS Layer</td> <tr> <td>IP Layer</td><td>IP Layer</td> <tr> <td>Data Link Layer</td><td> <tr> <td>Physical Layer</td><td> </td></tr> </td></tr> </tr></tr></td></tr></td></tr></td></tr>	MPLS Layer	<tr> <td>RSVP Layer</td><td> <tr> <td colspan="2"> <div>↙ IGP and IP Layers Functioning ↘</div> <tr> <td>OSPF Layer</td><td>IS-IS Layer</td> <tr> <td>IP Layer</td><td>IP Layer</td> <tr> <td>Data Link Layer</td><td> <tr> <td>Physical Layer</td><td> </td></tr> </td></tr> </tr></tr></td></tr></td></tr>	RSVP Layer	<tr> <td colspan="2"> <div>↙ IGP and IP Layers Functioning ↘</div> <tr> <td>OSPF Layer</td><td>IS-IS Layer</td> <tr> <td>IP Layer</td><td>IP Layer</td> <tr> <td>Data Link Layer</td><td> <tr> <td>Physical Layer</td><td> </td></tr> </td></tr> </tr></tr></td></tr>	<div>↙ IGP and IP Layers Functioning ↘</div> <tr> <td>OSPF Layer</td><td>IS-IS Layer</td> <tr> <td>IP Layer</td><td>IP Layer</td> <tr> <td>Data Link Layer</td><td> <tr> <td>Physical Layer</td><td> </td></tr> </td></tr> </tr></tr>		OSPF Layer	IS-IS Layer	IP Layer	IP Layer	Data Link Layer	<tr> <td>Physical Layer</td><td> </td></tr>	Physical Layer	
MPLS Layer	<tr> <td>RSVP Layer</td><td> <tr> <td colspan="2"> <div>↙ IGP and IP Layers Functioning ↘</div> <tr> <td>OSPF Layer</td><td>IS-IS Layer</td> <tr> <td>IP Layer</td><td>IP Layer</td> <tr> <td>Data Link Layer</td><td> <tr> <td>Physical Layer</td><td> </td></tr> </td></tr> </tr></tr></td></tr></td></tr>	RSVP Layer	<tr> <td colspan="2"> <div>↙ IGP and IP Layers Functioning ↘</div> <tr> <td>OSPF Layer</td><td>IS-IS Layer</td> <tr> <td>IP Layer</td><td>IP Layer</td> <tr> <td>Data Link Layer</td><td> <tr> <td>Physical Layer</td><td> </td></tr> </td></tr> </tr></tr></td></tr>	<div>↙ IGP and IP Layers Functioning ↘</div> <tr> <td>OSPF Layer</td><td>IS-IS Layer</td> <tr> <td>IP Layer</td><td>IP Layer</td> <tr> <td>Data Link Layer</td><td> <tr> <td>Physical Layer</td><td> </td></tr> </td></tr> </tr></tr>		OSPF Layer	IS-IS Layer	IP Layer	IP Layer	Data Link Layer	<tr> <td>Physical Layer</td><td> </td></tr>	Physical Layer			
RSVP Layer	<tr> <td colspan="2"> <div>↙ IGP and IP Layers Functioning ↘</div> <tr> <td>OSPF Layer</td><td>IS-IS Layer</td> <tr> <td>IP Layer</td><td>IP Layer</td> <tr> <td>Data Link Layer</td><td> <tr> <td>Physical Layer</td><td> </td></tr> </td></tr> </tr></tr></td></tr>	<div>↙ IGP and IP Layers Functioning ↘</div> <tr> <td>OSPF Layer</td><td>IS-IS Layer</td> <tr> <td>IP Layer</td><td>IP Layer</td> <tr> <td>Data Link Layer</td><td> <tr> <td>Physical Layer</td><td> </td></tr> </td></tr> </tr></tr>		OSPF Layer	IS-IS Layer	IP Layer	IP Layer	Data Link Layer	<tr> <td>Physical Layer</td><td> </td></tr>	Physical Layer					
<div>↙ IGP and IP Layers Functioning ↘</div> <tr> <td>OSPF Layer</td><td>IS-IS Layer</td> <tr> <td>IP Layer</td><td>IP Layer</td> <tr> <td>Data Link Layer</td><td> <tr> <td>Physical Layer</td><td> </td></tr> </td></tr> </tr></tr>		OSPF Layer	IS-IS Layer	IP Layer	IP Layer	Data Link Layer	<tr> <td>Physical Layer</td><td> </td></tr>	Physical Layer							
OSPF Layer	IS-IS Layer	IP Layer	IP Layer	Data Link Layer	<tr> <td>Physical Layer</td><td> </td></tr>	Physical Layer									
IP Layer	IP Layer	Data Link Layer	<tr> <td>Physical Layer</td><td> </td></tr>	Physical Layer											
Data Link Layer	<tr> <td>Physical Layer</td><td> </td></tr>	Physical Layer													
Physical Layer															

With this layer, you must check the encapsulation mode, for example, Point-to-Point Protocol (PPP) or Cisco High-Level Data Link Control (HDLC); PPP options, for example, header encapsulation; frame check sequence (FCS) size; and whether keepalive frames are enabled or disabled. Also, check the ingress, egress, and transit routers.

Figure 13 on page 105 illustrates the MPLS network used in this topic.

Figure 13: MPLS Network Broken at the Data Link Layer



The network shown in [Figure 13 on page 105](#) is a fully meshed configuration where every directly connected interface can receive and send packets to every other similar interface. The LSP in this network is configured to run from ingress router **R1**, through transit router **R3**, to egress router **R6**. In addition, a reverse LSP is configured to run from **R6** through **R3** to **R1**, creating bidirectional traffic.

However, in this example, the LSP is down without a path in either direction, from **R1** to **R6** or from **R6** to **R1**.

When you verify that the data link layer is not functioning correctly, you might find a mismatch with PPP or Cisco HDLC encapsulation, PPP options, or keepalive frames.

The cross shown in [Figure 13 on page 105](#) indicates where the LSP is broken because of a configuration error on ingress router **R1** that prevents the LSP from traversing the network as expected.

To check the data link layer, follow these steps:

1. [Verify the LSP on page 105](#)
2. [Verify Interfaces on page 106](#)
3. [Take Appropriate Action on page 111](#)
4. [Verify the LSP Again on page 111](#)

Verify the LSP

Purpose Typically, you use the **show mpls lsp extensive** command to verify the LSP. However for quick verification of the LSP state, use the **show mpls lsp** command. If the LSP is down, use the **extensive** option (**show mpls lsp extensive**) as a follow-up. If your network has numerous LSPs, you might consider specifying the name of the LSP, using the **name** option (**show mpls lsp name *name*** or **show mpls lsp name *name* extensive**).

Action To determine whether the LSP is up, enter the following command from the ingress router:

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

Sample Output 1

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

10.0.0.6
  From:10.0.0.1 , State: Dn, ActiveRoute: 0, LSPname: R1-to-R6
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary State: Dn
    Will be enqueued for recomputation in 15 second(s).
    140 Sep 30 12:01:12 CSPF failed: no route toward 10.0.0.6[26 times]
    139 Sep 30 11:48:57 Deselected as active
    138 Sep 30 11:48:56 CSPF failed: no route toward 10.0.0.6
    137 Sep 30 11:48:56 Clear Call
    136 Sep 30 11:48:56 CSPF: link down/deleted
    10.1.36.1(R3.00/10.0.0.3)->10.1.36.2(R6.00/10.0.0.6)
    135 Sep 30 11:48:56 ResvTear received
    134 Sep 30 11:48:56 Down
    133 Sep 30 11:48:56 CSPF failed: no route toward 10.0.0.6
    132 Sep 30 11:48:56 10.1.13.2: No Route toward dest
    [...Output truncated...]
  Created: Sat Jul 10 18:18:44 2004
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 from ingress router **R1** shows the LSPs within which it participates. The ingress LSP is down, without a path from **R1** to **R6**. Because a reverse LSP is configured in the network shown in [Figure 13 on page 105](#), we would expect an egress LSP session to be up. However, **R1** does not have any egress LSPs, indicating that the LSP from **R6** to **R1** is not functioning.

Verify Interfaces

Purpose From your network topology, determine the adjacent interfaces through which the LSP is meant to traverse, and examine the output for the encapsulation type, PPP options, FCS size, and whether keepalive frames are enabled or disabled



NOTE: Before you proceed with this step, check the physical layer to ensure that the problem is not in the physical layer.

Action To verify the functioning of adjacent interfaces, enter the following commands from the relevant routers:

```
user@host> show interfaces type-fpc/pic/port extensive
```

```
user@host> show interfaces type-fpc/pic/port
```

Sample Output 1

```

user@R6> show interfaces so-0/0/3 extensive
Physical interface: so-0/0/3, Enabled, Physical link is Up
  Interface index: 131, SNMP ifIndex: 27, Generation: 14
  Link-level type: Cisco-HDLC , MTU: 4474, Clocking: Internal, SONET mode, Speed:
OC3, Loopback: None,
  FCS:16 , Payload scrambler: Enabled
  Device flags   : Present Running
  Interface flags: Link-Layer-Down Point-To-Point SNMP-Traps 16384
  Link flags    : Keepalives
  Hold-times    : Up 0 ms, Down 0 ms
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive statistics:
    Input : 0 (last seen: never)
    Output: 357 (last sent 00:00:04 ago)
  CoS queues    : 4 supported
  Last flapped   : 2004-07-21 16:03:49 PDT (10w0d 07:01 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes   :          203368873          0 bps
    Output bytes  :          186714992         88 bps
    Input packets :           3641808          0 pps
    Output packets:           3297569          0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Bucket drops:
0,
    Policed discards: 1770, L3 incompletes: 0, L2 channel errors: 0, L2 mismatch
timeouts: 0,
    HS link CRC errors: 0, HS link FIFO overflows: 0
  Output errors:
    Carrier transitions: 1, Errors: 0, Drops: 0, Aged packets: 0, HS link FIFO
underflows: 0,
    MTU errors: 0
  Queue counters:      Queued packets  Transmitted packets      Dropped packets

    0 best-effort      197012                197012                0

    1 expedited-fo      0                      0                      0

    2 assured-forw      0                      0                      0

    3 network-cont     3100557               3100557               0

  SONET alarms : None
  SONET defects : None
  SONET PHY:
    Seconds      Count      State
    PLL Lock      0          0      OK
    PHY Light      0          0      OK
  SONET section:
    BIP-B1      0          0
    SEF          1          3      OK
    LOS          1          1      OK
    LOF          1          1      OK
    ES-S         1
    SES-S         1
    SEFS-S        1
  SONET line:
    BIP-B2      0          0
    REI-L       0          0

```



```

RDI-L          0          0 OK
AIS-L          0          0 OK
BERR-SF        0          0 OK
BERR-SD        0          0 OK
ES-L           1
SES-L           1
UAS-L           0
ES-LFE         0
SES-LFE         0
UAS-LFE         0
SONET path:
BIP-B3         0          0
REI-P          0          0
LOP-P          0          0 OK
AIS-P          0          0 OK
RDI-P          0          0 OK
UNEQ-P         0          0 OK
PLM-P          0          0 OK
ES-P           1
SES-P           1
UAS-P           0
ES-PFE         0
SES-PFE         0
UAS-PFE         0
Received SONET overhead:
F1      : 0x00, J0      : 0x00, K1      : 0x00, K2      : 0x00
S1      : 0x00, C2      : 0xcf, C2(cmp) : 0xcf, F2      : 0x00
Z3      : 0x00, Z4      : 0x00, S1(cmp) : 0x00
Transmitted SONET overhead:
F1      : 0x00, J0      : 0x01, K1      : 0x00, K2      : 0x00
S1      : 0x00, C2      : 0xcf, F2      : 0x00, Z3      : 0x00
Z4      : 0x00
Received path trace: R3 so-0/0/3
52 33 20 73 6f 2d 30 2f 30 2f 33 00 00 00 00 00 R3 so-0/0/3.. ...
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
00 00 00 00 00 00 00 00 00 00 00 00 00 00 0d 0a .....
Transmitted path trace: R6 so-0/0/3
52 36 20 73 6f 2d 30 2f 30 2f 33 00 00 00 00 00 R6 so-0/0/3 .....
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
HDLC configuration:
Policing bucket: Disabled
Shaping bucket : Disabled
Giant threshold: 4484, Runt threshold: 3
Packet Forwarding Engine configuration:
Destination slot: 0, PLP byte: 1 (0x00)
CoS transmit queue      Bandwidth      Buffer Priority  Limit
                        %      bps      %      bytes
0 best-effort            95    147744000 95      0      low  none
3 network-control        5     7776000  5      0      low  none

Logical interface so-0/0/3.0 (Index 71) (SNMP ifIndex 28) (Generation 16)
Flags: Device-Down Point-To-Point SNMP-Traps Encapsulation: Cisco-HDLC
Traffic statistics:
Input bytes : 406737746
Output bytes : 186714992
Input packets: 7283616
Output packets: 3297569
Local statistics:

```

```

Input bytes : 203368873
Output bytes : 186714992
Input packets: 3641808
Output packets: 3297569
Transit statistics:
Input bytes : 203368873 0 bps
Output bytes : 0 0 bps
Input packets: 3641808 0 pps
Output packets: 0 0 pps
Protocol inet, MTU: 4470, Generation: 46, Route table: 0
Flags: None
Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
Destination: 10.1.36.0/30, Local: 10.1.36.2, Broadcast: 10.1.36.3, Generation: 38
Protocol iso, MTU: 4469, Generation: 47, Route table: 0
Flags: None
Protocol mpls, MTU: 4458, Generation: 48, Route table: 0
Flags: None

```

Sample Output 2

```

user@R3> show interfaces so-0/0/3
Physical interface: so-0/0/3, Enabled, Physical link is Up
Interface index: 131, SNMP ifIndex: 24
Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC3,
Loopback: None, FCS: 16,
Payload scrambler: Enabled
Device flags : Present Running
Interface flags: Point-To-Point SNMP-Traps
Link flags : Keepalives
Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
Keepalive: Input: 736827 (00:00:03 ago), Output: 736972 (00:00:05 ago)
LCP state: Opened
NCP state: inet: Opened, inet6: Not-configured, iso: Opened, mpls: Opened
CHAP state: Not-configured
CoS queues : 4 supported
Last flapped : 2004-07-21 16:08:01 PDT (10w5d 19:57 ago)
Input rate : 40 bps (0 pps)
Output rate : 48 bps (0 pps)
SONET alarms : None
SONET defects : None

Logical interface so-0/0/3.0 (Index 70) (SNMP ifIndex 51)
Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
Protocol inet, MTU: 4470
Flags: None
Addresses, Flags: Is-Preferred Is-Primary
Destination: 10.1.36.0/30, Local: 10.1.36.1, Broadcast: 10.1.36.3
Protocol iso, MTU: 4470
Flags: None
Protocol mpls, MTU: 4458
Flags: None

```

Meaning Sample Output 1 from egress router **R6** shows that there are no SONET alarms or defects (**none**), the states are all **OK**, and the path trace shows the distant end (**R3 so-0.0.0**), indicating that the physical link is up. However, the logical link is down, and the link-level type is Cisco HDLC.

Sample Output 2 from transit router **R3** shows that the link-level type is PPP, indicating that the encapsulation types are mismatched, resulting in the LSP going down.

Take Appropriate Action

Problem Depending on the error you encountered in your investigation, you must take the appropriate action to correct the problem. In the example below, the encapsulation types are mismatched.

Solution To correct the error in this example, enter the following commands:

```
[edit interfaces so-0/0/3]
user@R1# show
user@R1# delete encapsulation
user@R1# show
user@R1# commit
```

Sample Output

```
[edit interfaces so-0/0/3]
user@R6# show
encapsulation cisco-hdlc;
unit 0 {
    family inet {
        address 10.1.36.2/30;
    }
    family iso;
    family mpls;
}

[edit interfaces so-0/0/3]
user@R6# delete encapsulation

[edit interfaces so-0/0/3]
user@R6# show
unit 0 {
    family inet {
        address 10.1.36.2/30;
    }
    family iso;
    family mpls;
}

[edit interfaces so-0/0/3]
user@R6# commit
commit complete
```

Meaning The sample output from egress router **R6** shows that the Cisco HDLC was incorrectly configured on interface **so-0/0/3** which prevented the LSP from using the intended path. The problem was corrected when the **encapsulation** statement was deleted and the configuration committed.

Verify the LSP Again

Purpose After taking the appropriate action to correct the error, the LSP needs to be checked again to confirm that the problem in the data link layer has been resolved.

Action From the ingress, egress, and transit routers, verify that the LSP is up and traversing the network as expected:

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

Sample Output 1

```

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

10.0.0.6
  From:10.0.0.1 , State: Up,  ActiveRoute:1, LSPname: R1-to-R6
  ActivePath: (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
  10.1.13.2 S 10.1.36.2 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

        10.1.13.2 10.1.36.2
145 Sep 30 12:25:01 Selected as active path
144 Sep 30 12:25:01 Record Route: 10.1.13.2 10.1.36.2
143 Sep 30 12:25:01 Up
142 Sep 30 12:25:01 Originate Call
141 Sep 30 12:25:01 CSPF: computation result accepted
140 Sep 30 12:24:32 CSPF failed: no route toward 10.0.0.6[74 times]
139 Sep 30 11:48:57 Deselected as active
138 Sep 30 11:48:56 CSPF failed: no route toward 10.0.0.6
137 Sep 30 11:48:56 Clear Call
136 Sep 30 11:48:56 CSPF: link down/deleted
10.1.36.1(R3.00/10.0.0.3)->10.1.36.2(R6.00/10.0.0.6)
[...Output truncated...]
Created: Sat Jul 10 18:18:43 2004
Total 1 displayed, Up 1 , Down 0

Egress LSP: 1 sessions

10.0.0.1
  From:10.0.0.6 , LSPstate: Up, ActiveRoute: 0
  LSPname: R6-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: 134, Since: Thu Sep 30 12:24:56 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 6 receiver 39024 protocol 0
  PATH rcvfrom: 10.1.13.2 (so-0/0/2.0) 7 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.1.36.2 10.1.13.2 <self>
Total 1 displayed, Up 1, Down 0

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

```

Sample Output 2

```

user@R6> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.1
  From:10.0.0.6, State: Up, ActiveRoute: 1, LSPname: R6-to-R1
  ActivePath: (primary)

```

```

LoadBalance: Random
Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary                               State: Up
  Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
10.1.36.1S10.1.13.1S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

    10.1.36.1 10.1.13.1
50 Sep 30 12:24:12 Selected as active path
49 Sep 30 12:24:12 Record Route: 10.1.36.1 10.1.13.1
48 Sep 30 12:24:12 Up
47 Sep 30 12:24:12 Originate Call
46 Sep 30 12:24:12 CSPF: computation result accepted
45 Sep 30 12:23:43 CSPF failed: no route toward 10.0.0.1[73 times]
44 Sep 30 11:48:12 Deselected as active
43 Sep 30 11:48:12 CSPF failed: no route toward 10.0.0.1
42 Sep 30 11:48:12 CSPF: link down/deleted
10.1.36.2(R6.00/10.0.0.6)->10.1.36.1(R3.00/10.0.0.3)
[...Output truncated...]
Created: Tue Aug 17 12:18:34 2004
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

10.0.0.6
From:10.0.0.1 , LSPstate: Up, ActiveRoute: 0
LSPname: R1-to-R6, 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: Thu Sep 30 12:24:16 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 19 receiver 44251 protocol 0
PATH rcvfrom: 10.1.36.1 (so-0/0/3.0) 4 pkts
Adspec: received MTU 1500
PATH sentto: localclient
RESV rcvfrom: localclient
Record route: 10.1.13.1 10.1.36.1 <self>
Total 1 displayed, Up 1, Down 0

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

```

Sample Output 3

```

user@R3> show mpls lsp extensive
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

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

Transit LSP: 2 sessions

10.0.0.1
From:10.0.0.6 , LSPstate: Up, ActiveRoute: 1
LSPname: R6-to-R1, LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: 3
Resv style: 1 FF, Label in: 100176, Label out: 3
Time left: 143, Since: Thu Sep 30 12:21:25 2004

```

```

Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 6 receiver 39024 protocol 0
PATH rcvfrom: 10.1.36.2 (so-0/0/3.0) 10 pkts
Adspec: received MTU 1500 sent MTU 1500
PATH sentto: 10.1.13.1 (so-0/0/2.0) 9 pkts
RESV rcvfrom: 10.1.13.1 (so-0/0/2.0) 9 pkts
Explct route: 10.1.13.1
Record route: 10.1.36.2 <self> 10.1.13.1

```

10.0.0.6

```

From:10.0.0.1 , LSPstate: Up, ActiveRoute: 1
LSPname: R1-to-R6, LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: 3
Resv style: 1 FF, Label in: 100192, Label out: 3
Time left: 148, Since: Thu Sep 30 12:21:30 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 19 receiver 44251 protocol 0
PATH rcvfrom: 10.1.13.1 (so-0/0/2.0) 9 pkts
Adspec: received MTU 1500 sent MTU 1500
PATH sentto: 10.1.36.2 (so-0/0/3.0) 9 pkts
RESV rcvfrom: 10.1.36.2 (so-0/0/3.0) 9 pkts
Explct route: 10.1.36.2
Record route: 10.1.13.1 <self> 10.1.36.2
Total 2 displayed, Up 2, Down 0

```

Sample Output 4

```

user@R1> show configuration protocols mpls
label-switched-path R1-to-R6 {
    to 10.0.0.6;
}
inactive: interface so-0/0/0.0;
inactive: interface so-0/0/1.0;
interface so-0/0/2.0;

user@R6> show configuration protocols mpls
label-switched-path R6-to-R1 {
    to 10.0.0.1;
}
inactive: interface so-0/0/0.0;
inactive: interface so-0/0/1.0;
inactive: interface so-0/0/2.0;
interface so-0/0/3.0;

user@R3> show configuration protocols mpls
interface fxp0.0 {
    disable;
}
inactive: interface so-0/0/0.0;
inactive: interface so-0/0/1.0;
interface so-0/0/2.0;
interface so-0/0/3.0;

```

Meaning Sample Outputs 1 and 2 from ingress router **R1** and egress router **R6**, respectively, show that the LSP is now traversing the network along the expected path, from **R1** through **R3** to **R6**, and the reverse LSP, from **R6** through **R3** to **R1**.

Sample Output 3 from transit router **R3** shows that there are two transit LSP sessions, one from **R1** to **R6** and the other from **R6** to **R1**.

Sample Output 4 shows the interfaces that were deactivated on the ingress, egress, and transit routers, forcing the LSP to take the intended path. If these interfaces were not deactivated, even though the configuration is now correct, the LSP would still traverse the network through the alternate path.

CHAPTER 9

Verifying the IP and IGP Layers

This chapter describes how to check the Internet Protocol (IP) and interior gateway protocol (IGP) layers of the layered Multiprotocol Label Switching (MPLS) model.

- [Checklist for Verifying the IP and IGP Layers on page 117](#)
- [Verifying the IP and IGP Layers on page 119](#)
- [Verifying the IP Layer on page 121](#)
- [Verifying the OSPF Protocol on page 132](#)
- [Verifying the IS-IS Protocol on page 144](#)

Checklist for Verifying the IP and IGP Layers

Problem This checklist provides the steps and commands for investigating a problem at the Internet Protocol (IP) and interior gateway protocol (IGP) layers of the layered Multiprotocol Label Switching (MPLS) model. The checklist provides links to an overview of the IP and IGP layers and more detailed information about the commands used to investigate the problem.

Solution [Table 13 on page 117](#) provides commands for verifying the IP and IGP layers.

Table 13: Checklist for Verifying the IP and IGP Layers

Tasks	Command or Action
“Verifying the IP and IGP Layers” on page 119	
“Verifying the IP Layer” on page 121	
1. Verify the LSP on page 122	<code>show mpls lsp extensive</code>
2. Verify IP Addressing on page 123	<code>show interfaces terse</code>
3. Verify Neighbors or Adjacencies at the IP Layer on page 125	<code>show ospf neighbor extensive</code> <code>show isis adjacency extensive</code>

Table 13: Checklist for Verifying the IP and IGP Layers (*continued*)

Tasks	Command or Action
4. Take Appropriate Action on page 128	<p>The following sequence of commands addresses the specific problem described in this topic:</p> <pre>[edit interfaces so-0/0/2] show rename unit 0 family inet address 10.1.13.2/30 to address 10.1.13.1/30 show commit</pre>
5. Verify the LSP Again on page 129	<code>show mpls lsp extensive</code>
“Verifying the OSPF Protocol” on page 132	
1. Verify the LSP on page 133	<code>show mpls lsp extensive</code>
2. Verify OSPF Interfaces on page 136	<code>show ospf interface</code>
3. Verify OSPF Neighbors on page 138	<code>show ospf neighbor</code>
4. Verify the OSPF Protocol Configuration on page 138	<code>show configuration protocols ospf</code>
5. Take Appropriate Action on page 139	<p>The following sequence of commands addresses the specific problem described in this topic:</p> <pre>[edit] edit protocols ospf area 0.0.0.0 [edit protocols ospf area 0.0.0.0] set interface lo0 set interface lo0 passive up [edit protocols ospf] set traffic-engineering show commit</pre>
6. Verify the LSP Again on page 141	<code>show mpls lsp extensive</code>
“Verifying the IS-IS Protocol” on page 144	
1. Verify the LSP on page 145	<code>show mpls lsp extensive</code>
2. Verify IS-IS Adjacencies and Interfaces on page 147	<pre>show isis adjacency show isis interface</pre>
3. Verify the IS-IS Configuration on page 149	<code>show configuration protocols isis</code>

Table 13: Checklist for Verifying the IP and IGP Layers (*continued*)

Tasks	Command or Action
4. Take Appropriate Action on page 150	<p>The following sequence of commands addresses the specific problem described in this topic:</p> <pre> edit [edit] edit protocols isis [edit protocols isis] show delete level 2 set level 1 disable show commit run show isis adjacency </pre>
5. Verify the LSP Again on page 151	<code>show mpls lsp extensive</code>

Verifying the IP and IGP Layers

Problem After you have configured the label-switched path (LSP), issued the **show mpls lsp extensive** command, and determined that there is an error, you might find that the error is not in the physical or data link layers. Continue investigating the problem at the IP and IGP layers of the network.

[Figure 14 on page 120](#) illustrates the IP and IGP layers of the layered MPLS model.

Figure 14: IP and IGP Layers

BGP Layer	tracroute <i>host-name</i> show bgp summary show configuration protocols bgp show route <i>destination-prefix</i> detail show route receive protocol bgp <i>neighbor-address</i>
MPLS Layer	show mpls lsp show mpls lsp extensive show route table mpls.0 show route <i>address</i> tracroute <i>address</i> ping mpls rsvp <i>lsp-name</i> detail
RSVP Layer	show rsvp session show rsvp neighbor show rsvp interface
↙ IGP and IP Layers Functioning ↘	
OSPF Layer	show ospf neighbor show configuration protocols ospf show ospf interface
IS-IS Layer	show isis adjacency show configuration protocols isis show isis interface
IP Layer	show ospf neighbor extensive show interfaces terse
IP Layer	show isis adjacency extensive show interfaces terse
Data Link Layer	show interfaces extensive "JUNOS Interfaces Operations Guide"
Physical Layer	show interfaces show interfaces terse ping <i>host</i>

g015545

Solution At the IP and IGP layers, you must check the following:

- Interfaces have correct IP addressing, and the IGP neighbors or adjacencies are established.
- Open Shortest Path First (OSPF) or Intermediate System-to-Intermediate System (IS-IS) protocols are configured and running correctly.
 - If the OSPF protocol is configured, check the IP layer first, then the OSPF configuration, making sure that the protocol, interfaces, and traffic engineering are configured correctly.
 - If the IS-IS protocol is configured, it doesn't matter whether you check IS-IS or IP first because both protocols are independent of each other. Verify that IS-IS adjacencies are up, and that the interfaces and IS-IS protocol are configured correctly.




NOTE: The IS-IS protocol has traffic engineering enabled by default.

If the network is not functioning at the IP or IGP layers, the LSP does not work as configured.

Figure 15 on page 121 illustrates the MPLS network used in this topic.

Figure 15: MPLS Network Broken at the IP and IGP Layers



The network shown in Figure 15 on page 121 is a fully meshed configuration where every directly connected interface can receive and send packets to every other similar interface. The LSP in this network is configured to run from ingress router R1, through transit router R3, to egress router R6. In addition, a reverse LSP is configured to run from R6, through R3, to R1, creating bidirectional traffic. The crosses in Figure 15 on page 121 indicate where the LSP is not working because of the following problems at the IP and IGP layer:

- An IP address is configured incorrectly on the ingress router (R1).
- The OSPF protocol is configured with a router ID (RID) but without the loopback (lo0) interface, and traffic engineering is missing from the transit router (R3).
- Levels in the IS-IS network are mismatched.

Related Documentation

To check the IP and IGP layers, follow these steps:

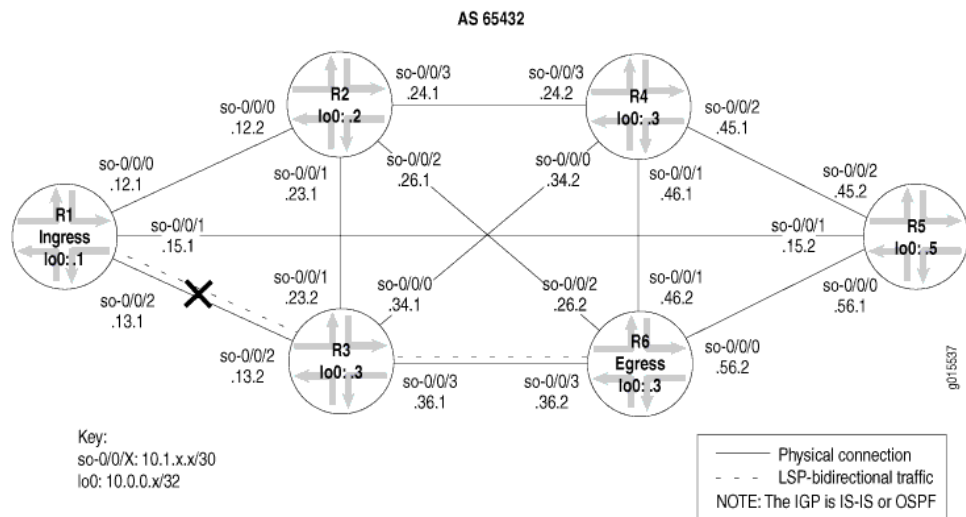
- [Verifying the IP Layer on page 121](#)
- [Verifying the OSPF Protocol on page 132](#)
- [Verifying the IS-IS Protocol on page 144](#)

Verifying the IP Layer

Purpose You can check the IP layer before or after you check the interior gateway protocol (IGP) layer, depending on whether you have OSPF or IS-IS configured as the IGP. If your MPLS network is configured with OSPF as the IGP, you must first verify the IP layer, checking that the interfaces have correct IP addressing and that the OSPF neighbors are established before you check the OSPF layer.

If you have IS-IS configured as the IGP in your MPLS network, you can verify either the IP layer or the IS-IS protocol layer first. The order in which you check the IP or IS-IS layer does not affect the results.

Figure 16: MPLS Network Broken at the IP Layer



The cross in Figure 16 on page 122 indicates where the LSP is broken because of the incorrect configuration of an IP address on ingress router R1.

1. [Verify the LSP on page 122](#)
2. [Verify IP Addressing on page 123](#)
3. [Verify Neighbors or Adjacencies at the IP Layer on page 125](#)
4. [Take Appropriate Action on page 128](#)
5. [Verify the LSP Again on page 129](#)

Verify the LSP

Purpose After configuring the LSP, you must verify that the LSP is up. LSPs can be ingress, transit, or egress. Use the **show mpls lsp** command for quick verification of the LSP state, with the **extensive** option (**show mpls lsp extensive**) as a follow-up if the LSP is down. If your network has numerous LSPs, you might consider specifying the name of the LSP, using the **name** option (**show mpls lsp name *name*** or **show mpls lsp name *name* extensive**).

Action To verify that the LSP is up, enter the following command from the ingress router:

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

Sample Output 1

```

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

10.0.0.6
  From: 10.0.0.1, State: Dn, ActiveRoute: 0, LSPname: R1-to-R6
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary State: Dn
    Will be enqueued for recomputation in 25 second(s).
  44 Oct 15 16:56:11 CSPF failed: no route toward 10.0.0.6 [2685 times]
  43 Oct 14 19:07:09 Clear Call
  42 Oct 14 19:06:56 Deselected as active
  41 Oct 14 19:06:56 10.1.12.1: MPLS label allocation failure
  40 Oct 14 19:06:56 Down
  39 Oct 14 18:43:43 Selected as active path
  38 Oct 14 18:43:43 Record Route: 10.1.13.2 10.1.36.2
  37 Oct 14 18:43:43 Up
  [...Output truncated...]
  Created: Thu Oct 14 16:04:33 2004
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 from ingress router **R1** shows that an MPLS label allocation failure occurred and the Constrained Shortest Path First (CSPF) algorithm failed, resulting in no route to destination **10.0.0.6** on **R6**.

Verify IP Addressing

Purpose When you investigate the IP layer, you verify that interfaces have correct IP addressing, and that OSPF neighbors or IS-IS adjacencies are established. In this example, an IP address is configured incorrectly on the ingress router (**R1**).

Action To verify IP addressing, enter the following command from the ingress, transit, and egress routers:

```
user@host> show interfaces terse
```

Sample Output

```

user@R1> show interfaces terse
Interface      Admin Link Proto Local Remote
so-0/0/0       up   up
so-0/0/0.0     up   up   inet 10.1.12.1/30
               up   up   iso
               up   up   mpls

so-0/0/1       up   up
so-0/0/1.0     up   up   inet 10.1.15.1/30
               up   up   iso
               up   up   mpls

so-0/0/2       up   up
so-0/0/2.0     up up inet 10.1.13.2 <<< Incorrect IP address
               up   up   iso
               up   up   mpls

lo0            up   up
lo0.0          up   up   inet 10.0.0.1
               up   up   iso 49.0004.1000.0000.0001.00

```

```

user@R3> show interfaces terse
Interface      Admin Link Proto Local Remote
so-0/0/0       up   up
so-0/0/0.0     up   up   inet 10.1.34.1/30
               up   up   iso
               up   up   mpls

so-0/0/1       up   up
so-0/0/1.0     up   up   inet 10.1.23.2/30
               up   up   iso
               up   up   mpls

so-0/0/2       up   up
so-0/0/2.0     up up inet 10.1.13.2/30 <<< Identical to R1
               up   up   iso
               up   up   mpls

so-0/0/3       up   up
so-0/0/3.0     up   up   inet 10.1.36.1/30
               up   up   iso
               up   up   mpls

lo0            up   up
lo0.0          up   up   inet 10.0.0.3
               up   up   iso 49.0004.1000.0000.0003.00

```

```

user@R6> show interfaces terse
Interface      Admin Link Proto Local Remote
so-0/0/0       up   up
so-0/0/0.0     up   up   inet 10.1.56.2/30
               up   up   iso
               up   up   mpls

so-0/0/1       up   up
so-0/0/1.0     up   up   inet 10.1.46.2/30
               up   up   iso
               up   up   mpls

so-0/0/2       up   up
so-0/0/2.0     up   up   inet 10.1.26.2/30
               up   up   iso
               up   up   mpls

so-0/0/3       up   up
so-0/0/3.0     up   up   inet 10.1.36.2/30
               up   up   iso
               up   up   mpls

```



```
100.0                up    up    inet 10.0.0.6
                      iso    49.0004.1000.0000.0006.00
```

Meaning The sample output shows that the IP addresses for interface **so-0/0/2.0** on **R1** and interface **so-0/0/2.0** on **R3** are identical. Interface IP addresses within a network must be unique for the interface to be identified correctly.

Verify Neighbors or Adjacencies at the IP Layer

Purpose If the IP addressing is configured incorrectly then the OSPF neighbors or IS-IS adjacencies both need to be checked to determine if one or both of them are established.

Action To verify neighbors (OSPF) or adjacencies (IS-IS), enter the following commands from the ingress, transit, and egress routers:

```
user@host> show ospf neighbor extensive
user@host> show isis adjacency extensive
```

Sample Output 1

```

user@R1> show ospf neighbor extensive
Address      Interface      State      ID              Pri  Dead
10.1.12.2    so-0/0/0.0     Full      10.0.0.2        128  34
  area 0.0.0.0, opt 0x42, DR 0.0.0.0, BDR 0.0.0.0
  Up 1d 04:45:20, adjacent 1d 04:45:20
10.1.15.2    so-0/0/1.0     Full      10.0.0.5        128  35
  area 0.0.0.0, opt 0x42, DR 0.0.0.0, BDR 0.0.0.0
  Up 1d 04:45:20, adjacent 1d 04:45:10 <<< no adjacency with R3 so-0/0/2

user@R3> show ospf neighbor extensive
Address      Interface      State      ID              Pri  Dead
10.1.23.1    so-0/0/1.0     Full      10.0.0.2        128  35
  area 0.0.0.0, opt 0x42, DR 0.0.0.0, BDR 0.0.0.0
  Up 1w2d 04:54:30, adjacent 1w2d 04:54:21
10.1.36.2    so-0/0/3.0     Full      10.0.0.6        128  39
  area 0.0.0.0, opt 0x42, DR 0.0.0.0, BDR 0.0.0.0
  Up 1w2d 04:54:30, adjacent 1w2d 04:54:30 <<< no adjacency with R1 so-0/0/2

user@R6> show ospf neighbor extensive
Address      Interface      State      ID              Pri  Dead
10.1.56.1    so-0/0/0.0     Full      10.0.0.5        128  39
  area 0.0.0.0, opt 0x42, DR 0.0.0.0, BDR 0.0.0.0
  Up 1d 02:59:35, adjacent 1d 02:59:35
10.1.26.1    so-0/0/2.0     Full      10.0.0.2        128  36
  area 0.0.0.0, opt 0x42, DR 0.0.0.0, BDR 0.0.0.0
  Up 1w2d 04:57:30, adjacent 1w2d 04:57:30
10.1.36.1    so-0/0/3.0     Full      10.0.0.3        128  36
  area 0.0.0.0, opt 0x42, DR 0.0.0.0, BDR 0.0.0.0
  Up 1w2d 04:56:11, adjacent 1w2d 04:56:11

```

Sample Output 2

```

user@R1> show isis adjacency extensive
R2
  Interface: so-0/0/0.0, Level:2, State:Up , Expires in 23 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 05:57:16 ago
  Circuit type: 2, Speaks:IP , IPv6
  Topologies: Unicast
  Restart capable: Yes
  IP addresses: 10.1.12.2
  Transition log:
    When              State      Reason
    Fri Oct 15 14:58:35 Up          Seenself

R5
  Interface: so-0/0/1.0, Level:2, State:Up, Expires in 26 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 05:56:52 ago
  Circuit type: 2, Speaks:IP , IPv6
  Topologies: Unicast
  Restart capable: Yes
  IP addresses: 10.1.15.2
  Transition log:
    When              State      Reason
    Fri Oct 15 14:59:00 Up          Seenself

R3
  Interface: so-0/0/2.0, Level: 2, State: Up, Expires in 26 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 05:56:51 ago

```

```

Circuit type: 2, Speaks:IP , IPv6
Topologies: Unicast
Restart capable: Yes
IP addresses:10.1.13.2
Transition log:
When                State      Reason
Fri Oct 15 14:59:01  Up        Seenself

```

user@R3> show isis adjacency extensive

```

R4
  Interface: so-0/0/0.0, Level:2, State:Up , Expires in 25 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 1w1d 00:22:51 ago
  Circuit type: 2, Speaks:IP , IPv6
  Topologies: Unicast
  Restart capable: Yes
  IP addresses:10.1.34.2
  Transition log:
When                State      Reason
Thu Oct 28 15:13:12  Up        Seenself

```

```

R2
  Interface: so-0/0/1.0, Level:2, State:Up , Expires in 25 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 2w2d 18:02:48 ago
  Circuit type: 2, Speaks:IP , IPv6
  Topologies: Unicast
  Restart capable: Yes
  IP addresses:10.1.23.1
  Transition log:
When                State      Reason
Tue Oct 19 21:33:15  Up        Seenself

```

```

R1
  Interface: so-0/0/2.0, Level:2, State:Up , Expires in 22 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 2w2d 17:24:06 ago
  Circuit type: 2, Speaks:IP , IPv6
  Topologies: Unicast
  Restart capable: Yes
  IP addresses:10.1.13.1
  Transition log:
When                State      Reason
Tue Oct 19 22:11:57  Up        Seenself

```

```

R6
  Interface: so-0/0/3.0, Level:2, State:Up , Expires in 21 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 2w1d 00:07:00 ago
  Circuit type: 2, Speaks:IP , IPv6
  Topologies: Unicast
  Restart capable: Yes
  IP addresses:10.1.36.2
  Transition log:
When                State      Reason
Thu Oct 21 15:29:03  Up        Seenself

```

user@R6> show isis adjacency extensive

```

R5
  Interface: so-0/0/0.0, Level:2, State:Up , Expires in 23 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 1w2d 01:10:03 ago
  Circuit type: 2, Speaks:IP , IPv6
  Topologies: Unicast
  Restart capable: Yes
  IP addresses:10.1.56.1

```

```
Transition log:
When          State      Reason
Wed Oct 27 14:35:32 Up        SeenseIf
```

R4

```
Interface: so-0/0/1.0, Level:2, State:Up , Expires in 25 secs
Priority: 0, Up/Down transitions: 1, Last transition: 1w1d 00:26:50 ago
Circuit type: 2, Speaks:IP , IPv6
Topologies: Unicast
Restart capable: Yes
IP addresses: 10.1.46.1
Transition log:
When          State      Reason
Thu Oct 28 15:18:45 Up        SeenseIf
```

R2

```
Interface: so-0/0/2.0, Level:2, State:Up , Expires in 24 secs
Priority: 0, Up/Down transitions: 1, Last transition: 2w1d 00:11:40 ago
Circuit type: 2, Speaks:IP , IPv6
Topologies: Unicast
Restart capable: Yes
IP addresses: 10.1.26.1
Transition log:
When          State      Reason
Thu Oct 21 15:33:55 Up        SeenseIf
```

R3

```
Interface: so-0/0/3.0, Level:2, State:Up , Expires in 19 secs
Priority: 0, Up/Down transitions: 1, Last transition: 2w1d 00:11:40 ago
Circuit type: 2, Speaks:IP , IPv6
Topologies: Unicast
Restart capable: Yes
IP addresses: 10.1.36.1
Transition log:
When          State      Reason
Thu Oct 21 15:33:55 Up        SeenseIf
```

Meaning Sample Output 1 from the ingress, transit, and egress routers shows that **R1** and **R3** are not established OSPF neighbors. Considering that the two interfaces **so-0/0/2.0** (**R1** and **R3**) are configured with identical IP addresses, you would expect this. The OSPF protocol routes IP packets based solely on the destination IP address contained in the IP packet header. Therefore, identical IP addresses in the autonomous system (AS) result in neighbors not establishing.

Sample Output 2 from the ingress, transit, and egress routers shows that **R1** and **R3** have established an IS-IS adjacency despite the identical IP addresses configured on interfaces **so-0/0/2.0** on **R1** and **R3**. The IS-IS protocol behaves differently from the OSPF protocol because it does not rely on IP to establish an adjacency. However, if the LSP is not up, it is still useful to check the IP subnet addressing in case there is a mistake in that layer. Correcting the addressing error might bring the LSP back up.

Take Appropriate Action

Problem Depending on the error you encountered in your investigation, you must take the appropriate action to correct the problem. In this example, the IP address of an interface on transit router **R2** is incorrectly configured.

Solution To correct the error in this example, enter the following commands:

```
[edit interfaces so-0/0/2]
user@R1# show
user@R1# rename unit 0 family inet address 10.1.13.2/30 to address 10.1.13.1/30
user@R1# show
user@R1# commit
```

Sample Output

```
[edit interfaces so-0/0/2]
user@R1# show
unit 0 {
    family inet {
        address 10.1.13.2/30; <<< Incorrect IP address
    }
    family iso;
    family mpls;
}

[edit interfaces so-0/0/2]
user@R1# rename unit 0 family inet address 10.1.13.2/30 to address 10.1.13.1/30

[edit interfaces so-0/0/2]
user@R1# show
unit 0 {
    family inet {
        address 10.1.13.1/30; <<< Correct IP address.
    }
    family iso;
    family mpls;
}

[edit interfaces so-0/0/2]
user@R1# commit
commit complete
```

Meaning The sample output shows that interface **so-0/0/2** on ingress router **R1** is now configured with the correct IP address. This correction results in unique subnet IP addresses for all interfaces in the MPLS network in [Figure 15 on page 121](#), and the possibility that the LSP might come up.

Verify the LSP Again

Purpose After taking the appropriate action to correct the error, the LSP needs to be checked again to confirm that the problem in the OSPF protocol has been resolved.

Action To verify the LSP again, enter the following command on the ingress, transit, and egress routers:

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

Sample Output 1

```

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

10.0.0.6
  From: 10.0.0.1, State: Up,  ActiveRoute:1 , LSPname: R1-to-R6
  ActivePath: (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State:Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
10.1.13.2 S 10.1.36.2 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

    10.1.13.2 10.1.36.2
54 Oct 15 21:28:16 Selected as active path
53 Oct 15 21:28:16 Record Route:  10.1.13.2 10.1.36.2
52 Oct 15 21:28:16 Up
51 Oct 15 21:28:16 10.1.15.1: MPLS label allocation failure[2 times]
50 Oct 15 21:28:11 CSPF: computation result accepted
49 Oct 15 21:27:42 10.1.15.1: MPLS label allocation failure
48 Oct 15 21:27:42 CSPF: computation result accepted
47 Oct 15 21:27:31 10.1.15.1: MPLS label allocation failure[4 times]
46 Oct 15 21:27:13 Originate Call
45 Oct 15 21:27:13 CSPF: computation result accepted
[...Output truncated...]
Created: Thu Oct 14 16:04:34 2004
Total 1 displayed,  Up1 , Down 0

Egress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6,  LSPstate:Up , ActiveRoute: 0
  LSPname: R6-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: 149, Since: Fri Oct 15 21:28:13 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 13 receiver 39024 protocol 0
  PATH rcvfrom: 10.1.13.2 (so-0/0/2.0) 10 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.1.36.2 10.1.13.2 <self>
Total 1 displayed,  Up1 , Down 0

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

```

Sample Output 2

```

user@R3> show mpls lsp extensive
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

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

```

```

Transit LSP: 2 sessions

10.0.0.1
  From: 10.0.0.6,  LSPstate: Up , ActiveRoute: 1
  LSPname: R6-to-R1 , LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: 100336, Label out: 3
  Time left: 156, Since: Fri Oct 15 21:15:47 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 13 receiver 39024 protocol 0
  PATH rcvfrom: 10.1.36.2 (so-0/0/3.0) 11 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 10.1.13.1 (so-0/0/2.0) 11 pkts
  RESV rcvfrom: 10.1.13.1 (so-0/0/2.0) 11 pkts
  Explct route: 10.1.13.1
  Record route: 10.1.36.2 <self> 10.1.13.1

10.0.0.6
  From: 10.0.0.1,  LSPstate: Up , ActiveRoute: 1
  LSPname: R1-to-R6 , 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: 159, Since: Fri Oct 15 21:15:50 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 5 receiver 47901 protocol 0
  PATH rcvfrom: 10.1.13.1 (so-0/0/2.0) 11 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 10.1.36.2 (so-0/0/3.0) 11 pkts
  RESV rcvfrom: 10.1.36.2 (so-0/0/3.0) 11 pkts
  Explct route: 10.1.36.2
  Record route: 10.1.13.1 <self> 10.1.36.2
Total 2 displayed,  Up 2 , Down 0

```

Sample Output 3

```

user@R6> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6,  State: Up , ActiveRoute: 1,  LSPname: R6-to-R1
  ActivePath: (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary          State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
  10.1.36.1 S 10.1.13.1 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

      10.1.36.1 10.1.13.1
    187 Oct 15 21:20:05 Selected as active path
    186 Oct 15 21:20:05 Record Route: 10.1.36.1 10.1.13.1
    185 Oct 15 21:20:05 Up
    184 Oct 15 21:20:05 Clear Call
    183 Oct 15 21:20:05 CSPF: computation result accepted
    182 Oct 15 21:20:05 CSPF: link down/deleted
  10.1.13.2(R3.00/10.0.0.3)->10.1.13.2(R1.00/10.0.0.1)
  [...Output truncated...]
  Created: Tue Aug 17 12:18:33 2004
Total 1 displayed,  Up 1 , Down 0

```

Egress LSP: 1 sessions

10.0.0.6

```
From: 10.0.0.1, LSPstate: Up, ActiveRoute: 0
  LSPname: R1-to-R6 , 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: 144, Since: Fri Oct 15 21:20:08 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 5 receiver 47901 protocol 0
  PATH rcvfrom: 10.1.36.1 (so-0/0/3.0) 11 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.1.13.1 10.1.36.1 <self>
Total 1 displayed, Up 1 , Down 0
```

Transit LSP: 0 sessions

Total 0 displayed, Up 0, Down 0

Meaning Sample Output 1 from ingress router **R1** shows that LSP **R1-to-R6** has an active route to **R6** and the state is up. The output shows that the egress LSP session **R6-to-R1** received and sent a recovery label.

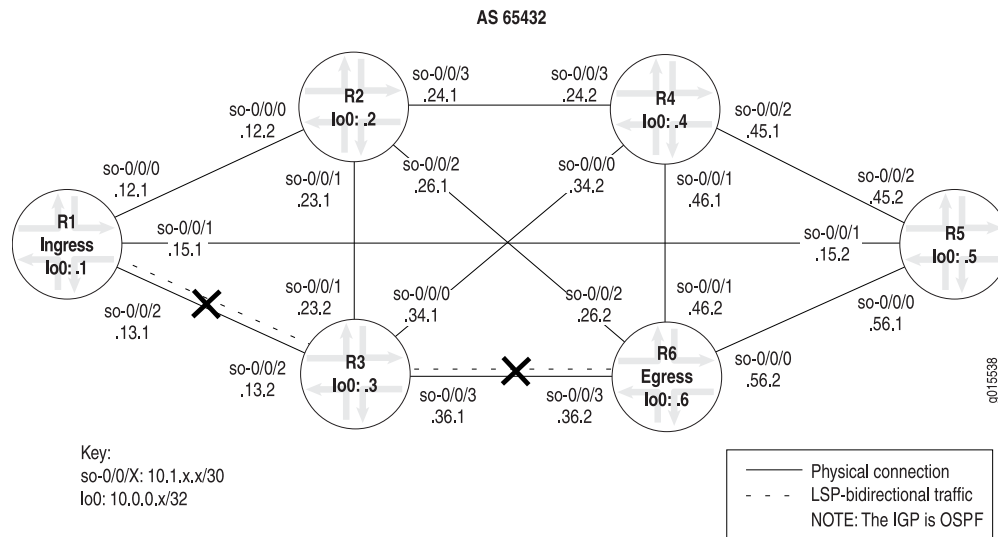
Sample Output 2 from transit router **R3** shows that there are two transit LSP sessions, one from **R1** to **R6** and the other from **R6** to **R1**. Both LSPs are up.

Sample Output 3 from egress router **R6** shows that the LSP is up and the active route is the primary route. The LSP is now traversing the network along the expected path, from **R1** through **R3** to **R6**, and the reverse LSP, from **R6** through **R3** to **R1**.

Verifying the OSPF Protocol

Purpose After you have verified that the LSP is down, and the cause is not in the physical, datalink, or IP layer, verify the OSPF configuration. Check the routers in your network to ensure that the interfaces and the OSPF protocol are configured correctly, and that the neighbors are established.

Figure 17: MPLS Network Broken at the OSPF Protocol Layer



1. [Verify the LSP on page 133](#)
2. [Verify OSPF Interfaces on page 136](#)
3. [Verify OSPF Neighbors on page 138](#)
4. [Verify the OSPF Protocol Configuration on page 138](#)
5. [Take Appropriate Action on page 139](#)
6. [Verify the LSP Again on page 141](#)

Verify the LSP

Purpose Confirm that interfaces are configured for OSPF, the OSPF protocol is configured correctly and that neighbors are established.

Action To verify the LSP, enter the following command on the ingress, transit, and egress routers:

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

Sample Output 1

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

10.0.0.6
  From: 10.0.0.1, State: Dn, ActiveRoute: 0,  LSPname: R1-to-R6
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary                               State: Dn
    11 Oct 19 18:06:04 No Route toward dest[78 times]
    10 Oct 19 17:08:09 Deselected as active
  Created: Mon Oct 18 21:48:42 2004
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
```

Sample Output 2

```
user@R3> show mpls lsp extensive
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

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

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

Sample Output 3

```
user@R6> show mpls lsp extensive
Ingress LSP: 1 sessions

```

To	From	State	Rt	ActivePath	P	LSPname
10.0.0.1	10.0.0.6	Dn	0	-		R6-to-R1

```
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
```

Sample Output 4

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

10.0.0.6
  From: 10.0.0.1, State: Up, ActiveRoute: 1,  LSPname: R1-to-R6
  ActivePath: (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
```

```

Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

      10.1.13.2 10.1.36.2
    5 Oct 19 10:37:55 Selected as active path
    4 Oct 19 10:37:55 Record Route: 10.1.13.2 10.1.36.2
    3 Oct 19 10:37:55 Up
    2 Oct 19 10:37:10 No Route toward dest[1029 times]
    1 Oct 18 21:48:42 Originate Call
Created: Mon Oct 18 21:48:42 2004
Total 1 displayed, Up1 , Down 0

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

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

```

Sample Output 5

```

user@R3> show mpls lsp extensive
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

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

Transit LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, LSPstate: Up, ActiveRoute: 1
  LSPname: R1-to-R6 , LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: 100368, Label out: 3
  Time left: 154, Since: Tue Oct 19 10:25:24 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 47933 protocol 0
  PATH rcvfrom: 10.1.13.1 (so-0/0/2.0) 209 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 10.1.36.2 (so-0/0/3.0) 209 pkts
  RESV rcvfrom: 10.1.36.2 (so-0/0/3.0) 209 pkts
  Record route: 10.1.13.1 <self> 10.1.36.2
Total 1 displayed, Up 1, Down 0

```

Sample Output 6

```

user@R6> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, State: Dn, ActiveRoute: 0, LSPname: R6-to-R1
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary State: Dn
    2 Oct 19 13:01:54 10.1.56.2: MPLS label allocation failure [9 times]
    1 Oct 19 12:57:51 Originate Call
Created: Tue Oct 19 12:57:51 2004
Total 1 displayed, Up 0, Down1

```

Egress LSP: 1 sessions

10.0.0.6

```
From: 10.0.0.1, LSPstate: Up, ActiveRoute: 0
LSPname: R1-to-R6, LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: -
Resv style: 1 FF, Label in: 3, Label out: -
Time left: 148, Since: Tue Oct 19 10:30:03 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 1 receiver 47933 protocol 0
PATH rcvfrom: 10.1.36.1 (so-0/0/3.0) 206 pkts
Adspec: received MTU 1500
PATH sentto: localclient
RESV rcvfrom: localclient
Record route: 10.1.13.1 10.1.36.1 <self>
Total 1 displayed, Up 1, Down 0
```

Transit LSP: 0 sessions

Total 0 displayed, Up 0, Down 0

Meaning Sample Outputs 1, 2, and 3 show that the LSP and the reverse LSP are down:

- Sample Output 1 from ingress router **R1** shows that LSP **R1-to-R6** does not have a route towards the destination (**R6**).
- Sample Output 2 from transit router **R3** shows that there are no LSP sessions.
- Sample Output 3 from egress router **R6** also shows that reverse LSP **R6-to-R1** is down.

Sample Outputs 4, 5, and 6 show that the LSP is up and the reverse LSP is down:

- Sample Output 4 from ingress router **R1** shows that LSP **R1-to-R6** is up and there are no egress LSP sessions.
- Sample Output 5 from transit router **R3** shows that there is one ingress LSP session (**R1-to-R6**) and no egress LSP sessions.
- Sample Output 6 from egress router **R6** shows that LSP **R6-to-R1** is down due to an MPLS label allocation failure.

Verify OSPF Interfaces

Purpose After you have verified that the LSP is down, and the cause is not in the physical, data link, or IP layer, check the routers in your network to determine that all relevant OSPF interfaces are configured correctly.

Action To verify OSPF interfaces, enter the following commands from the ingress, transit, and egress routers:

```
user@host> show ospf interface
```

Sample Output 1

```
user@R1> show ospf interface
Interface      State   Area      DR ID      BDR ID      Nbrs
so-0/0/0.0     PtToPt  0.0.0.0   0.0.0.0    0.0.0.0     1
so-0/0/1.0     PtToPt  0.0.0.0   0.0.0.0    0.0.0.0     1
so-0/0/2.0     PtToPt  0.0.0.0   0.0.0.0    0.0.0.0     1
```

```
user@R3> show ospf interface
Interface      State   Area      DR ID      BDR ID      Nbrs
so-0/0/0.0     PtToPt  0.0.0.0   0.0.0.0    0.0.0.0     1
so-0/0/1.0     PtToPt  0.0.0.0   0.0.0.0    0.0.0.0     1
so-0/0/2.0     PtToPt  0.0.0.0   0.0.0.0    0.0.0.0     1
so-0/0/3.0     PtToPt  0.0.0.0   0.0.0.0    0.0.0.0     1
```

```
user@R6> show ospf interface
Interface      State   Area      DR ID      BDR ID      Nbrs
so-0/0/0.0     PtToPt  0.0.0.0   0.0.0.0    0.0.0.0     1
so-0/0/1.0     PtToPt  0.0.0.0   0.0.0.0    0.0.0.0     1
so-0/0/2.0     PtToPt  0.0.0.0   0.0.0.0    0.0.0.0     1
so-0/0/3.0     PtToPt  0.0.0.0   0.0.0.0    0.0.0.0     1
```

Sample Output 2

```
user@R1> show ospf interface
Interface      State   Area      DR ID      BDR ID      Nbrs
lo0.0          DR      0.0.0.0   10.0.0.1   0.0.0.0     0
so-0/0/0.0     PtToPt  0.0.0.0   0.0.0.0    0.0.0.0     1
so-0/0/1.0     PtToPt  0.0.0.0   0.0.0.0    0.0.0.0     1
so-0/0/2.0     PtToPt  0.0.0.0   0.0.0.0    0.0.0.0     1
```

```
user@R3> show ospf interface
Interface      State   Area      DR ID      BDR ID      Nbrs
lo0.0          DR      0.0.0.0   10.0.0.3   0.0.0.0     0
so-0/0/0.0     Down    0.0.0.0   0.0.0.0    0.0.0.0     0
so-0/0/1.0     PtToPt  0.0.0.0   0.0.0.0    0.0.0.0     1
so-0/0/2.0     PtToPt  0.0.0.0   0.0.0.0    0.0.0.0     1
so-0/0/3.0     PtToPt  0.0.0.0   0.0.0.0    0.0.0.0     1
```

```
user@R6> show ospf interface
Interface      State   Area      DR ID      BDR ID      Nbrs
lo0.0          DR      0.0.0.0   10.0.0.6   0.0.0.0     0
so-0/0/0.0     PtToPt  0.0.0.0   0.0.0.0    0.0.0.0     1
so-0/0/1.0     Down    0.0.0.0   0.0.0.0    0.0.0.0     0
so-0/0/2.0     PtToPt  0.0.0.0   0.0.0.0    0.0.0.0     1
so-0/0/3.0     PtToPt  0.0.0.0   0.0.0.0    0.0.0.0     1
```

Meaning Sample Output 1 shows that all interfaces on all routers are in the correct area (0.0.0.0), and the loopback (lo0) interface is missing from the list of interfaces on all routers. The missing loopback (lo0) interface is a problem in this configuration.

In an MPLS network configured with OSPF as the IGP, when you manually configure the RID, it is important to explicitly configure the loopback interface at the `[edit protocols ospf]` hierarchy level. If the RID is not manually configured, OSPF automatically advertises the loopback (lo0) interface. In the configuration of all the routers in this network, the RID is configured manually, therefore, the loopback (lo0) interface must be explicitly configured at the `[edit protocols ospf]` hierarchy level. In addition, the loopback (lo0)

interface is configured with the **passive** statement to ensure that the protocols are not run over the loopback (**lo0**) interface and it is correctly advertised throughout the network.

Sample Output 2 shows that all the relevant interfaces on the ingress, transit, and egress routers, including the loopback (**lo0**) interface, are in the correct area (**0.0.0.0**). Because the configuration of the interfaces is correct, further investigation is required to determine the reason for the LSP problem.

Verify OSPF Neighbors

Purpose After you have checked OSPF interfaces, check your network topology to determine that all relevant neighbors are established.

Action To verify OSPF neighbors, enter the following commands from the ingress, transit, and egress routers:

```
user@host> show ospf neighbor
```

Sample Output

```
user@R1> show ospf neighbor
  Address      Interface      State      ID              Pri  Dead
  10.1.12.2     so-0/0/0.0     Full      10.0.0.2        128  39
  10.1.15.2     so-0/0/1.0     Full      10.0.0.5        128  39
  10.1.13.2     so-0/0/2.0     Full      10.0.0.3        128  33
```

```
user@R3> show ospf neighbor
  Address      Interface      State      ID              Pri  Dead
  10.1.34.2     so-0/0/0.0     Full      10.0.0.4        128  33
  10.1.23.1     so-0/0/1.0     Full      10.0.0.2        128  33
  10.1.13.1     so-0/0/2.0     Full      10.0.0.1        128  33
  10.1.36.2     so-0/0/3.0     Full      10.0.0.6        128  33
```

```
user@R6> show ospf neighbor
  Address      Interface      State      ID              Pri  Dead
  10.1.56.1     so-0/0/0.0     Full      10.0.0.5        128  30
  10.1.46.1     so-0/0/1.0     Full      10.0.0.4        128  38
  10.1.26.1     so-0/0/2.0     Full      10.0.0.2        128  34
  10.1.36.1     so-0/0/3.0     Full      10.0.0.3        128  35
```

Meaning The sample output shows that all neighbors are fully adjacent, indicating that each router has exchanged a full copy of its link-state database with the other routers, passed through several neighbor states, and become fully adjacent. These adjacencies are created by router link and network link advertisements.

Verify the OSPF Protocol Configuration

Purpose After you have checked interfaces and neighbors, verify the OSPF protocol configuration.

Action To verify the OSPF protocol configuration, enter the following command from the ingress, transit, and egress routers:

```
user@host> show configuration protocols ospf
```

Sample Output 1

```
user@R1> show configuration protocols ospf
traffic-engineering;
area 0.0.0.0 {
  interface so-0/0/0.0;
  interface so-0/0/1.0;
  interface so-0/0/2.0;      <<< The loopback interface (lo0) is missing
}
```

Sample Output 2

```
user@R3> show configuration protocols ospf
area 0.0.0.0 { <<< traffic engineering is missing
  interface so-0/0/0.0;
  interface so-0/0/1.0;
  interface so-0/0/2.0;
  interface so-0/0/3.0;      <<< The loopback interface (lo0) is missing
}
```

Sample Output 3

```
user@R6> show configuration protocols ospf
traffic-engineering;
area 0.0.0.0 {
  interface so-0/0/0.0;
  interface so-0/0/1.0;
  interface so-0/0/2.0;
  interface so-0/0/3.0;      <<< The loopback interface (lo0) is missing
}
```

Meaning All three sample outputs show that the loopback interface is not included on any of the routers. Including the loopback (**lo0**) interface is important when you have the RID manually configured.

In addition, Sample Output 2 from transit router **R3** shows that traffic engineering is not configured. Traffic engineering must be manually enabled when you configure OSPF for an MPLS network.

Because the loopback interface and traffic engineering are missing from the OSPF protocol configuration, the LSP does not work as expected.

Take Appropriate Action

Problem Depending on the error you encountered in your investigation, you must take the appropriate action to correct the problem. In this example, the loopback (**lo0**) interface is missing from all routers, and traffic engineering is missing from the transit router (**R3**).

Solution To correct the errors in this example, follow these steps:

1. Include the loopback (**lo0**) interface on all routers that have the RID manually configured. Enter the following configuration mode commands:

```
[edit]
user@R3# edit protocols ospf area 0.0.0.0
```

```
[edit protocols ospf area 0.0.0.0]
user@R3# set interface lo0
user@R3# set interface lo0 passive
```

2. Move up one level of the configuration hierarchy:

```
[edit protocols ospf area 0.0.0.0]
user@R3# up
[edit protocols ospf]
user@R3#
```

3. Include traffic engineering on the transit router (**R3**). Enter the following configuration mode command:

```
[edit protocols ospf]
user@R3# set traffic-engineering
```

4. On all routers, verify and commit the configuration:

```
user@R3# show
user@R3# commit
```

Sample Output

```
user@R3> edit
Entering configuration mode

[edit]
user@R3# edit protocols ospf area 0.0.0.0

[edit protocols ospf area 0.0.0.0]
user@R3# set interface lo0

[edit protocols ospf area 0.0.0.0]
user@R3# set interface lo0 passive

[edit protocols ospf area 0.0.0.0]
user@R3# up

[edit protocols ospf]
user@R3# set traffic-engineering

[edit protocols ospf]
user@R3# show
traffic-engineering;
area 0.0.0.0 {
  interface so-0/0/0.0;
  interface so-0/0/1.0;
  interface so-0/0/2.0;
  interface lo0.0; {
    passive
  }
}

[edit protocols ospf]
user@R3# commit
commit complete
```

Meaning The sample output shows that the loopback (**lo0**) interface and traffic engineering are now correctly configured on transit router **R3**. When traffic engineering is configured, OSPF advertises the traffic engineering capabilities of the links.

In the OSPF configuration, you must manually include the loopback (**lo0**) interface and set it to passive when you manually configure an RID. Setting the loopback (**lo0**) interface to passive ensures that protocols are not run over the loopback (**lo0**) interface and the loopback (**lo0**) interface is advertised correctly throughout the network.. If you do not manually configure an RID, there is no need to explicitly include the loopback interface because the OSPF protocol automatically includes the loopback (**lo0**) interface.

For more information about configuring LSPs and MPLS, see the *Junos MPLS Applications Configuration Guide*.

Verify the LSP Again

Purpose After taking the appropriate action to correct the error, the LSP needs to be checked again to confirm that the problem in the IS-IS protocol has been resolved.

Action To verify that the LSP is up and traversing the network as expected, enter the following command from the ingress, egress, and transit routers:

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

Sample Output

```

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

10.0.0.6
  From: 10.0.0.1, State: Up , ActiveRoute: 1, LSPname: R1-to-R6
  ActivePath: (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary State: Up
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

      10.1.13.2 10.1.36.2
    4 Oct 19 21:22:54 Selected as active path
    3 Oct 19 21:22:53 Record Route: 10.1.13.2 10.1.36.2
    2 Oct 19 21:22:53 Up
    1 Oct 19 21:22:53 Originate Call
  Created: Tue Oct 19 21:22:53 2004
Total 1 displayed, Up1 , Down 0

Egress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, LSPstate: Up , ActiveRoute: 0
  LSPname: R6-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: 117, Since: Tue Oct 19 21:17:42 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 2 receiver 39064 protocol 0
  PATH rcvfrom: 10.1.13.2 (so-0/0/2.0) 10 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.1.36.2 10.1.13.2 <self>
Total 1 displayed, Up1 , Down 0

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

user@R3> show mpls lsp extensive
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

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

Transit LSP: 2 sessions

10.0.0.1
  From: 10.0.0.6, LSPstate: Up, ActiveRoute: 1
  LSPname: R6-to-R1 , LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: 100416, Label out: 3
  Time left: 139, Since: Tue Oct 19 21:05:11 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 2 receiver 39064 protocol 0

```

```

PATH rcvfrom: 10.1.36.2 (so-0/0/3.0) 11 pkts
Adspec: received MTU 1500 sent MTU 1500
PATH sentto: 10.1.13.1 (so-0/0/2.0) 11 pkts
RESV rcvfrom: 10.1.13.1 (so-0/0/2.0) 11 pkts
Explct route: 10.1.13.1
Record route: 10.1.36.2 <self> 10.1.13.1

10.0.0.6
  From: 10.0.0.1, LSPstate: Up, ActiveRoute: 1
  LSPname: R1-to-R6 , 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: 135, Since: Tue Oct 19 21:10:22 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 47951 protocol 0
  PATH rcvfrom: 10.1.13.1 (so-0/0/2.0) 4 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 10.1.36.2 (so-0/0/3.0) 4 pkts
  RESV rcvfrom: 10.1.36.2 (so-0/0/3.0) 4 pkts
  Record route: 10.1.13.1 <self> 10.1.36.2
Total 2 displayed, Up2 , Down 0

user@R6> run show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, State: Up, ActiveRoute: 1, LSPname: R6-to-R1
  ActivePath: (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 2)
    10.1.36.1 S 10.1.13.1 S
      Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

          10.1.36.1 10.1.13.1
19 Oct 19 21:09:52 Selected as active path
18 Oct 19 21:09:52 Record Route: 10.1.36.1 10.1.13.1
17 Oct 19 21:09:52 Up
16 Oct 19 21:09:52 Originate Call
15 Oct 19 21:09:52 CSPF: computation result accepted
Created: Tue Oct 19 18:30:09 2004
Total 1 displayed, Up1 , Down 0

Egress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, LSPstate: Up, ActiveRoute: 0
  LSPname: R1-to-R6 , 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: 120, Since: Tue Oct 19 21:15:03 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 47951 protocol 0
  PATH rcvfrom: 10.1.36.1 (so-0/0/3.0) 4 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.1.13.1 10.1.36.1 <self>

```

Total 1 displayed, Up 1 , Down 0

Transit LSP: 0 sessions

Total 0 displayed, Up 0, Down 0

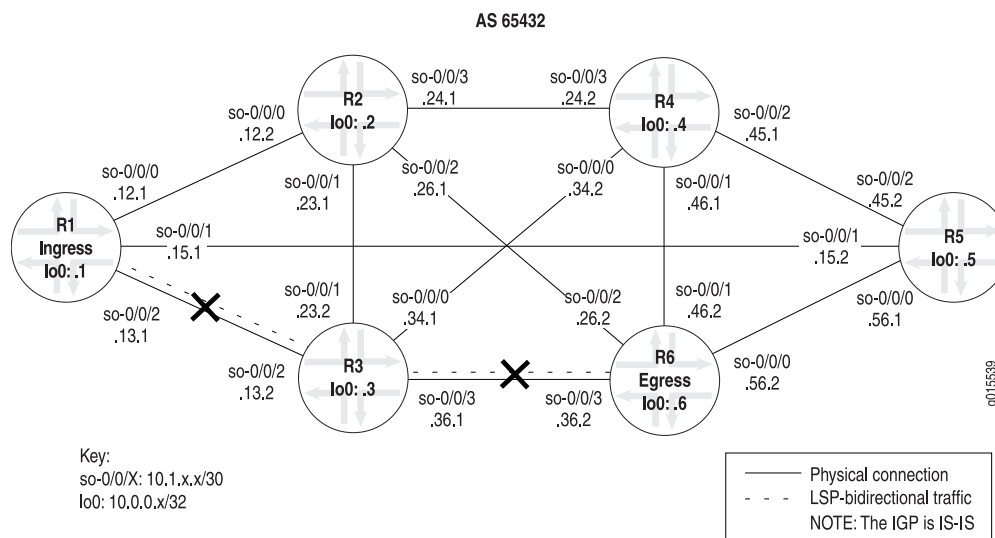
Meaning The sample output from ingress router **R1** and egress router **R6** shows that the LSP is now traversing the network along the expected path, from **R1** through **R3** to **R6**, and the reverse LSP, from **R6** through **R3** to **R1**. In addition, the sample output from transit router **R3** shows that there are two transit LSP sessions, one from **R1** to **R6**, and the other from **R6** to **R1**.

Verifying the IS-IS Protocol

Purpose If your MPLS network is configured with IS-IS as the interior gateway protocol (IGP), and the output of the **show mpls lsp extensive** command shows that there is a problem, check the IP and IS-IS layers. Because IS-IS and IP are independent of each other, you can check either layer first. For more information about checking the IP layer, see [“Verifying the IP Layer” on page 121](#).

After you have checked the IP layer and determined that there is still a problem, check the IS-IS layer, verify that IS-IS adjacencies are up, and make sure that the interfaces and IS-IS protocol are configured correctly.

Figure 18: MPLS Network Broken at the IS-IS Protocol Layer



To check the IS-IS protocol, follow these steps:

1. [Verify the LSP on page 145](#)
2. [Verify IS-IS Adjacencies and Interfaces on page 147](#)
3. [Verify the IS-IS Configuration on page 149](#)
4. [Take Appropriate Action on page 150](#)
5. [Verify the LSP Again on page 151](#)

Verify the LSP

Purpose Confirm that interfaces are configured for IS-IS, that the IS-IS protocol is configured correctly, and that adjacencies are established.

Action To verify the label-switched path (LSP), enter the following command on the ingress, transit, and egress routers:

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

Sample Output 1

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

10.0.0.6
  From: 10.0.0.1, State: Dn,  ActiveRoute: 0 , LSPname: R1-to-R6
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary                               State: Dn
    24 Oct 21 13:48:01  No Route toward dest [3 times]
    23 Oct 21 13:47:44 Deselected as active
    22 Oct 21 13:47:43 No Route toward dest[2 times]
    21 Oct 21 13:47:43 ResvTear received
    20 Oct 21 13:47:43 Down
    19 Oct 21 13:47:43 10.1.13.2: No Route toward dest[2 times]
    18 Oct 21 13:47:38 Record Route:  10.1.13.2 10.1.36.2
  [...Output truncated...]
  Created: Tue Oct 19 21:22:53 2004
Total 1 displayed, Up 0,  Down1

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

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

Sample Output 2

```
user@R3> show mpls lsp extensive
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

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

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

Sample Output 3

```
user@R6> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, State: Dn,  ActiveRoute: 0 , LSPname:  R6-to-R1
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary                               State: Dn
    Will be enqueued for recomputation in 3 second(s).
    13 Oct 21 14:23:33 CSPF failed: no route toward 10.0.0.1[90 times]
    12 Oct 21 13:39:56 Deselected as active
    11 Oct 21 13:39:56 CSPF: could not determine self
  [...Output truncated...]
  Created: Tue Oct 19 22:28:30 2004
Total 1 displayed, Up 0,  Down1
```

```
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 shows that LSP **R1-to-R6** and the reverse LSP **R6-to-R1** are down, and there are no LSP sessions on transit router R3.

Verify IS-IS Adjacencies and Interfaces

Purpose When you check the IS-IS layer, you verify that IS-IS adjacencies are up and that the IS-IS interfaces are included at the protocol level.

Action To verify the functioning of adjacent interfaces, enter the following commands from the relevant routers:

```
user@host> show isis adjacency  
user@host> show isis interface
```

Sample Output 1

```
user@R1> show isis adjacency
Interface      System      L State      Hold (secs) SNPA
so-0/0/0.0     R2           2 Up         20
so-0/0/1.0     R5           2 Up         23
so-0/0/2.0     R3           2 Up         26
```

```
user@R3> show isis adjacency
Interface      System      L State      Hold (secs) SNPA
so-0/0/0.0     R4           2 Up         23
so-0/0/1.0     R2           2 Up         21
so-0/0/2.0     R1           2 Up         19
so-0/0/3.0     R6           2 Down       0
```

```
user@R6> show isis adjacency
IS-IS instance is not running
```

Sample Output 2

```
user@R1> show isis interface
IS-IS interface database:
Interface      L CirID Level 1 DR      Level 2 DR      L1/L2 Metric
lo0.0          0 0x1 Passive           Passive          0/0
so-0/0/0.0     2 0x1 Disabled         Point to Point   10/10
so-0/0/1.0     2 0x1 Disabled         Point to Point   10/10
so-0/0/2.0     2 0x1 Disabled         Point to Point   10/10
```

```
user@R3> show isis interface
IS-IS interface database:
Interface      L CirID Level 1 DR      Level 2 DR      L1/L2 Metric
lo0.0          0 0x1 Passive           Passive          0/0
so-0/0/0.0     2 0x1 Disabled         Point to Point   10/10
so-0/0/1.0     2 0x1 Disabled         Point to Point   10/10
so-0/0/2.0     2 0x1 Disabled         Point to Point   10/10
so-0/0/3.0     2 0x1 Disabled         Point to Point   10/10
```

```
user@R6> show isis interface
IS-IS interface database:
Interface      L CirID Level 1 DR      Level 2 DR      L1/L2 Metric
lo0.0          0 0x1 Passive           Passive          0/0
so-0/0/0.0     1 0x1 Point to Point    Disabled         10/10
so-0/0/1.0     1 0x1 Down              Disabled         10/10
so-0/0/2.0     1 0x1 Point to Point    Disabled         10/10
so-0/0/3.0     1 0x1 Point to Point    Disabled         10/10
```

Meaning Sample Output 1 shows that ingress router R1 has established adjacencies with the relevant routers. Transit router R3 does not have an adjacency with egress router R6, and egress router R6 has no adjacencies established in the network shown in [Figure 15 on page 121](#), indicating that the problem might be at the IS-IS protocol level.

Sample Output 2 shows that R1 and R2 are Level 2 routers, in contrast to R6 which is a Level 1 router. When a router is configured explicitly as a Level 1 or Level 2 router, it does not communicate with routers configured at a different level. Level 1 routers communicate with other Level 1 routers within their area, while Level 2 routers communicate with other Level 2 routers, and toward other autonomous systems. Because all the routers in this

network are configured for Level 2, they cannot form an adjacency with R6, which is incorrectly configured as a Level 1 router.

Verify the IS-IS Configuration

Purpose When you have determined that the problem is probably at the IS-IS protocol level, check the IS-IS configuration of the routers in your network.

Action To verify the IS-IS configuration, enter the following command from the relevant routers:

```
user@host> show configuration protocols isis
```

Sample Output

```
user@R1> show configuration protocols isis
level 1 disable;
interface so-0/0/0.0;
interface so-0/0/1.0;
interface so-0/0/2.0;
interface lo0.0; {
    passive
```

```
user@R3> show configuration protocols isis
level 1 disable;
interface all {
    level 2 metric 10;
}
interface fxp0.0 {
    disable;
}
interface lo0.0; {
    passive
```

```
user@R6> show configuration protocols isis
level 2 disable; <<< Incorrect level disabled
interface all {
    level 2 metric 10;
}
interface fxp0.0 {
    disable;
}
interface lo0.0; {
    passive
```

Meaning The sample output shows that R6 has Level 2 disabled, while R1 and R3 have Level 1 disabled. For IS-IS adjacencies to establish, routers need to be at the same level. Another common configuration error is to omit the loopback interface (lo0) from the configuration at the **[edit protocols isis]** hierarchy level. IS-IS does not function correctly if the loopback interface (lo0) is not configured at this level. In addition, including the **passive** statement ensures that protocols are not run over the loopback interface (lo0) and that the loopback interface (lo0) is advertised correctly throughout the network.

Take Appropriate Action

Problem Depending on the error you encountered in your investigation, you must take the appropriate action to correct the problem. In the example below, the routers are configured to function at different levels of the IS-IS protocol.

Solution To correct the error in this example, enter the following commands:

Sample Output

```
[edit protocols isis]
user@R6# show
level 2 disable;
interface all {
    level 2 metric 10;
}
interface fxp0.0 {
    disable;
}
interface lo0.0; {
    passive

[edit protocols isis]
user@R6# delete level 2

[edit protocols isis]
user@R6# set level 1 disable

[edit protocols isis]
user@R6# show
level 1 disable;
interface all {
    level 2 metric 10;
}
interface fxp0.0 {
    disable;
}
interface lo0.0; {
    passive

[edit protocols isis]
user@R6# commit
commit complete

[edit protocols isis]
user@R6# run show isis adjacency
Interface          System      L State      Hold (secs) SNPA
so-0/0/0.0         R5          2 Up         22
so-0/0/1.0         R4          2 Up         22
so-0/0/2.0         R2          2 Up         22
so-0/0/3.0         R3          2 Up         22
```

Meaning The sample output shows that the configuration error on egress router R6 has been corrected, and IS-IS adjacencies are now established.

Verify the LSP Again

Purpose After taking the appropriate action to correct the error, the label-switched path (LSP) needs to be checked again to confirm that the problem in the RSVP layer has been resolved.

Action To verify that the LSP is up and traversing the network as expected, enter the following command from the ingress, egress, and transit routers:

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

Sample Output 1

```

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

10.0.0.6
  From: 10.0.0.1, State: Up,  ActiveRoute: 1, LSPname: R1-to-R6
  ActivePath: (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
  10.1.13.2 S 10.1.36.2 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

        10.1.13.2 10.1.36.2
      5 Oct 21 15:52:07 Selected as active path
      4 Oct 21 15:52:07 Record Route:  10.1.13.2 10.1.36.2
      3 Oct 21 15:52:07 Up
      2 Oct 21 15:52:07 Originate Call
      1 Oct 21 15:52:07 CSPF: computation result accepted
    Created: Thu Oct 21 15:52:06 2004
Total 1 displayed,  Up1 , Down 0

Egress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, LSPstate: Up, ActiveRoute: 0
  LSPname: R6-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: 142, Since: Thu Oct 21 15:41:59 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 2 receiver 39082 protocol 0
  PATH rcvfrom: 10.1.13.2 (so-0/0/2.0) 17 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.1.36.2 10.1.13.2 <self>
Total 1 displayed,  Up1 , Down 0

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

```

Sample Output 2

```

user@R3> show mpls lsp extensive
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

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

Transit LSP: 2 sessions

10.0.0.1
  From: 10.0.0.6, LSPstate: Up, ActiveRoute: 1
  LSPname: R6-to-R1 , LSPpath: Primary
  Suggested label received: -, Suggested label sent: -

```

```

Recovery label received: -, Recovery label sent: 3
Resv style: 1 FF, Label in: 100528, Label out: 3
Time left: 125, Since: Thu Oct 21 15:29:26 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 2 receiver 39082 protocol 0
PATH rcvfrom: 10.1.36.2 (so-0/0/3.0) 17 pkts
Adspec: received MTU 1500 sent MTU 1500
PATH sentto: 10.1.13.1 (so-0/0/2.0) 17 pkts
RESV rcvfrom: 10.1.13.1 (so-0/0/2.0) 17 pkts
Explct route: 10.1.13.1
Record route: 10.1.36.2 <self> 10.1.13.1

10.0.0.6
From: 10.0.0.1, LSPstate: Up, ActiveRoute: 1
  LSPname: R1-to-R6 , LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: 3
Resv style: 1 FF, Label in: 100544, Label out: 3
Time left: 147, Since: Thu Oct 21 15:39:33 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 1 receiver 47963 protocol 0
PATH rcvfrom: 10.1.13.1 (so-0/0/2.0) 4 pkts
Adspec: received MTU 1500 sent MTU 1500
PATH sentto: 10.1.36.2 (so-0/0/3.0) 4 pkts
RESV rcvfrom: 10.1.36.2 (so-0/0/3.0) 4 pkts
Explct route: 10.1.36.2
Record route: 10.1.13.1 <self> 10.1.36.2
Total 2 displayed, Up 2, Down 0

```

Sample Output 3

```

user@R6> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.1
From: 10.0.0.6, State: Up, ActiveRoute: 1,  LSPname: R6-to-R1
ActivePath: (primary)
LoadBalance: Random
Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary                               State: Up
  Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
10.1.36.1S 10.1.13.1S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

      10.1.36.1 10.1.13.1
18 Oct 21 15:34:18 Selected as active path
17 Oct 21 15:34:17 Record Route: 10.1.36.1 10.1.13.1
16 Oct 21 15:34:17 Up
15 Oct 21 15:34:17 Originate Call
14 Oct 21 15:34:17 CSPF: computation result accepted
[...Output truncated...]
Created: Tue Oct 19 22:28:30 2004
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

10.0.0.6
From: 10.0.0.1, LSPstate: Up, ActiveRoute: 0
  LSPname: R1-to-R6 , 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: 126, Since: Thu Oct 21 15:44:25 2004  
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500  
Port number: sender 1 receiver 47963 protocol 0  
PATH rcvfrom: 10.1.36.1 (so-0/0/3.0) 4 pkts  
Adspec: received MTU 1500  
PATH sentto: localclient  
RESV rcvfrom: localclient  
Record route: 10.1.13.1 10.1.36.1 <self>  
Total 1 displayed, Up 1, Down 0  
  
Transit LSP: 0 sessions  
Total 0 displayed, Up 0, Down 0
```

Meaning Sample Outputs 1 and 3 from ingress router R1 and egress router R6 show that the LSP is now traversing the network along the expected path, from R1 through R3 to R6, and the reverse LSP, from R6 through R3 to R1. In addition, Sample Output 2 from transit router R3 shows that there are two transit LSP sessions, one from R1 to R6, and the other from R6 to R1.

Checking the RSVP Layer

This chapter describes how to check the Resource Reservation Protocol (RSVP) layer of the layered Multiprotocol Label Switching (MPLS) model.

- [Checklist for Checking the RSVP Layer on page 155](#)
- [Checking the RSVP Layer on page 156](#)

Checklist for Checking the RSVP Layer

Problem This checklist provides the steps and commands for checking the Resource Reservation Protocol (RSVP) layer of the layered Multiprotocol Label Switching (MPLS) model. The checklist provides links to an overview of the RSVP layer and more detailed information about the commands used to investigate the problem.

[Table 14 on page 155](#) provides commands for checking the RSVP layer.

Table 14: Checklist for Checking the RSVP Layer

Tasks	Command or Action
"Checking the RSVP Layer" on page 156	
1. Verify the LSP on page 158	<code>show mpls lsp extensive</code>
2. Verify RSVP Sessions on page 160	<code>show rsvp session</code>
3. Verify RSVP Neighbors on page 162	<code>show rsvp neighbor</code>
4. Verify RSVP Interfaces on page 163	<code>show rsvp interface</code>
5. Verify the RSVP Protocol Configuration on page 166	<code>show configuration protocols rsvp</code>

Table 14: Checklist for Checking the RSVP Layer (*continued*)

Tasks	Command or Action
6. Take Appropriate Action on page 167	<p>The following sequence of commands addresses the specific problem described in this topic:</p> <pre> [edit] edit protocols rsvp [edit protocols rsvp] show set interface <i>type-fpc/plc/port</i> show commit </pre>
7. Verify the LSP Again on page 168	<code>show mpls lsp extensive</code>

Checking the RSVP Layer

Purpose After you have configured the label-switched path (LSP), issued the `show mpls lsp extensive` command, and determined that there is an error, you might find that the error is not in the physical, data link, or Internet Protocol (IP) and interior gateway protocol (IGP) layers. Continue investigating the problem at the RSVP layer of the network.

[Figure 19 on page 156](#) illustrates the RSVP layer of the layered MPLS model.

Figure 19: Checking the RSVP Layer

BGP Layer	<traceroute <i="">host-name show bgp summary show configuration protocols bgp show route <i>destination-prefix</i> detail show route receive protocol bgp <i>neighbor-address</i> </traceroute>
MPLS Layer	show mpls lsp show mpls lsp extensive show route table mpls.0 show route <i>address</i> traceroute <i>address</i> ping mpls rsvp <i>lsp-name</i> detail
RSVP Layer	show rsvp session show rsvp neighbor show rsvp interface
<div> <div>IGP and IP Layers Functioning</div> <div> <div> <div> <div>OSPF Layer</div> <div>show ospf neighbor</div> <div>show configuration protocols ospf</div> <div>show ospf interface</div> </div> <div> <div>IS-IS Layer</div> <div>show isis adjacency</div> <div>show configuration protocols isis</div> <div>show isis interface</div> </div> </div> <div> <div> <div>IP Layer</div> <div>show ospf neighbor extensive</div> <div>show interfaces terse</div> </div> <div> <div>IP Layer</div> <div>show isis adjacency extensive</div> <div>show interfaces terse</div> </div> </div> </div> </div>	
Data Link Layer	show interfaces extensive <i>"JUNOS Interfaces Operations Guide"</i>
Physical Layer	show interfaces show interfaces terse ping <i>host</i>

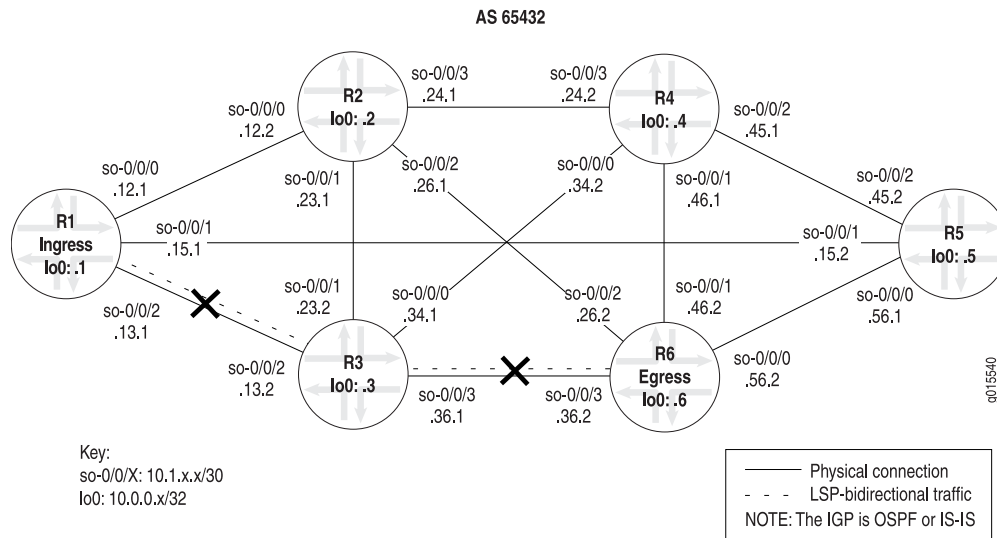
g015546

With this layer, you check that dynamic RSVP signaling is occurring as expected, neighbors are connected, and interfaces are configured correctly for RSVP. Check the ingress, egress, and transit routers.

If the network is not functioning at this layer, the LSP does not work as configured.

Figure 20 on page 157 illustrates the MPLS network used in this topic.

Figure 20: MPLS Network Broken at the RSVP Layer



The network shown in Figure 20 on page 157 is a fully meshed configuration where every directly connected interface can receive and send packets to every other similar interface. The LSP in this network is configured to run from ingress router **R1**, through transit router **R3**, to egress router **R6**. In addition, a reverse LSP is configured to run from **R6** through **R3** to **R1**, creating bidirectional traffic.

However, in this example, the LSP is down without a path in either direction, from **R1** to **R6** or from **R6** to **R1**.

The crosses shown in Figure 20 on page 157 indicate where the LSP is broken. Some possible reasons the LSP is broken might include that dynamic RSVP signaling is not occurring as expected, neighbors are not connected, or interfaces are incorrectly configured for RSVP.

In the network in Figure 20 on page 157, a configuration error on transit router **R3** prevents the LSP from traversing the network as expected.

To check the RSVP layer, follow these steps:

1. [Verify the LSP on page 158](#)
2. [Verify RSVP Sessions on page 160](#)
3. [Verify RSVP Neighbors on page 162](#)
4. [Verify RSVP Interfaces on page 163](#)
5. [Verify the RSVP Protocol Configuration on page 166](#)

6. [Take Appropriate Action on page 167](#)
7. [Verify the LSP Again on page 168](#)

Verify the LSP

Purpose Typically, you use the **show mpls lsp extensive** command to verify the LSP. However for quick verification of the LSP state, use the **show mpls lsp** command. If the LSP is down, use the **extensive** option (**show mpls lsp extensive**) as a follow-up. If your network has numerous LSPs, you might consider specifying the name of the LSP, using the **name** option (**show mpls lsp name *name*** or **show mpls lsp name *name* extensive**).

Action To determine whether the LSP is up, enter the following command from the ingress router:

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

Sample Output 1

```

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

10.0.0.6
  From: 10.0.0.1, State: Dn, ActiveRoute: 0, LSPname: R1-to-R6
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary State: Dn
    2 Oct 27 15:06:05 10.1.13.2: No Route toward dest [4 times]
    1 Oct 27 15:05:56 Originate Call
  Created: Wed Oct 27 15:05:55 2004
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@R3> show mpls lsp extensive
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

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

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

user@R6> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, State: Dn, ActiveRoute: 0, LSPname: R6-to-R1
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary State: Dn
    Will be enqueued for recomputation in 22 second(s).
    1 Oct 27 14:59:12 CSPF failed: no route toward 10.0.0.1 [4 times]
  Created: Wed Oct 27 14:57:44 2004
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 shows that the LSP is down in both directions, from **R1** to **R6**, and from **R6** to **R1**. The output from **R1** shows that **R1** is using a no-cspf LSP since it tried to originate the call without being able to reach the destination. The output from **R6** shows that the Constrained Shortest Path First (CSPF) algorithm failed, resulting in no route to destination **10.0.0.1**.

Verify RSVP Sessions

Purpose When an RSVP session is successfully created, the LSP is set up along the paths created by the RSVP session. If the RSVP session is unsuccessful, the LSP does not work as configured.

Action To verify currently active RSVP sessions, enter the following command from the ingress, transit, and egress routers:

```
user@host> show rsvp session
```

Sample Output 1

```

user@R1> show rsvp session
Ingress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

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

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

user@R3> show rsvp session
Ingress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

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

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

user@R6> show rsvp session
Ingress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

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

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

```

Sample Output 2

```

user@R1> show rsvp session
Ingress RSVP: 1 sessions
  To      From      State Rt Style Labelin Labelout LSPName
10.0.0.6  10.0.0.1  Up    1 1 FF      -    100768  R1-to-R6
Total 1 displayed, Up1 , Down 0

Egress RSVP: 1 sessions
  To      From      State Rt Style Labelin Labelout LSPName
10.0.0.1  10.0.0.6  Up    0 1 FF      3      -    R6-to-R1
Total 1 displayed, Up1 , Down 0

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

user@R3> show rsvp session
Ingress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

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

Transit RSVP: 2 sessions
  To      From      State Rt Style Labelin Labelout LSPName
10.0.0.1  10.0.0.6  Up    1 1 FF    100784      3  R6-to-R1
10.0.0.6  10.0.0.1  Up    1 1 FF    100768      3  R1-to-R6
Total 2 displayed, Up2 , Down 0

```

```
user@R6> show rsvp session
Ingress RSVP: 1 sessions
To          From          State Rt Style Labelin Labelout LSPname
10.0.0.1    10.0.0.6    Up    1 1 FF      -    100784  R6-to-R1
Total 1 displayed, Up 1 , Down 0

Egress RSVP: 1 sessions
To          From          State Rt Style Labelin Labelout LSPname
10.0.0.6    10.0.0.1    Up    0 1 FF      3      -    R1-to-R6
Total 1 displayed, Up 1 , Down 0

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

Meaning Sample Output 1 from all routers shows that no RSVP sessions were successfully created, even though the LSP **R6-to-R1** is configured. Continue investigating the problem in [“Verify RSVP Neighbors” on page 162](#).

In contrast to Sample Output 1 and to illustrate correct output, Sample Output 2 shows the output from the ingress, transit, and egress routers when the RSVP configuration is correct, and the LSP is traversing the network as configured. **R1** and **R6** both show an ingress and egress RSVP session, with the LSP **R1-to-R6**, and the reverse LSP **R6-to-R1**. Transit router **R3** shows two transit RSVP sessions.

Verify RSVP Neighbors

Purpose Display a list of RSVP neighbors that were learned dynamically when exchanging RSVP packets. Once a neighbor is learned, it is never removed from the list of RSVP neighbors unless the RSVP configuration is removed from the router.

Action To verify RSVP neighbors, enter the following command from the ingress, transit, and egress routers:

```
user@host> show rsvp neighbor
```

Sample Output 1

```

user@R1> show rsvp neighbor
RSVP neighbor: 1 learned
Address          Idle Up/Dn LastChange HelloInt HelloTx/Rx MsgRcvd
10.1.13.2         10  1/0      9:22         9    64/64    32

user@R3> show rsvp neighbor
RSVP neighbor: 2 learned
Address          Idle Up/Dn LastChange HelloInt HelloTx/Rx MsgRcvd
10.1.13.1         0  1/0      28:20        9   190/190   41
10.1.36.2        16:50 1/1      15:37        9   105/78    38

user@R6> show rsvp neighbor
RSVP neighbor: 1 learned
Address          Idle Up/Dn LastChange HelloInt HelloTx/Rx MsgRcvd
10.1.36.1        17:30 1/1      16:15        9   104/78    39

```

Sample Output 2

```

user@R3> show rsvp neighbor
RSVP neighbor: 2 learned
Address          Idle Up/Dn LastChange HelloInt HelloTx/Rx MsgRcvd
10.1.13.1         5  1/0      9:14         9    63/63    33
10.1.36.2         5  1/0      9:05         9    62/62    32

user@R6> show rsvp neighbor
RSVP neighbor: 1 learned
Address          Idle Up/Dn LastChange HelloInt HelloTx/Rx MsgRcvd
10.1.36.1         5  1/0      8:54         9    61/61    32

```

Meaning Sample Output 1 shows that **R1** and **R6** have one RSVP neighbor each, **R3**. However, the values in the **Up/Dn** field are different. **R1** has a value of **1/0** and **R6** has a value of **1/1**, indicating that **R1** is an active neighbor with **R3**, but **R6** is not. When the up count is one more than the down count, the neighbor is active; if the values are equal, the neighbor is down. The values for **R6** are equal, **1/1**, indicating that the neighbor **R3** is down.

Transit router **R3** knows about two neighbors, **R1** and **R6**. The **Up/Dn** field indicates that **R1** is an active neighbor and **R6** is down. At this point it is not possible to determine if the problem resides with **R3** or **R6**, because both neighbors are not active. Continue investigating the problem in [“Verify RSVP Interfaces” on page 163](#).

In contrast to Sample Output 1 and to illustrate correct output, Sample Output 2 shows the correct neighbor relationship between transit router **R3** and egress router **R6**. The **Up/Dn** field shows the up count to be one more than the down count, **1/0**, indicating that the neighbors are active.

Verify RSVP Interfaces

- Purpose** Display the status of each interface on which RSVP is enabled to determine where the configuration error occurred.
- Action** To verify the status of RSVP interfaces, enter the following command from the ingress, transit, and egress routers:

```
user@host> show rsvp interface
```


Sample Output 1

```
user@R1> show rsvp interface
```

```
RSVP interface: 3 active
```

Interface	State	Active resv	Subscr- ption	Static BW	Available BW	Reserved BW	Highwater mark
so-0/0/0.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/1.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/2.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps

```
user@R3> show rsvp interface
```

```
RSVP interface: 3 active
```

Interface	State	Active resv	Subscr- ption	Static BW	Available BW	Reserved BW	Highwater mark
so-0/0/0.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/1.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/2.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps

```
<<< Missing interface so-0/0/3.0
```

```
user@R6> show rsvp interface
```

```
RSVP interface: 4 active
```

Interface	State	Active resv	Subscr- ption	Static BW	Available BW	Reserved BW	Highwater mark
so-0/0/0.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/1.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/2.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/3.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps

Sample Output 2

```
user@R1> show rsvp interface
```

```
RSVP interface: 3 active
```

Interface	State	Active resv	Subscr- ption	Static BW	Available BW	Reserved BW	Highwater mark
so-0/0/0.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/1.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/2.0	Up	1	100%	155.52Mbps	155.52Mbps	0bps	0bps

```
user@R3> show rsvp interface
```

```
RSVP interface: 4 active
```

Interface	State	Active resv	Subscr- ption	Static BW	Available BW	Reserved BW	Highwater mark
so-0/0/0.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/1.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/2.0	Up	1	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/3.0	Up	1	100%	155.52Mbps	155.52Mbps	0bps	0bps

```
user@R6> show rsvp interface
```

```
RSVP interface: 4 active
```

Interface	State	Active resv	Subscr- ption	Static BW	Available BW	Reserved BW	Highwater mark
so-0/0/0.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/1.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/2.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps

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

Meaning Sample Output 1 shows that even though each router has interfaces that are up and have RSVP active, there are no reservations (**Active resv**) on any of the routers. In this example, we would expect at least one reservation on the ingress and egress routers, and two reservations on the transit router.

In addition, interface **so-0/0/3** on transit router **R3** is not included in the configuration. The inclusion of this interface is critical to the success of the LSP.

In contrast to Sample Output 1 and to illustrate correct output, Sample Output 2 shows the relevant interfaces with active reservations.

Verify the RSVP Protocol Configuration

Purpose After you have checked RSVP sessions, interfaces, neighbors, and determined that there might be a configuration error, verify the RSVP protocol configuration.

Action To verify the RSVP configuration, enter the following command from the ingress, transit, and egress routers:

```
user@host> show configuration protocols rsvp
```

Sample Output

```
user@R1> show configuration protocols rsvp
```

```
interface so-0/0/0.0;
interface so-0/0/1.0;
interface so-0/0/2.0;
interface fxp0.0 {
    disable;
}
```

```
user@R3> show configuration protocols rsvp
```

```
interface so-0/0/0.0;
interface so-0/0/1.0;
interface so-0/0/2.0; <<< Missing interface so-0/0/3.0
interface fxp0.0 {
    disable;
}
```

```
user@R6> show configuration protocols rsvp
```

```
interface so-0/0/0.0;
interface so-0/0/1.0;
interface so-0/0/2.0;
interface so-0/0/3.0;
interface fxp0.0 {
    disable;
}
```

Meaning The sample output shows that **R3** has interface **so-0/0/3.0** missing from the RSVP protocol configuration. This interface is critical for the correct functioning of the LSP.

Take Appropriate Action

Problem Depending on the error you encountered in your investigation, you must take the appropriate action to correct the problem. In this example, an interface is missing from the configuration of router R3.

Solution To correct the error in this example, follow these steps:

1. Include the missing interface in the configuration of transit router R3:

```
user@R3> edit
user@R3# edit protocols rsvp
[edit protocols rsvp]
user@R3# show
user@R3# set interface so-0/0/3.0
```

2. Verify and commit the configuration:

```
[edit protocols rsvp]
user@R3# show
user@R3# commit
```

Sample Output

```
user@R3> edit
Entering configuration mode

[edit]
user@R3# edit protocols rsvp

[edit protocols rsvp]
user@R3# show
interface so-0/0/0.0;
interface so-0/0/1.0;
interface so-0/0/2.0; <<< Missing interface so-0/0/3.0
interface fxp0.0 {
    disable;
}
[edit protocols rsvp]
user@R3# set interface so-0/0/3.0

[edit protocols rsvp]
user@R3# show
interface so-0/0/0.0;
interface so-0/0/1.0;
interface so-0/0/2.0;
interface fxp0.0 {
    disable;
}
interface so-0/0/3.0; <<< Interface now included in the configuration

[edit protocols rsvp]
user@R3# commit
commit complete
```

Meaning The sample output shows that the missing interface **so-0/0/3.0** on transit router **R3** is now correctly included at the **[edit protocols rsvp]** hierarchy level. This results in the possibility that the LSP might come up.

Verify the LSP Again

Purpose After taking the appropriate action to correct the error, the LSP needs to be checked again to confirm that the problem in the MPLS layer has been resolved.

Action To verify the LSP again, enter the following command on the ingress, transit, and egress routers:

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

Sample Output 1

```

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

10.0.0.6
  From: 10.0.0.1, State: Up,  ActiveRoute: 1, LSPname: R1-to-R6
  ActivePath: (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

      10.1.13.2 10.1.36.2
    5 Oct 27 15:28:57 Selected as active path
    4 Oct 27 15:28:57 Record Route: 10.1.13.2 10.1.36.2
    3 Oct 27 15:28:57 Up
    2 Oct 27 15:28:44 10.1.13.2: No Route toward dest[35 times]
    1 Oct 27 15:05:56 Originate Call
    Created: Wed Oct 27 15:05:56 2004
  Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, LSPstate: Up, ActiveRoute: 0
  LSPname: R6-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: 136, Since: Wed Oct 27 15:29:20 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 39092 protocol 0
  PATH rcvfrom: 10.1.13.2 (so-0/0/2.0) 6 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.1.36.2 10.1.13.2 <self>
  Total 1 displayed, Up 1, Down 0

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

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

```

Sample Output 2

```

user@R3> show mpls lsp extensive
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

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

Transit LSP: 2 sessions

10.0.0.1
  From: 10.0.0.6, LSPstate: Up, ActiveRoute: 1
  LSPname: R6-to-R1, LSPpath: Primary

```

```

Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: 3
Resv style: 1 FF, Label in: 100672, Label out: 3
Time left: 152, Since: Wed Oct 27 15:16:39 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 1 receiver 39092 protocol 0
PATH rcvfrom: 10.1.36.2 (so-0/0/3.0) 7 pkts
Adspec: received MTU 1500 sent MTU 1500
PATH sentto: 10.1.13.1 (so-0/0/2.0) 7 pkts
RESV rcvfrom: 10.1.13.1 (so-0/0/2.0) 7 pkts
Explct route: 10.1.13.1
Record route: 10.1.36.2 <self> 10.1.13.1

```

10.0.0.6

```

From: 10.0.0.1, LSPstate: Up, ActiveRoute: 1
LSPname: R1-to-R6, LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: 3
Resv style: 1 FF, Label in: 100656, Label out: 3
Time left: 129, Since: Wed Oct 27 14:53:14 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 1 receiver 47977 protocol 0
PATH rcvfrom: 10.1.13.1 (so-0/0/2.0) 40 pkts
Adspec: received MTU 1500 sent MTU 1500
PATH sentto: 10.1.36.2 (so-0/0/3.0) 7 pkts
RESV rcvfrom: 10.1.36.2 (so-0/0/3.0) 7 pkts
Record route: 10.1.13.1 <self> 10.1.36.2

```

Total 2 displayed, Up 2, Down 0

Sample Output 3

```
user@R6> show mpls lsp extensive
```

Ingress LSP: 1 sessions

10.0.0.1

```

From: 10.0.0.6, State: Up, ActiveRoute:1, LSPname: R6-to-R1
ActivePath: (primary)
LoadBalance: Random
Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary State: Up
Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
10.1.36.1 S 10.1.13.1 S
Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

```

```

10.1.36.1 10.1.13.1
6 Oct 27 15:22:06 Selected as active path
5 Oct 27 15:22:06 Record Route: 10.1.36.1 10.1.13.1
4 Oct 27 15:22:06 Up
3 Oct 27 15:22:06 Originate Call
2 Oct 27 15:22:06 CSPF: computation result accepted
1 Oct 27 15:21:36 CSPF failed: no route toward 10.0.0.1[50 times]
Created: Wed Oct 27 14:57:45 2004

```

Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

10.0.0.6

```

From: 10.0.0.1, LSPstate: Up, ActiveRoute: 0
LSPname: R1-to-R6, 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: Wed Oct 27 15:21:43 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 1 receiver 47977 protocol 0
PATH rcvfrom: 10.1.36.1 (so-0/0/3.0) 7 pkts
Adspec: received MTU 1500
PATH sentto: localclient
RESV rcvfrom: localclient
Record route: 10.1.13.1 10.1.36.1 <self>
Total 1 displayed, Up 1, Down 0

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

Meaning Sample Output 1 from ingress router **R1** shows that LSP **R1-to-R6** has an active route to **R6** and the state is up.

Sample Output 2 from transit router **R3** shows that there are two transit LSP sessions, one from **R1** to **R6** and the other from **R6** to **R1**. Both LSPs are up.

Sample Output 3 from egress router **R6** shows that the LSP is up and the active route is the primary route. The LSP is now traversing the network along the expected path, from **R1** through **R3** to **R6**, and the reverse LSP, from **R6** through **R3** to **R1**.

CHAPTER 11

Checking the MPLS Layer

This chapter describes how to check the Multiprotocol Label Switching (MPLS) layer of the layered MPLS model.

- [Checklist for Checking the MPLS Layer on page 173](#)
- [Checking the MPLS Layer on page 174](#)

Checklist for Checking the MPLS Layer

Problem This checklist provides the steps and commands for checking the Multiprotocol Label Switching (MPLS) layer of the layered MPLS model. The checklist provides links to an overview of the MPLS layer and more detailed information about the commands used to investigate the problem.

[Table 15 on page 173](#) provides commands for checking the MPLS layer.

Table 15: Checklist for Checking the MPLS Layer

Tasks	Command or Action
“Checking the MPLS Layer” on page 174	
1. Verify the LSP on page 176	<code>show mpls lsp</code> <code>show mpls lsp extensive</code> <code>show mpls lsp name <i>name</i></code> <code>show mpls lsp name <i>name</i> extensive</code>
2. Verify the LSP Route on the Transit Router on page 179	<code>show route table mpls.0</code>
3. Verify the LSP Route on the Ingress Router on page 180	<code>show route <i>destination</i></code>
4. Verify MPLS Labels with the traceroute Command on page 182	<code>traceroute <i>hostname</i></code>
5. Verify MPLS Labels with the ping Command on page 183	On the ingress router: <code>ping mpls rsvp <i>lsp-name</i> detail</code>
6. Verify the MPLS Configuration on page 184	<code>show configuration protocols mpls</code> <code>show configuration interfaces</code>

Table 15: Checklist for Checking the MPLS Layer (*continued*)

Tasks	Command or Action
7. Take Appropriate Action on page 187	<p>The following sequence of commands addresses the specific problem described in this topic:</p> <pre>edit edit protocols mpls [edit protocols mpls] show activate interface so-0/0/3.0 show commit</pre>
8. Verify the LSP Again on page 188	<pre>show mpls lsp extensive</pre>

Checking the MPLS Layer

Purpose After you have configured the label-switched path (LSP), issued the **show mpls lsp** command, and determined that there is an error, you might find that the error is not in the physical, data link, Internet Protocol (IP), interior gateway protocol (IGP), or Resource Reservation Protocol (RSVP) layers. Continue investigating the problem at the MPLS layer of the network.

[Figure 21 on page 175](#) illustrates the MPLS layer of the layered MPLS model.

Figure 21: Checking the MPLS Layer


BGP Layer	traceroute <i>host-name</i> show bgp summary show configuration protocols bgp show route <i>destination-prefix</i> detail show route receive protocol bgp <i>neighbor-address</i>
MPLS Layer	show mpls lsp show mpls lsp extensive show route table mpls.0 show route <i>address</i> traceroute <i>address</i> ping mpls rsvp <i>lsp-name</i> detail
RSVP Layer	show rsvp session show rsvp neighbor show rsvp interface
↙ IGP and IP Layers Functioning ↘	
OSPF Layer	show ospf neighbor show configuration protocols ospf show ospf interface
IS-IS Layer	show isis adjacency show configuration protocols isis show isis interface
IP Layer	show ospf neighbor extensive show interfaces terse
IP Layer	show isis adjacency extensive show interfaces terse
Data Link Layer	show interfaces extensive <i>"JUNOS Interfaces Operations Guide"</i>
Physical Layer	show interfaces show interfaces terse ping <i>host</i>

g015547

With the MPLS layer, you check whether the LSP is up and functioning correctly. If the network is not functioning at this layer, the LSP does not work as configured.

Figure 22 on page 175 illustrates the MPLS network used in this topic.

Figure 22: MPLS Network Broken at the MPLS Layer



g015541

The network shown in Figure 22 on page 175 is a fully meshed configuration where every directly connected interface can receive and send packets to every other similar interface.

The LSP in this network is configured to run from ingress router **R1**, through transit router **R3**, to egress router **R6**. In addition, a reverse LSP is configured to run from **R6** through **R3** to **R1**, creating bidirectional traffic.

However, in this example, the reverse LSP is down without a path from **R6** to **R1**.

The cross shown in [Figure 22 on page 175](#) indicates where the LSP is broken. Some possible reasons the LSP is broken might include an incorrectly configured MPLS protocol, or interfaces that are incorrectly configured for MPLS.

In the network shown in [Figure 22 on page 175](#), a configuration error on egress router **R6** prevents the LSP from traversing the network as expected.

To check the MPLS layer, follow these steps:

1. [Verify the LSP on page 176](#)
2. [Verify the LSP Route on the Transit Router on page 179](#)
3. [Verify the LSP Route on the Ingress Router on page 180](#)
4. [Verify MPLS Labels with the traceroute Command on page 182](#)
5. [Verify MPLS Labels with the ping Command on page 183](#)
6. [Verify the MPLS Configuration on page 184](#)
7. [Take Appropriate Action on page 187](#)
8. [Verify the LSP Again on page 188](#)

Verify the LSP

Purpose Typically, you use the **show mpls lsp extensive** command to verify the LSP. However for quick verification of the LSP state, use the **show mpls lsp** command. If the LSP is down, use the **extensive** option (**show mpls lsp extensive**) as a follow-up. If your network has numerous LSPs, you might consider specifying the name of the LSP, using the **name** option (**show mpls lsp name *name*** or **show mpls lsp name *name* extensive**).

Action To verify that the LSP is up, enter some or all of the following commands from the ingress router:

```
user@host> show mpls lsp
user@host> show mpls lsp extensive
user@host> show mpls lsp name name
user@host> show mpls lsp name name extensive
```

Sample Output 1

```

user@R1> show mpls lsp
Ingress LSP: 1 sessions
  To          From          State Rt ActivePath      P      LSPname
  10.0.0.6     10.0.0.1      Dn    0  -              P      R1-to-R6
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@R3> show mpls lsp
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

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

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

user@R6> show mpls lsp
Ingress LSP: 1 sessions
  To          From          State Rt ActivePath      P      LSPname
  10.0.0.1     10.0.0.6      Dn    0  -              P      R6-to-R1
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

```

Sample Output 2

```

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

10.0.0.6
  From: 10.0.0.1, State: Dn, ActiveRoute: 0, LSPname: R1-to-R6
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary          State: Dn
    Will be enqueued for recomputation in 22 second(s).
    1 Nov  2 14:43:38  CSPF failed: no route toward 10.0.0.6 [175 times]
  Created: Tue Nov  2 13:18:39 2004
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@R3> show mpls lsp extensive
Ingress LSP: 0 sessions

```

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

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

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

user@R6> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, State: Dn, ActiveRoute: 0, LSPname: R6-to-R1
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary                               State: Dn
    Will be enqueued for recomputation in 13 second(s).
    1 Nov  2 14:38:12  CSPF failed: no route toward 10.0.0.1 [177 times]
  Created: Tue Nov  2 13:12:22 2004
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
```

Sample Output 3

```
user@R1> show mpls lsp name R1-to-R6
Ingress LSP: 1 sessions
To          From          State Rt ActivePath      P    LSPname
10.0.0.6    10.0.0.1    Dn    0  -                R1-to-R6
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
```

Sample Output 4

```
user@R1> show mpls lsp name R1-to-R6 extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, State: Dn, ActiveRoute: 0, LSPname: R1-to-R6
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary                               State: Dn
    Will be enqueued for recomputation in 10 second(s).
    1 Nov  2 14:51:53 CSPF failed: no route toward 10.0.0.6[192 times]
  Created: Tue Nov  2 13:18:39 2004
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 Sample Output 1 shows a brief description of the state of the LSP for the ingress, transit, and egress routers. Output from ingress router **R1** and egress router **R6** shows that both LSPs are down, **R1-to-R6** and **R6-to-R1**. With the configured LSPs on **R1** and **R6**, we would expect egress LSP sessions on both **R1** and **R6**. In addition, transit router **R3** has no transit sessions.

Sample Output 2 shows all information about the LSPs, including all past state history and the reason why an LSP failed. Output from **R1** and **R6** indicates that there is no route to the destination because the Constrained Shortest Path First (CSPF) algorithm failed.

Sample Outputs 3 and 4 show examples of the output for the **show mpls lsp name** command with the **extensive** option. In this instance, the output is very similar to the **show mpls lsp** command because only one LSP is configured in the example network in [Figure 22 on page 175](#). However, in a large network with many LSPs configured, the results would be quite different between the two commands.

Verify the LSP Route on the Transit Router

Purpose If the LSP is up, the LSP route should appear in the **mpls.0** routing table. MPLS maintains an MPLS path routing table (**mpls.0**), which contains a list of the next label-switched router in each LSP. This routing table is used on transit routers to route packets to the next router along an LSP. If routes are not present in the output for the transit router, check the MPLS protocol configuration on the ingress and egress routers.

Action To verify the LSP route on the transit router, enter the following command from the transit router:

```
user@host> show route table mpls.0
```

Sample Output 1

```
user@R3> show route table mpls.0
mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
0          *[MPLS/0] 16w2d 21:52:40, metric 1
           Receive
1          *[MPLS/0] 16w2d 21:52:40, metric 1
           Receive
2          *[MPLS/0] 16w2d 21:52:40, metric 1
           Receive
```

Sample Output 2

```
user@R3> show route table mpls.0
mpls.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
0          *[MPLS/0] 16w2d 22:26:08, metric 1
           Receive
1          *[MPLS/0] 16w2d 22:26:08, metric 1
           Receive
2          *[MPLS/0] 16w2d 22:26:08, metric 1
           Receive
100864     *[RSVP/7] 00:07:23, metric 1
           > via so-0/0/2.0, label-switched-path R6-to-R1
100864(S=0) *[RSVP/7] 00:07:23, metric 1
           > via so-0/0/2.0, label-switched-path R6-to-R1
100880     *[RSVP/7] 00:07:01, metric 1
           > via so-0/0/3.0, label-switched-path R1-to-R6
100880(S=0) *[RSVP/7] 00:07:01, metric 1
           > via so-0/0/3.0, label-switched-path R1-to-R6
```

Meaning Sample Output 1 from transit router **R3** shows three route entries in the form of MPLS label entries. These MPLS labels are reserved MPLS labels defined in RFC 3032, and are always present in the **mpls.0** routing table, regardless of the state of the LSP. The incoming labels assigned by RSVP to the upstream neighbor are missing from the output, indicating that the LSP is down. For more information on MPLS label entries, see [“Checklist for Verifying LSP Use” on page 77](#).

In contrast, Sample Output 2 shows the MPLS labels and routes for a correctly configured LSP. The three reserved MPLS labels are present, and the four other entries represent the incoming labels assigned by RSVP to the upstream neighbor. These four entries represent two routes. There are two entries per route because the stack values in the MPLS header may be different. For each route, the second entry **100864 (S=0)** and **100880 (S=0)** indicates that the stack depth is not 1, and additional label values are included in the packet. In contrast, the first entry, **100864** and **100880** has an inferred **S=1** value which indicates a stack depth of 1 and makes each label the last label in that particular packet. The dual entries indicate that this is the penultimate router. For more information on MPLS label stacking, see RFC 3032, *MPLS Label Stack Encoding*.

Verify the LSP Route on the Ingress Router

Purpose Check whether the LSP route is included in the active entries in the **inet.3** routing table for the specified address.

Action To verify the LSP route, enter the following command from the ingress router:

```
user@host> show route destination
```

Sample Output 1

```
user@R1> show route 10.0.0.6
inet.0 : 27 destinations, 27 routes (27 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.0.0.6/32      *[IS-IS/18] 6d 01:41:37, metric 20
                  to 10.1.12.2 via so-0/0/0.0
                  > to 10.1.15.2 via so-0/0/1.0
                  to 10.1.13.2 via so-0/0/2.0

user@R6> show route 10.0.0.1

inet.0 : 28 destinations, 28 routes (28 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.0.0.1/32      *[IS-IS/18] 5d 01:01:38, metric 20
                  to 10.1.56.1 via so-0/0/0.0
                  > to 10.1.26.1 via so-0/0/2.0
                  to 10.1.36.1 via so-0/0/3.0
```

Sample Output 2

```
user@R1> show route 10.0.0.6
inet.0: 28 destinations, 28 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.6/32      *[IS-IS/18] 6d 02:13:42, metric 20
                  to 10.1.12.2 via so-0/0/0.0
                  > to 10.1.15.2 via so-0/0/1.0
                  to 10.1.13.2 via so-0/0/2.0

inet.3 : 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.6/32      *[RSVP/7] 00:08:07, metric 20
                  > via so-0/0/2.0, label-switched-path R1-to-R6

user@R6> show route 10.0.0.1

inet.0: 29 destinations, 29 routes (28 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.1/32      *[IS-IS/18] 5d 01:34:03, metric 20
                  to 10.1.56.1 via so-0/0/0.0
                  > to 10.1.26.1 via so-0/0/2.0
                  to 10.1.36.1 via so-0/0/3.0

inet.3 : 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.1/32      *[RSVP/7] 00:10:39, metric 20
                  > via so-0/0/3.0, label-switched-path R6-to-R1
```

Meaning Sample Output 1 shows entries in the **inet.0** routing table only. The **inet.3** routing table is missing from the output because the LSP is not working. The **inet.0** routing table is used by interior gateway protocols (IGPs) and Border Gateway Protocol (BGP) to store routing

information. In this case, the IGP is Intermediate System-to-Intermediate System (IS-IS). For more information on the **inet.0** routing table, see the *Junos MPLS Applications Configuration Guide*.

If the LSP was working, we would expect to see entries that include the LSP in the **inet.3** routing table. The **inet.3** routing table is used on ingress routers to route BGP packets to the destination egress router. BGP uses the **inet.3** routing table on the ingress router to help resolve next-hop addresses. BGP is configured in the example network shown in [Figure 22 on page 175](#).

Sample Output 2 shows output you should receive when the LSP is up. The output shows both the **inet.0** and **inet.3** routing tables, indicating that LSPs **R1-to-R6** and **R6-to-R1** are available.

Verify MPLS Labels with the traceroute Command

Purpose Display the route packets take to a BGP destination where the BGP next hop for that route is the LSP egress address. By default, BGP uses the **inet.0** and the **inet.3** routing tables to resolve the next-hop address. When the next-hop address of the BGP route is not the router ID of the egress router, traffic is mapped to IGP routes, not to the LSP. Use the **traceroute** command as a debugging tool to determine whether the LSP is being used to forward traffic.

Action To verify MPLS labels, enter the following command from the ingress router:

```
user@host> traceroute hostname
```

Sample Output 1

```
user@R1> traceroute 100.100.6.1
traceroute to 100.100.6.1 (100.100.6.1), 30 hops max, 40 byte packets
 1  10.1.12.2 (10.1.12.2)  0.627 ms  0.561 ms  0.520 ms
 2  10.1.26.2 (10.1.26.2)  0.570 ms !N  0.558 ms !N  4.879 ms !N

user@R6> traceroute 100.100.1.1
traceroute to 100.100.1.1 (100.100.1.1), 30 hops max, 40 byte packets
 1  10.1.26.1 (10.1.26.1)  0.630 ms  0.545 ms  0.488 ms
 2  10.1.12.1 (10.1.12.1)  0.551 ms !N  0.557 ms !N  0.526 ms !N
```

Sample Output 2

```
user@R1> traceroute 100.100.6.1
to 100.100.6.1 (100.100.6.1), 30 hops max, 40 byte packets
 1  10.1.13.2 (10.1.13.2)  0.866 ms  0.746 ms  0.724 ms
    MPLS Label=100912 CoS=0 TTL=1 S=1
 2  10.1.36.2 (10.1.36.2)  0.577 ms !N  0.597 ms !N  0.546 ms !N

user@R6> traceroute 100.100.1.1
traceroute to 100.100.1.1 (100.100.1.1), 30 hops max, 40 byte packets
 1  10.1.36.1 (10.1.36.1)  0.802 ms  0.716 ms  0.688 ms
    MPLS Label=100896 CoS=0 TTL=1 S=1
 2  10.1.13.1 (10.1.13.1)  0.570 ms !N  0.568 ms !N  0.546 ms !N
```

Meaning Sample Output 1 shows that BGP traffic is not using the LSP, consequently MPLS labels do not appear in the output. Instead of using the LSP, BGP traffic is using the IGP (IS-IS,

in the example network in [Figure 22 on page 175](#)) to reach the BGP next-hop LSP egress address. The Junos OS default behavior uses LSPs for BGP traffic when the BGP next hop equals the LSP egress address.

Sample Output 2 is an example of output for a correctly configured LSP. The output shows MPLS labels, indicating that BGP traffic is using the LSP to reach the BGP next hop.

Verify MPLS Labels with the ping Command

Purpose When you ping a specific LSP, you check that echo requests are sent over the LSP as MPLS packets.

Action To verify MPLS labels, enter the following command from the ingress router to ping the egress router:

```
user@host> ping mpls rsvp lsp-name detail
```

For example:

```
user@R1> ping mpls rsvp R1-to-R6 detail
```

Sample Output 1

```
user@R1> ping mpls rsvp R1-to-R6 detail
LSP R1-to-R6 - LSP has no active path, exiting.
```

```
user@R6> ping mpls rsvp R6-to-R1 detail
LSP R6-to-R1 - LSP has no active path, exiting.
```

Sample Output 2

```
user@R1> traceroute 10.0.0.6
traceroute to 10.0.0.6 (10.0.0.6), 30 hops max, 40 byte packets
 1 10.1.15.2 (10.1.15.2) 0.708 ms 0.613 ms 0.576 ms
 2 10.0.0.6 (10.0.0.6) 0.763 ms 0.708 ms 0.700 ms
```

```
user@R1> ping mpls rsvp R1-to-R6 detail
Request for seq 1, to interface 69, label 100880
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 69, label 100880
Reply for seq 2, return code: Egress-ok
Request for seq 3, to interface 69, label 100880
Reply for seq 3, return code: Egress-ok
Request for seq 4, to interface 69, label 100880
Reply for seq 4, return code: Egress-ok
Request for seq 5, to interface 69, label 100880
Reply for seq 5, return code: Egress-ok
```

```
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

```
user@R6> ping mpls rsvp R6-to-R1 detail
Request for seq 1, to interface 70, label 100864
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 70, label 100864
Reply for seq 2, return code: Egress-ok
Request for seq 3, to interface 70, label 100864
Reply for seq 3, return code: Egress-ok
Request for seq 4, to interface 70, label 100864
Reply for seq 4, return code: Egress-ok
Request for seq 5, to interface 70, label 100864
Reply for seq 5, return code: Egress-ok
```

```
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

Meaning Sample Output 1 shows that the LSP does not have an active path to forward echo requests, indicating that the LSP is down.

Sample Output 2 is an example of output you should receive when the LSP is up and forwarding packets.

Verify the MPLS Configuration

Purpose After you have checked the transit and ingress routers, used the **traceroute** command to verify the BGP next hop, and used the **ping** command to verify the active path, you can check for problems with the MPLS configuration at the **[edit protocols mpls]** and **[edit interfaces]** hierarchy levels.

Action To verify the MPLS configuration, enter the following commands from the ingress, transit, and egress routers:

```
user@host> show configuration protocols mpls
```

```
user@host> show configuration interfaces
```

Sample Output 1

```
user@R1> show configuration protocols mpls
label-switched-path R1-to-R6 {
    to 10.0.0.6;
}
inactive: interface so-0/0/0.0;
inactive: interface so-0/0/1.0;
interface so-0/0/2.0;
interface fxp0.0 {
    disable;
}

user@R3> show configuration protocols mpls
interface fxp0.0 {
    disable;
}
inactive: interface so-0/0/0.0;
inactive: interface so-0/0/1.0;
interface so-0/0/2.0;
interface so-0/0/3.0;

user@R6> show configuration protocols mpls
label-switched-path R6-to-R1 {
    to 10.0.0.1;
}
inactive: interface so-0/0/0.0;
inactive: interface so-0/0/1.0;
inactive: interface so-0/0/2.0;
inactive: interface so-0/0/3.0; <<< Incorrectly configured
```

Sample Output 2

```
user@R6> show configuration interfaces
so-0/0/0 {
    unit 0 {
        family inet {
            address 10.1.56.2/30;
        }
        family iso;
        family mpls;
    }
}
so-0/0/1 {
    unit 0 {
        family inet {
            address 10.1.46.2/30;
        }
        family iso;
        family mpls;
    }
}
so-0/0/2 {
    unit 0 {
        family inet {
            address 10.1.26.2/30;
        }
        family iso;
        family mpls;
    }
}
```

```

}
so-0/0/3 {
  unit 0 {
    family inet {
      address 10.1.36.2/30;
    }
    family iso;
    family mpls;
  }
}
fxp0 {
  unit 0 {
    family inet {
      address 192.168.70.148/21;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.0.0.6/32;
      address 127.0.0.1/32;
    }
    family iso {
      address 49.0003.1000.0000.0006.00;
    }
  }
}
}

```

Meaning Sample Output 1 from the ingress, transit, and egress routers shows that the configuration of interfaces on egress router **R6** is incorrect. Interface **so-0/0/3.0** is included as inactive at the **[edit protocols mpls]** hierarchy level, when it should be active because it is the interface through which the LSP travels.

Sample Output 2 shows that interfaces are correctly configured for MPLS on egress router **R6**. The interfaces are also correctly configured on the ingress and transit routers (not shown).

Take Appropriate Action

Problem Depending on the error you encountered in your investigation, you must take the appropriate action to correct the problem. In this example, an interface is incorrectly configured at the **[edit protocols mpls]** hierarchy level on egress router **R6**.

Solution To correct the error in this example, follow these steps:

1. Activate the interface in the MPLS protocol configuration on egress router **R6**:

```

user@R6> edit
user@R6# edit protocols mpls
[edit protocols mpls]
user@R6# show
user@R6# activate interface so-0/0/3.0

```

2. Verify and commit the configuration:

```

[edit protocols mpls]

```

```
user@R6# show
user@R6# commit
```

Sample Output

```
user@R6> edit
Entering configuration mode

[edit]
user@R6# edit protocols mpls

[edit protocols mpls]
user@R6# show
label-switched-path R6-to-R1 {
    to 10.0.0.1;
}
inactive: interface so-0/0/0.0;
inactive: interface so-0/0/1.0;
inactive: interface so-0/0/2.0;
inactive: interface so-0/0/3.0; <<< Incorrectly configured interface

[edit protocols mpls]
user@R6# activate interface so-0/0/3

[edit protocols mpls]
user@R6# show
label-switched-path R6-to-R1 {
    to 10.0.0.1;
}
inactive: interface so-0/0/0.0;
inactive: interface so-0/0/1.0;
inactive: interface so-0/0/2.0;
interface so-0/0/3.0; <<< Correctly configured interface

[edit protocols mpls]
user@R6# commit
commit complete
```

Meaning The sample output shows that the incorrectly configured interface **so-0/0/3.0** on egress router **R6** is now activated at the **[edit protocols mpls]** hierarchy level. The LSP can now come up.

Verify the LSP Again

Purpose After taking the appropriate action to correct the error, the LSP needs to be checked again to confirm that the problem in the BGP layer has been resolved.

Action To verify the LSP again, enter the following command from the ingress, transit, and egress routers:

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


Sample Output

```

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

10.0.0.6
  From: 10.0.0.1,  State: Up,  ActiveRoute: 1,  LSPname: R1-to-R6
  ActivePath:  (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                      State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
10.1.13.2 S 10.1.36.2 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

      10.1.13.2 10.1.36.2
6 Nov  2 15:48:52 Selected as active path
5 Nov  2 15:48:52 Record Route:  10.1.13.2 10.1.36.2
4 Nov  2 15:48:52 Up
3 Nov  2 15:48:52 Originate Call
2 Nov  2 15:48:52 CSPF: computation result accepted
1 Nov  2 15:48:22 CSPF failed: no route toward 10.0.0.6[308 times]
Created: Tue Nov  2 13:18:39 2004
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6,  LSPstate: Up ,  ActiveRoute: 0
  LSPname: R6-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: Tue Nov  2 15:48:30 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 39106 protocol 0
  PATH rcvfrom: 10.1.13.2 (so-0/0/2.0) 10 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.1.36.2 10.1.13.2 <self>
Total 1 displayed, Up 1, Down 0

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

user@R3> show mpls lsp extensive
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

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

Transit LSP: 2 sessions

10.0.0.1
  From: 10.0.0.6, LSPstate: Up, ActiveRoute: 1
  LSPname: R6-to-R1, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3

```

```

Resv style: 1 FF, Label in: 100864, Label out: 3
Time left: 123, Since: Tue Nov  2 15:35:41 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 1 receiver 39106 protocol 0
PATH rcvfrom: 10.1.36.2 (so-0/0/3.0) 10 pkts
Adspec: received MTU 1500 sent MTU 1500
PATH sentto: 10.1.13.1 (so-0/0/2.0) 10 pkts
RESV rcvfrom: 10.1.13.1 (so-0/0/2.0) 10 pkts
Explct route: 10.1.13.1
Record route: 10.1.36.2 <self> 10.1.13.1

```

10.0.0.6

```

From: 10.0.0.1, LSPstate: Up, ActiveRoute: 1
LSPname: R1-to-R6, LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: 3
Resv style: 1 FF, Label in: 100880, Label out: 3
Time left: 145, Since: Tue Nov  2 15:36:03 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 1 receiver 48015 protocol 0
PATH rcvfrom: 10.1.13.1 (so-0/0/2.0) 10 pkts
Adspec: received MTU 1500 sent MTU 1500
PATH sentto: 10.1.36.2 (so-0/0/3.0) 10 pkts
RESV rcvfrom: 10.1.36.2 (so-0/0/3.0) 10 pkts
Explct route: 10.1.36.2
Record route: 10.1.13.1 <self> 10.1.36.2
Total 2 displayed, Up 2, Down 0

```

user@R6> show mpls lsp extensive

Ingress LSP: 1 sessions

10.0.0.1

```

From: 10.0.0.6, State: Up, ActiveRoute: 1, LSPname: R6-to-R1
ActivePath: (primary)
LoadBalance: Random
Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary State: Up
Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
10.1.36.1 S 10.1.13.1 S
Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

10.1.36.1 10.1.13.1
6 Nov  2 15:41:44 Selected as active path
5 Nov  2 15:41:44 Record Route: 10.1.36.1 10.1.13.1
4 Nov  2 15:41:44 Up
3 Nov  2 15:41:44 Originate Call
2 Nov  2 15:41:44 CSPF: computation result accepted
1 Nov  2 15:41:14 CSPF failed: no route toward 10.0.0.1[306 times]
Created: Tue Nov  2 13:12:21 2004
Total 1 displayed, Up 1, Down 0

```

Egress LSP: 1 sessions

10.0.0.6

```

From: 10.0.0.1, LSPstate: Up, ActiveRoute: 0
LSPname: R1-to-R6, LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: -
Resv style: 1 FF, Label in: 3, Label out: -
Time left: 157, Since: Tue Nov  2 15:42:06 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500

```

```
Port number: sender 1 receiver 48015 protocol 0
PATH rcvfrom: 10.1.36.1 (so-0/0/3.0) 11 pkts
Adspec: received MTU 1500
PATH sentto: localclient
RESV rcvfrom: localclient
Record route: 10.1.13.1 10.1.36.1 <self>
Total 1 displayed, Up 1, Down 0

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

Meaning Sample Output 1 from ingress router **R1** shows that LSP **R1-to-R6** has an active route to **R6** and the state is up.

Sample Output 2 from transit router **R3** shows that there are two transit LSP sessions, one from **R1** to **R6** and the other from **R6** to **R1**. Both LSPs are up.

Sample Output 3 from egress router **R6** shows that the LSP is up and the active route is the primary route. The LSP is now traversing the network along the expected path, from **R1** through **R3** to **R6**, and the reverse LSP, from **R6** through **R3** to **R1**.

CHAPTER 12

Checking the BGP Layer

This chapter describes how to check the Border Gateway Protocol (BGP) layer of the layered Multiprotocol Label Switching (MPLS) model.

- [Checklist for Checking the BGP Layer on page 193](#)
- [Checking the BGP Layer on page 194](#)

Checklist for Checking the BGP Layer

Problem This checklist provides the steps and commands for checking the BGP configuration of the Multiprotocol Label Switching (MPLS) network. The checklist provides links to an overview of the BGP configuration and more detailed information about the commands used to configure BGP. (See [Table 16 on page 193](#).)

Table 16: Checklist for Checking the BGP Layer

Tasks	Command or Action
“Checking the BGP Layer” on page 194	
1. Check That BGP Traffic Is Using the LSP on page 195	<code>traceroute <i>hostname</i></code>
2. Check BGP Sessions on page 196	<code>show bgp summary</code>
3. Verify the BGP Configuration on page 198	<code>show configuration</code>
4. Examine BGP Routes on page 204	<code>show route <i>destination-prefix</i> detail</code>
5. Verify Received BGP Routes on page 205	<code>show route receive protocol bgp <i>neighbor-address</i></code>
6. Take Appropriate Action on page 206	<p>The following sequence of commands addresses the specific problem described in this topic:</p> <pre>[edit] edit protocols bgp [edit protocols bgp] show set local-address 10.0.0.1 delete group internal neighbor 10.1.36.2 show commit</pre>

Table 16: Checklist for Checking the BGP Layer (*continued*)

Tasks	Command or Action
7. Check That BGP Traffic Is Using the LSP Again on page 207	<code>traceroute hostname</code>

Checking the BGP Layer

Purpose After you have configured the label-switched path (LSP) and determined that it is up, and configured BGP and determined that sessions are established, ensure that BGP is using the LSP to forward traffic.

Figure 23 on page 194 illustrates the BGP layer of the layered MPLS model.

Figure 23: Checking the BGP Layer


BGP Layer	traceroute <i>host-name</i> show bgp summary show configuration protocols bgp show route <i>destination-prefix</i> detail show route receive protocol bgp <i>neighbor-address</i>
MPLS Layer	show mpls lsp show mpls lsp extensive show route table mpls.0 show route <i>address</i> traceroute <i>address</i> ping mpls rsvp <i>lsp-name</i> detail
RSVP Layer	show rsvp session show rsvp neighbor show rsvp interface
<div>↙</div> IGP and IP Layers Functioning <div>↘</div>	
OSPF Layer show ospf neighbor show configuration protocols ospf show ospf interface	IS-IS Layer show isis adjacency show configuration protocols isis show isis interface
IP Layer show ospf neighbor extensive show interfaces terse	IP Layer show isis adjacency extensive show interfaces terse
Data Link Layer	show interfaces extensive <i>"JUNOS Interfaces Operations Guide"</i>
Physical Layer	show interfaces show interfaces terse ping <i>host</i>

g015548

When you check the BGP layer, you verify that the route is present and active, and more importantly, you ensure that the next hop is the LSP. There is no point in checking the BGP layer unless the LSP is established, because BGP uses the MPLS LSP to forward traffic. If the network is not functioning at the BGP layer, the LSP does not work as configured.

Figure 24 on page 195 illustrates the MPLS network used in this topic.

Figure 24: MPLS Network Broken at the BGP Layer



The network shown in [Figure 24 on page 195](#) is a fully meshed configuration where every directly connected interface can receive and send packets to every other similar interface. The LSP in this network is configured to run from ingress router **R1**, through transit router **R3**, to egress router **R6**. In addition, a reverse LSP is configured to run from **R6** through **R3** to **R1**, creating bidirectional traffic.

The cross shown in [Figure 24 on page 195](#) indicates where BGP is not being used to forward traffic through the LSP. Possible reasons for the LSP not working correctly are that the destination IP address of the LSP does not equal the BGP next hop or that BGP is not configured properly.

To check the BGP layer, follow these steps:

1. [Check That BGP Traffic Is Using the LSP on page 195](#)
2. [Check BGP Sessions on page 196](#)
3. [Verify the BGP Configuration on page 198](#)
4. [Examine BGP Routes on page 204](#)
5. [Verify Received BGP Routes on page 205](#)
6. [Take Appropriate Action on page 206](#)
7. [Check That BGP Traffic Is Using the LSP Again on page 207](#)

Check That BGP Traffic Is Using the LSP

Purpose At this level of the troubleshooting model, BGP and the LSP may be up, however BGP traffic might not be using the LSP to forward traffic.

Action To verify that BGP traffic is using the LSP, enter the following Junos OS command-line interface (CLI) operational mode command from the ingress router:

```
user@host> traceroute hostname
```

Sample Output

```
user@R1> traceroute 100.100.6.1
traceroute to 100.100.6.1 (100.100.6.1), 30 hops max, 40 byte packets
 1  10.1.13.2 (10.1.13.2)  0.653 ms  0.590 ms  0.543 ms
 2  10.1.36.2 (10.1.36.2)  0.553 ms !N  0.552 ms !N  0.537 ms !N

user@R6> traceroute 100.100.1.1
traceroute to 100.100.1.1 (100.100.1.1), 30 hops max, 40 byte packets
 1  10.1.36.1 (10.1.36.1)  0.660 ms  0.551 ms  0.526 ms
 2  10.1.13.1 (10.1.13.1)  0.568 ms !N  0.553 ms !N  0.536 ms !N
```

Meaning The sample output shows that BGP traffic is not using the LSP, consequently MPLS labels do not appear in the output. Instead of using the LSP, BGP traffic is using the interior gateway protocol (IGP) to reach the BGP next-hop LSP egress address for **R6** and **R1**. The Junos OS default is to use LSPs for BGP traffic when the BGP next hop equals the LSP egress address.

Check BGP Sessions

Purpose Display summary information about BGP and its neighbors to determine if routes are received from peers in the autonomous system (AS). When a BGP session is established, the peers are exchanging update messages.

Action To check that BGP sessions are up, enter the following Junos OS CLI operational mode command from the ingress router:

```
user@host> show bgp summary
```


Sample Output 1

```

user@R1> show bgp summary
Groups: 1 Peers: 6 Down peers: 1
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet.0 1 1 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Damped...
10.0.0.2 65432 11257 11259 0 0 3d 21:49:57 0/0/0
0/0/0
10.0.0.3 65432 11257 11259 0 0 3d 21:49:57 0/0/0
0/0/0
10.0.0.4 65432 11257 11259 0 0 3d 21:49:57 0/0/0
0/0/0
10.0.0.5 65432 11257 11260 0 0 3d 21:49:57 0/0/0
0/0/0
10.0.0.6 65432 4 4572 0 13d 21:46:59 Active
10.1.36.2 65432 11252 11257 0 0 3d 21:46:49 1/1/0
0/0/0

```

Sample Output 2

```

user@R1> show bgp summary
Groups: 1 Peers: 5 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet.0 1 1 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Damped...
10.0.0.2 65432 64 68 0 0 32:18 0/0/0
0/0/0
10.0.0.3 65432 64 67 0 0 32:02 0/0/0
0/0/0
10.0.0.4 65432 64 67 0 0 32:10 0/0/0
0/0/0
10.0.0.5 65432 64 67 0 0 32:14 0/0/0
0/0/0
10.0.0.6 65432 38 39 0 1 18:02 1/1/0
0/0/0

```

Meaning Sample Output 1 shows that one peer (egress router **10.0.0.6**) is not established, as indicated by the **Down Peers: 1** field. The last column (**State|#Active/Received/Damped**) shows that peer **10.0.0.6** is active, indicating that it is not established. All other peers are established as indicated by the number of active, received, and damped routes. For example, **0/0/0** for peer **10.0.0.2** indicates that no BGP routes were active or received in the routing table, and no BGP routes were damped; **1/1/0** for peer **10.1.36.2** indicates that one BGP route was active and received in the routing table, and no BGP routes were damped.

If the output of the **show bgp summary** command of an ingress router shows that a neighbor is down, check the BGP configuration. For information on checking the BGP configuration, see [“Verify the BGP Configuration” on page 198](#).

Sample Output 2 shows output from ingress router **R1** after the BGP configurations on **R1** and **R6** were corrected in [“Take Appropriate Action” on page 206](#). All BGP peers are established and one route is active and received. No BGP routes were damped.

If the output of the **show bgp summary** command shows that a neighbor is up but packets are not being forwarded, check for received routes from the egress router. For information on checking the egress router for received routes, see [“Verify Received BGP Routes” on page 205](#).

Verify the BGP Configuration

Purpose For BGP to run on the router, you must define the local AS number, configure at least one group, and include information about at least one peer in the group (the peer's IP address and AS number). When BGP is part of an MPLS network, you must ensure that the LSP is configured with a destination IP address equal to the BGP next hop in order for BGP routes to be installed with the LSP as the next hop for those routes.

Action To verify the BGP configuration, enter the following Junos OS CLI operational mode command:

```
user@host> show configuration
```

Sample Output 1

```

user@R1> show configuration
[...Output truncated...]
interfaces {
  so-0/0/0 {
    unit 0 {
      family inet {
        address 10.1.12.1/30;
      }
      family iso;
      family mpls;
    }
  }
  so-0/0/1 {
    unit 0 {
      family inet {
        address 10.1.15.1/30;
      }
      family iso;
      family mpls;
    }
  }
  so-0/0/2 {
    unit 0 {
      family inet {
        address 10.1.13.1/30;
      }
      family iso;
      family mpls;
    }
  }
  fxp0 {
    unit 0 {
      family inet {
        address 192.168.70.143/21;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.0.0.1/32;
      }
      family iso {
        address 49.0004.1000.0000.0001.00;
      }
    }
  }
}
routing-options {
  [...Output truncated...]
  route 100.100.1.0/24 reject;
}
router-id 10.0.0.1;
autonomous-system 65432;
}
protocols {
  rsvp {
    interface so-0/0/0.0;
  }
}

```

```

        interface so-0/0/1.0;
        interface so-0/0/2.0;
        interface fxp0.0 {
            disable;
        }
    }
    mpls {
        label-switched-path R1-to-R6 {
            to 10.0.0.6; <<< destination address of the LSP
        }
        inactive: interface so-0/0/0.0;
        inactive: interface so-0/0/1.0;
        interface so-0/0/2.0;
        interface fxp0.0 {
            disable;
        }
    }
    bgp {
        export send-statics; <<< missing local-address statement
        group internal {
            type internal;
            neighbor 10.0.0.2;
            neighbor 10.0.0.5;
            neighbor 10.0.0.4;
            neighbor 10.0.0.6;
            neighbor 10.0.0.3;
            neighbor 10.1.36.2; <<< incorrect interface address
        }
    }
    isis {
        level 1 disable;
        interface so-0/0/0.0;
        interface so-0/0/1.0;
        interface so-0/0/2.0;
        interface all {
            level 2 metric 10;
        }
        interface fxp0.0 {
            disable;
        }
        interface lo0.0 {
            passive;
        }
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface so-0/0/0.0;
            interface so-0/0/1.0;
            interface so-0/0/2.0;
            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

```

user@R6> show configuration
[...Output truncated...]
interfaces {
  so-0/0/0 {
    unit 0 {
      family inet {
        address 10.1.56.2/30;
      }
      family iso;
      family mpls;
    }
  }
  so-0/0/1 {
    unit 0 {
      family inet {
        address 10.1.46.2/30;
      }
      family iso;
      family mpls;
    }
  }
  so-0/0/2 {
    unit 0 {
      family inet {
        address 10.1.26.2/30;
      }
      family iso;
      family mpls;
    }
  }
  so-0/0/3 {
    unit 0 {
      family inet {
        address 10.1.36.2/30;
      }
      family iso;
      family mpls;
    }
  }
  fxp0 {
    unit 0 {
      family inet {
        address 192.168.70.148/21;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.0.0.6/32;
        address 127.0.0.1/32;
      }
      family iso {

```

```

        address 49.0004.1000.0000.0006.00;
    }
}
}
routing-options {
    [...Output truncated...]
    route 100.100.6.0/24 reject;
}
    router-id 10.0.0.6;
    autonomous-system 65432;
}
protocols {
    rsvp {
        interface so-0/0/0.0;
        interface so-0/0/1.0;
        interface so-0/0/2.0;
        interface so-0/0/3.0;
        interface fxp0.0 {
            disable;
        }
    }
    mpls {
        label-switched-path R6-to-R1 {
            to 10.0.0.1; <<< destination address of the reverse LSP
        }
        inactive: interface so-0/0/0.0;
        inactive: interface so-0/0/1.0;
        inactive: interface so-0/0/2.0;
        interface so-0/0/3.0;
    }
    bgp {
        group internal {
            type internal;
            export send-statics; <<< missing local-address statement
            neighbor 10.0.0.2;
            neighbor 10.0.0.3;
            neighbor 10.0.0.4;
            neighbor 10.0.0.5;
            neighbor 10.0.0.1;
            neighbor 10.1.13.1; <<< incorrect interface address
        }
    }
    isis {
        level 1 disable;
        interface all {
            level 2 metric 10;
        }
        interface fxp0.0 {
            disable;
        }
        interface lo0.0 {
            passive;
        }
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface so-0/0/0.0;
            interface so-0/0/1.0;
            interface so-0/0/2.0;

```

```

        interface so-0/0/3.0;
        interface lo0.0 {
            passive;
        }
    }
}
policy-options {
    policy-statement send-statics {
        term statics {
            from {
                route-filter 100.100.6.0/24 exact;
            }
            then accept;
        }
    }
}

```

Meaning The sample output shows the BGP configurations on ingress router **R1** and egress router **R6**. Both configurations show the local AS (**65432**), one group (**internal**), and six peers configured. The underlying interior gateway protocol is IS-IS, and the relevant interfaces are configured to run IS-IS.



NOTE: In this configuration, the RID is manually configured to avoid any duplicate RID problems, and all interfaces configured with BGP include the family inet statement at the [edit interfaces *type-fpc/pic/port* unit *logical-unit-number*] hierarchy level.

Sample output for ingress router **R1** and egress router **R6** shows that the BGP protocol configuration is missing the **local-address** statement for the internal group. When the **local-address** statement is configured, BGP packets are forwarded from the local router loopback (**lo0**) interface address, which is the address to which BGP peers are peering. If the **local-address** statement is not configured, BGP packets are forwarded from the outgoing interface address, which does not match the address to which BGP peers are peering, and BGP does not come up.

On the ingress router, the IP address (**10.0.0.1**) in the **local-address** statement should be the same as the address configured for the LSP on the egress router (**R6**) in the **to** statement at the [edit protocols mpls label-switched-path *lsp-path-name*] hierarchy level. BGP uses this address, which is identical to the LSP address, to forward BGP traffic through the LSP.

In addition, the BGP configuration on **R1** includes two IP addresses for **R6**, an interface address (**10.1.36.2**) and a loopback (**lo0**) interface address (**10.0.0.6**), resulting in the LSP destination address (**10.0.0.6**) not matching the BGP next-hop address (**10.1.36.2**). The BGP configuration on **R6** also includes two IP addresses for **R1**, an interface address (**10.1.13.1**) and a loopback (**lo0**) interface address, resulting in the reverse LSP destination address (**10.0.0.1**) not matching the BGP next-hop address (**10.1.13.1**).

In this instance, because the **local-address** statement is missing in the BGP configurations of both routers and the LSP destination address does not match the BGP next-hop address, BGP is not using the LSP to forward traffic.

Examine BGP Routes

Purpose You can examine the BGP path selection process to determine the single, active path when BGP receives multiple routes to the same destination. In this step, we examine the reverse LSP **R6-to-R1**, making **R6** the ingress router for that LSP.

Action To examine BGP routes and route selection, enter the following Junos OS CLI operational mode command:

```
user@host> show route destination-prefix detail
```

Sample Output 1

```
user@R6> show route 100.100.1.1 detail
inet.0: 30 destinations, 46 routes (29 active, 0 holddown, 1 hidden)
100.100.1.0/24 (1 entry, 1 announced)
  *BGP      Preference: 170/-101
             Source: 10.1.13.1
             Next hop: via so-0/0/3.0, selected
             Protocol next hop: 10.1.13.1 Indirect next hop: 8671594 304
             State: <Active Int Ext>
             Local AS: 65432 Peer AS: 65432
             Age: 4d 5:15:39 Metric2: 2
             Task: BGP_65432.10.1.13.1+3048
             Announcement bits (2): 0-KRT 4-Resolve inet.0
             AS path: I
             Localpref: 100
             Router ID: 10.0.0.1
```

Sample Output 2

```
user@R6> show route 100.100.1.1 detail
inet.0: 30 destinations, 46 routes (29 active, 0 holddown, 1 hidden)
100.100.1.0/24 (1 entry, 1 announced)
  *BGP      Preference: 170/-101
             Source: 10.0.0.1
             Next hop: via so-0/0/3.0 weight 1, selected
             Label-switched-path R6-to-R1
             Label operation: Push 100000
             Protocol next hop: 10.0.0.1 Indirect next hop: 8671330 301
             State: <Active Int Ext>
             Local AS: 65432 Peer AS: 65432
             Age: 24:35 Metric2: 2
             Task: BGP_65432.10.0.0.1+179
             Announcement bits (2): 0-KRT 4-Resolve inet.0
             AS path: I
             Localpref: 100
             Router ID: 10.0.0.1
```

Meaning Sample Output 1 shows that the BGP next hop (10.1.13.1) does not equal the LSP destination address (10.0.0.1) in the **to** statement at the `[edit protocols mpls`

`label-switched-path label-switched-path-name]` hierarchy level when the BGP configuration of **R6** and **R1** is incorrect.

Sample Output 2, taken after the configurations on **R1** and **R6** are corrected, shows that the BGP next hop (**10.0.0.1**) and the LSP destination address (**10.0.0.1**) are the same, indicating that BGP can use the LSP to forward BGP traffic.

Verify Received BGP Routes

Purpose Display the routing information received on router **R6**, the ingress router for the reverse LSP **R6-to-R1**.

Action To verify that a particular BGP route is received on the egress router, enter the following Junos OS CLI operational mode command:

```
user@host> show route receive protocol bgp neighbor-address
```

Sample Output 1

```
user@R6> show route receive-protocol bgp 10.0.0.1
inet.0: 30 destinations, 46 routes (29 active, 0 holddown, 1 hidden)
<<< missing route
inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

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

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

__juniper_private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0
hidden)
```

Sample Output 2

```
user@R6> show route receive-protocol bgp 10.0.0.1
inet.0: 30 destinations, 46 routes (29 active, 0 holddown, 1 hidden)
  Prefix                Nexthop          MED      Lclpref    AS path
*100.100.1.0/24    10.0.0.1          100      I
inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

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

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

__juniper_private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0
hidden)
```

Meaning Sample Output 1 shows that ingress router **R6** (reverse LSP **R6-to-R1**) does not receive any BGP routes into the **inet.0** routing table when the BGP configurations of **R1** and **R6** are incorrect.

Sample Output 2 shows a BGP route installed in the **inet.0** routing table after the BGP configurations on **R1** and **R6** are corrected using [“Take Appropriate Action” on page 206](#).

Take Appropriate Action

Problem Depending on the error you encountered in your investigation, you must take the appropriate action to correct the problem. In this example, the ingress and egress routers are incorrectly configured for BGP to forward traffic using the LSP.

Solution To correct the errors in this example, follow these steps:

1. On ingress router R1, include the **local-address** statement and delete the incorrect interface address (repeat these steps on egress router R6):

```
[edit]
user@R1# edit protocols bgp
[edit protocols bgp]
user@R1# show
user@R1# set local-address 10.0.0.1
user@R1# delete group internal neighbor 10.1.36.2
```

2. Verify and commit the configuration:

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

Sample Output

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

[edit protocols bgp]
user@R1# show
export send-statics;
group internal {
    type internal;
    neighbor 10.0.0.2;
    neighbor 10.0.0.5;
    neighbor 10.0.0.4;
    neighbor 10.0.0.6;
    neighbor 10.0.0.3;
    neighbor 10.1.36.2;
}

[edit protocols bgp]
user@R1# set local-address 10.0.0.1

[edit protocols bgp]
user@R1# delete group internal neighbor 10.1.36.2

[edit protocols bgp]
user@R1# show
local-address 10.0.0.1;
export send-statics;
group internal {
    type internal;
    neighbor 10.0.0.2;
    neighbor 10.0.0.5;
    neighbor 10.0.0.4;
    neighbor 10.0.0.6;
    neighbor 10.0.0.3;
}

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

Meaning The sample output shows that the configuration of BGP on ingress router **R1** is now correct. BGP can now forward BGP traffic through the LSP.

Check That BGP Traffic Is Using the LSP Again

Purpose After taking the appropriate action to correct the error, the LSP needs to be checked again to confirm that BGP traffic is using the LSP and that the problem in the BGP layer has been resolved.

Action To verify that BGP traffic is using the LSP, enter the following Junos OS CLI operational mode command from the ingress router:

```
user@host> traceroute hostname
```

Sample Output

```
user@R1> traceroute 100.100.6.1
traceroute to 100.100.6.1 (100.100.6.1), 30 hops max, 40 byte packets
 1 10.1.13.2 (10.1.13.2) 0.858 ms 0.740 ms 0.714 ms
    MPLS Label=100016 CoS=0 TTL=1 S=1
 2 10.1.36.2 (10.1.36.2) 0.592 ms !N 0.564 ms !N 0.548 ms !N

user@R6> traceroute 100.100.1.1
traceroute to 100.100.1.1 (100.100.1.1), 30 hops max, 40 byte packets
 1 10.1.36.1 (10.1.36.1) 0.817 ms 0.697 ms 0.771 ms
    MPLS Label=100000 CoS=0 TTL=1 S=1
 2 10.1.13.1 (10.1.13.1) 0.581 ms !N 0.567 ms !N 0.544 ms !N
```

Meaning The sample output shows that MPLS labels are used to forward packets through the LSP. Included in the output is a label value (**MPLS Label=100016**), the time-to-live value (**TTL=1**), and the stack bit value (**S=1**).

The **MPLS Label** field is used to identify the packet to a particular LSP. It is a 20-bit field, with a maximum value of ($2^{20}-1$), approximately 1,000,000.

The time-to-live (TTL) value contains a limit on the number of hops that this MPLS packet can travel through the network (1). It is decremented at each hop, and if the TTL value drops below one, the packet is discarded.

The bottom of the stack bit value (**S=1**) indicates that is the last label in the stack and that this MPLS packet has one label associated with it. The MPLS implementation in the Junos OS supports a stacking depth of 3 on the M-series routers and up to 5 on the T-series routing platforms. For more information on MPLS label stacking, see RFC 3032, *MPLS Label Stack Encoding*.

MPLS labels appear in the sample output because the **traceroute** command is issued to a BGP destination where the BGP next hop for that route is the LSP egress address. The Junos OS by default uses LSPs for BGP traffic when the BGP next hop equals the LSP egress address.

If the BGP next hop does not equal the LSP egress address, the BGP traffic does not use the LSP, and consequently MPLS labels do not appear in the output for the **traceroute** command, as indicated in the sample output in [“Check BGP Sessions” on page 196](#).

PART 3

Index

- [Index on page 211](#)

Index

Symbols

#, comments in configuration statements.....	xvii
(), in syntax descriptions.....	xvii
< >, in syntax descriptions.....	xvii
[], in configuration statements.....	xvii
{ }, in configuration statements.....	xvii
(pipe), in syntax descriptions.....	xvii

A

activate interface command	187
adjacencies	
IP layer, verifying.....	125
IS-IS	
area.....	13
establishing	10
Level 2	9
verifying	14, 147
IS-IS, verifying	14, 147
management backbone	12
OSPF	
network	16
verifying.....	19
administrative groups, MPLS	49
area, OSPF	
backbone	137
configuring.....	17
autonomous-system statement	21

B

backbone, OSPF area	16
BGP.....	6
AS, defining.....	20
components of example network.....	20
configuration, verifying.....	198
delete group internal neighbor	
command.....	206
edit protocol bgp command.....	206
external neighbors.....	22
full-mesh network.....	7
group.....	20
local address, configuring.....	22

network topology, figure	20
peers and groups, configuring	21
peers, verifying	25
routing policy.....	20
sessions, checking	196
set local-address command.....	206
setting up.....	19
show bgp summary command.....	25, 196
show configuration command.....	198
show route detail command.....	204
show route receive protocol bgp	
command.....	205
traceroute command.....	207
traffic, verifying	207

BGP layer

broken network topology, figure	195
---------------------------------------	-----

Border Gateway Protocol See BGP

braces, in configuration statements.....	xvii
--	------

brackets

angle, in syntax descriptions.....	xvii
square, in configuration statements.....	xvii

C

checklists.....	117
data link layer.....	103
LSP	
state, determining	59
use, verifying	77
MPLS layered model.....	85
physical layer	93
RSVP protocol	
signal processing, verifying	69
Cisco High-level Data Link Control See HDLC	
comments, in configuration statements.....	xvii
Constrained Shortest Path First See CSPF	
conventions	
text and syntax.....	xvi
CSPF	6
curly braces, in configuration statements.....	xvii
customer support.....	xviii
contacting JTAC.....	xviii

D

data link layer	
broken network topology, figure	105
checklist for verifying.....	103
delete encapsulation command.....	111
encapsulation statement.....	111

interfaces	
deactivated.....	116
verifying.....	106
problems.....	105
show configuration protocols mpls	
command.....	115
show interfaces command.....	107
show interfaces extensive command	107
show mpls extensive command.....	112
show mpls lsp extensive command.....	106
database, Level 2.....	13
deactivate traceoptions command	74
delete encapsulation command	111
delete group internal neighbor command	206
documentation	
comments on.....	xviii

E

edit interface command.....	10
edit interfaces command	8
edit protocols bgp command.....	21, 206
edit protocols isis command.....	10, 17
edit protocols mpls command	187
edit protocols rsvp command	167
edit protocols rsvp traceoptions command	73
edit routing-options command.....	21, 23
egress router.....	6
encapsulation	
mode.....	104
type.....	106
encapsulation statement	111
export policy.....	23
external neighbors, BGP	22

F

families, show interfaces terse command.....	52
family inet statement.....	8
family iso statement.....	12, 14
family mpls statement.....	28, 99
family statement.....	8
font conventions.....	xvi
frame check sequence See FCS	
fxp0.0 statement.....	10, 17

G

graceful restart	50
group, configuring for BGP.....	21

H

history log, examining.....	72
host addresses.....	78

I

IBGP topology.....	20
IGP	9, 15
IGP layer	
checklist for verifying.....	117
inet.3 routing table.....	79
ingress router.....	6
inet.3 routing table, examining	79
show bgp summary command.....	197
interfaces	
configuration, incorrect for MPLS.....	49
configuring	
AS.....	21
IS-IS.....	10
MPLS.....	28
data link layer, verifying	106
deactivating	116
IS-IS, verifying.....	147
MPLS, verifying	47
OSPF, verifying.....	136
RSVP protocol, verifying.....	50
interfaces statement.....	8
interior gateway protocol See IGP	
interior gateway protocol IGP See IGP	
Intermediate System-to-Intermediate System See	
IS-IS protocol	
International Organization for Standardization See	
ISO	
IP addresses	
configuring.....	8
IP layer	
correcting	128
incorrectly configured.....	125, 128
verifying.....	123
IS-IS	128
OSPF.....	128
IP and IGP layers	117
broken network topology, figure	121
model, figure	120
problems.....	121
verifying.....	119
IP layer	
adjacencies, verifying	125
broken network topology, figure	122
checklist for verifying	117

-
- IP addresses
 - correcting.....128
 - incorrect128
 - verifying123
 - LSP up.....132
 - neighbors, verifying125
 - rename unit 0 family inet address
 - command.....129
 - show interfaces terse command.....123
 - show isis adjacency extensive command.....125
 - show mpls lsp extensive command.....129
 - show ospf neighbor extensive command.....125
 - verifying.....121
 - IS-IS protocol6
 - adjacencies
 - establishing.....10
 - verifying.....14, 147
 - broken network topology, figure144
 - configuration, verifying.....149
 - configuring.....14
 - interfaces, verifying147
 - IP addresses128
 - Level 111, 13
 - Level 2.....11, 13
 - Level 2 adjacencies.....9
 - metric.....12
 - network topology, figure9
 - on routers, enabling.....10
 - passive statement.....10
 - best practice.....12
 - show configuration protocols isis
 - command.....149
 - show isis adjacency command.....14, 147
 - show isis interface command147
 - show mpls lsp extensive command.....145, 151
 - verifying.....144
 - ISO10
 - address, configuring.....12
 - reception and transmission, enabling.....13
- K**
- keepalive
 - frames.....104, 106
 - multiplier50
- L**
- label-switched path See LSP
 - layered model
 - BGP layer, figure194
 - checklist85
 - data link layer, figure104
 - figure86
 - MPLS layer, figure175
 - physical layer, figure94
 - RSVP layer, figure156
 - Level 1, disabling.....10
 - level statement.....10
 - link-state database.....138
 - local address
 - configuring for BGP22
 - local-address statement.....22
 - log files, RSVP
 - configuring.....73
 - viewing74
 - loopback interface.....12, 13
 - configuring
 - for IS-IS.....10
 - for OSPF.....17, 139
 - NET address, configuring.....12
 - LSP6
 - configuration.....95
 - ingress router, verifying78
 - route, checking.....179
 - show mpls lsp command.....60
 - show mpls lsp extensive command.....61
 - show route table inet.3 command.....79
 - show route table mpls.0 command80
 - show rsvp session command.....65
 - state, checklist for determining.....59
 - statistics, determining.....65
 - transit router, verifying79
 - use, checklist for verifying.....77
 - verifying, general.....30, 122
- M**
- management backbone, establishing
 - adjacencies.....12
 - management interface
 - IS-IS, disabling.....10
 - MPLS, disabling26
 - OSPF, disabling.....17
 - RSVP, disabling.....26
 - manuals
 - comments on.....xviii
 - metric.....12
 - metric statement.....10

model, checklist for.....	85	OSPF protocol.....	6
MPLS layer		adjacencies.....	16
broken network topology, figure	175	verifying	19
checking	174	area, configuring	17, 137
MPLS protocol.....	3	backbone.....	16
activate interface command.....	187	broken network topology, figure	133
administrative groups.....	49	components of example network.....	16
configuration, incorrect.....	49	configuration, verifying.....	138
edit protocols mpls command.....	187	enabling.....	17
enabling	25	interfaces, verifying	136
fxp0.0 statement.....	26	IP addresses.....	128
interfaces, verifying.....	47	LSP, verifying.....	133
labels, verifying.....	183	neighbors, verifying.....	138
network topology, figure	7, 46	network topology, figure	16
on routers, enabling	26	passive statement.....	17
ping command.....	183	best practice.....	19
ping mpls rsvp lsp-name detail		RID, configuring.....	17
command.....	183	set traffic-engineering command.....	140
routing table	80, 179	show configuration protocols ospf	
show configuration command.....	32	command.....	138
show configuration interfaces command.....	185	show mpls lsp extensive command.....	133, 141
show configuration protocols mpls		show ospf interface command.....	136
command.....	185	show ospf neighbor command.....	19, 138
show mpls interface command.....	47	traffic engineering.....	16, 139
show mpls lsp extensive command.....	176, 188	verifying.....	132
show route command.....	181		
show route table mpls.0 command.....	179	P	
verifying	46	parentheses, in syntax descriptions.....	xvii
Multiprotocol Label Switching See MPLS protocol		passive statement	
		IS-IS	10, 12
N		OSPF	17, 19
neighbors		path messages, RSVP protocol.....	70
configuring BGP.....	21	peers	
IP layer verifying	125	configuring BGP	21
NET address.....	12	verifying BGP.....	25
network		physical layer	
configuring.....	6	broken network topology, figure	95
example	6	checklist for verifying.....	93
example configurations for routers.....	32	family mpls statement.....	99
MPLS, configuring.....	7	ping command.....	98
problems.....	92	set family mpls command.....	99
network entity title See NET		show configuration interfaces command.....	98
network topology		show interfaces terse command.....	98
LSP status, figure	60	show mpls lsp extensive command	96, 100
LSP use, figure	78	ping command	98, 183
		ping mpls rsvp lsp-name detail command	183
O		Point-to-Point Protocol See PPP	
Open Shortest Path First OSPF See OSPF protocol		policy, applying	23
		preemption	50

R

refresh timer49
 rename unit 0 family inet address command129
 Resource Reservation Protocol See RSVP protocol
 restart helper mode50
 restart time50
 Resv messages.....70
 RID16, 203
 OSPF, configuring17
 problems.....18
 route export policy23
 router ID See RID
 routing policy.....19
 applying.....23
 defining.....23
 routing table, MPLS.....80, 179
 RSVP layer
 broken network topology, figure157
 RSVP protocol6
 checklist for
 signal processing, verifying.....69
 configuration, verifying.....166
 deactivate traceoptions.....74
 edit protocols rsvp command.....167
 edit protocols rsvp traceoptions
 command.....73
 enabling.....25, 26
 fxp0.0 statement.....26
 graceful restart.....50
 interfaces, verifying50, 163
 keepalive multiplier.....50
 log files
 configuring73
 monitoring.....74
 viewing.....74
 neighbor state, displaying.....72
 neighbors, verifying.....162
 path messages70
 preemption.....50
 refresh timer.....49
 restart helper mode.....50
 restart time.....50
 RSVP layer, checking.....156
 sessions, verifying.....160
 set flag packets command.....74
 set interface command.....167
 show configuration protocols rsvp
 command.....166
 show mpls lsp extensive command.....158, 168

 show rsvp interface command.....50
 show rsvp interfaces command.....164
 show rsvp neighbor command72, 73, 162
 show rsvp session command.....160
 show rsvp statistics command70
 show rsvp version command.....49
 tracing operations, configuring.....74
 verifying.....46, 49
 run show log rsvp.log command.....74

S

sessions, checking BGP.....196
 set area command.....17
 set autonomous-system command21
 set export policy command.....23
 set family inet command8
 set family mpls command99
 set flag packets command74
 set group command.....22
 set interface command.....10, 17, 167
 set local-address command206
 set mpls interface command.....26
 set router-id command.....17
 set rsvp interface command.....26
 set traffic-engineering command.....17, 140
 show bgp summary command25, 196
 show configuration command32, 198
 show configuration interfaces command98, 185
 show configuration protocols isis command149
 show configuration protocols mpls
 command115
 show configuration protocols mpls command
 185
 show configuration protocols ospf command138
 show configuration protocols rsvp command166
 show interfaces command107
 show interfaces extensive command107
 show interfaces terse command52, 98, 123
 show isis adjacency command14, 147
 show isis adjacency extensive command125
 show isis interface command147
 show mpls interface command47
 show mpls lsp command60, 92
 show mpls lsp extensive command
 data link layer106, 112
 history log, displaying.....72
 IP layer129
 IS-IS protocol145, 151
 LSP, verifying.....30, 61, 122

MPLS layer	176
MPLS protocol	188
OSPF protocol	133, 141
physical layer.....	96, 100
RSVP protocol	158, 168
show mpls lsp name command.....	91
show mpls lsp name extensive command.....	91
show ospf interface command	136
show ospf neighbor command	19, 138
show ospf neighbor extensive command	125
show route command	181
show route detail command	204
show route receive protocol bgp command.....	205
show route table inet.3 command	79
show route table mpls.0 command	80, 179
show rsvp interface command	50
show rsvp interfaces command	164
show rsvp neighbor command	72, 73, 162
show rsvp session command	160
show rsvp session detail command	65
show rsvp statistics command	70
show rsvp version command	49
stack bit value.....	208
static route	
configuring.....	23
export policy	23
support, technical See technical support	
syntax conventions.....	xvi

T

technical support	
contacting JTAC.....	xviii
time-to-live.....	208
traceroute command	207
tracing operations, configuring	74
traffic engineering	
configuring.....	17
OSPF	16, 139
traffic, verifying BGP	207
transit router	
show route table mpls.0 command	80, 179

U

unit number.....	11
unit statement.....	12