



Junos[®] OS

RSVP Configuration Guide

Release
13.1



Published: 2013-02-28

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 RSVP Configuration Guide

13.1

Copyright © 2013, Juniper Networks, Inc.
All rights reserved.

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.

Table of Contents

About the Documentation	xi
Documentation and Release Notes	xi
Supported Platforms	xi
Using the Examples in This Manual	xi
Merging a Full Example	xii
Merging a Snippet	xii
Documentation Conventions	xiii
Documentation Feedback	xv
Requesting Technical Support	xv
Self-Help Online Tools and Resources	xv
Opening a Case with JTAC	xvi

Part 1

Chapter 1

Overview

Introduction to RSVP	3
RSVP Introduction	4
Junos OS RSVP Protocol Implementation	4
RSVP Operation Overview	5
RSVP Authentication	5
RSVP and IGP Hello Packets and Timers	5
RSVP Message Types	6
Path Messages	6
Resv Messages	7
PathTear Messages	7
ResvTear Messages	7
PathErr Messages	7
ResvErr Messages	7
ResvConfirm Messages	8
RSVP Reservation Styles	8
RSVP Refresh Reduction	9
MTU Signaling in RSVP	9
How the Correct MTU Is Signaled in RSVP	10
Determining an Outgoing MTU Value	11
MTU Signaling in RSVP Limitations	12
Link Protection	12
Fast Reroute, Node Protection, and Link Protection	13
Multiple Bypass LSPs	14
Node Protection	14
RSVP Graceful Restart	16
RSVP Graceful Restart Operation	16
Processing the Restart Cap Object	17

Part 2

Chapter 2

Configuration

RSVP Configuration Guidelines	21
Minimum RSVP Configuration	21
Configuring RSVP and MPLS	22
Example: Configuring RSVP and MPLS	23
Configuring RSVP Interfaces	23
Configuring RSVP Refresh Reduction	23
Determining the Refresh Reduction Capability of RSVP Neighbors	25
Configuring the RSVP Hello Interval	25
Configuring RSVP Authentication	26
Configuring the Bandwidth Subscription for Class Types	26
Configuring the RSVP Update Threshold on an Interface	27
Configuring RSVP for Unnumbered Interfaces	28
Configuring RSVP Node ID Hellos	28
Configuring Hello Acknowledgments for Nonsession RSVP Neighbors	29
Configuring Node Protection or Link Protection for LSPs	29
Switching LSPs Away from a Network Node	30
Configuring Inter-AS Node and Link Protection	31
Configuring Link Protection on Interfaces Used by LSPs	32
Configuring Bypass LSPs	33
Configuring the Next-Hop or Next-Next-Hop Node Address for Bypass LSPs	33
Configuring Administrative Groups for Bypass LSPs	34
Configuring the Bandwidth for Bypass LSPs	34
Configuring Class of Service for Bypass LSPs	35
Configuring the Hop Limit for Bypass LSPs	35
Configuring the Maximum Number of Bypass LSPs	36
Disabling CSPF for Bypass LSPs	36
Disabling Node Protection for Bypass LSPs	36
Configuring the Optimization Interval for Bypass LSPs	37
Configuring an Explicit Path for Bypass LSPs	37
Configuring the Amount of Bandwidth Subscribed for Bypass LSPs	38
Configuring Priority and Preemption for Bypass LSPs	38
Configuring RSVP Setup Protection	39
Configuring RSVP Graceful Restart	40
Enabling Graceful Restart for All Routing Protocols	40
Disabling Graceful Restart for RSVP	40
Disabling RSVP Helper Mode	41
Configuring the Maximum Helper Recovery Time	41
Configuring the Maximum Helper Restart Time	41
Configuring Load Balancing Across RSVP LSPs	41
Configuring RSVP Automatic Mesh	42
Configuring Timers for RSVP Refresh Messages	44
Preempting RSVP Sessions	45
Configuring MTU Signaling in RSVP	45
Enabling MTU Signaling in RSVP	46
Enabling Packet Fragmentation	46
Configuring Ultimate-Hop Popping for LSPs	46

	Configuring RSVP to Pop the Label on the Ultimate-Hop Router	50
	Enabling Ultimate-Hop Popping on Point-to-Multipoint LSPs	50
	Tracing RSVP Protocol Traffic	52
	Examples: Tracing RSVP Protocol Traffic	52
Part 3	Administration	
Chapter 3	RSVP Standards and Terminology	57
	Supported RSVP Standards	57
	RSVP Graceful Restart Standard	58
	RSVP Graceful Restart Terminology	58
Chapter 4	Summary of RSVP Configuration Statements	61
	admin-group	61
	aggregate (Protocols RSVP)	62
	authentication-key (Protocols RSVP)	63
	bandwidth (Protocols RSVP)	64
	bypass (Signaled LSP)	65
	bypass (Static LSP)	66
	class-of-service (Protocols RSVP)	67
	destination-networks	68
	devices	69
	disable (Protocols RSVP)	70
	fast-reroute (Protocols RSVP)	71
	graceful-deletion-timeout	71
	graceful-restart (Protocols RSVP)	72
	hello-acknowledgements	73
	hello-interval (Protocols RSVP)	73
	hop-limit	74
	interface (Protocols RSVP)	75
	keep-multiplier	76
	label-switched-path-template	77
	link-protection (RSVP)	78
	load-balance (Protocols RSVP)	79
	max-bypasses	80
	no-local-reversion	81
	node-hello	82
	no-adjacency-down-notification (Protocols IS-IS)	83
	no-cspf (Protocols RSVP)	84
	no-interface-hello	84
	no-neighbor-down-notification	85
	no-node-id-subobject	85
	no-p2mp-sublsp	86
	node-link-protection (Protocols MPLS)	86
	optimize-timer (Protocols RSVP)	87
	path (Protocols RSVP)	88
	peer-interface (Protocols RSVP)	89
	preemption	90
	priority (Protocols RSVP)	91
	refresh-time	92

reliable	92
rsvp	93
rsvp-te (Routing Options)	94
setup-protection	94
soft-preemption (Protocols RSVP)	95
static-label-switched-path	96
subscription	97
traceoptions (Protocols RSVP)	98
transit	100
tunnel-services (RSVP)	101
ultimate-hop-popping	102
update-threshold	103

Part 4

Index

Index	107
-------------	-----

List of Figures

Part 1	Overview	
Chapter 1	Introduction to RSVP	3
	Figure 1: Link Protection Creating a Bypass LSP for the Protected Interface	13
	Figure 2: Node Protection Creating a Next-Next-Hop Bypass LSP	15
Part 2	Configuration	
Chapter 2	RSVP Configuration Guidelines	21
	Figure 3: Penultimate-Hop Popping for an LSP	47
	Figure 4: Ultimate-Hop Popping for an LSP	47

List of Tables

	About the Documentation	xi
	Table 1: Notice Icons	xiii
	Table 2: Text and Syntax Conventions	xiii
Part 1	Overview	
Chapter 1	Introduction to RSVP	3
	Table 3: One-to-One Backup Compared with Facility Backup	13

About the Documentation

- Documentation and Release Notes on page xi
- Supported Platforms on page xi
- Using the Examples in This Manual on page xi
- Documentation Conventions on page xiii
- Documentation Feedback on page xv
- Requesting Technical Support on page xv

Documentation and Release Notes

To obtain the most current version of all Juniper Networks[®] technical documentation, see the product documentation page on the Juniper Networks website at <http://www.juniper.net/techpubs/>.

If the information in the latest release notes differs from the information in the documentation, follow the product Release Notes.

Juniper Networks Books publishes books by Juniper Networks engineers and subject matter experts. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration. The current list can be viewed at <http://www.juniper.net/books>.

Supported Platforms

For the features described in this document, the following platforms are supported:

- [ACX Series](#)
- [MX Series](#)
- [T Series](#)
- [M Series](#)
- [PTX Series](#)

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.xsl;
    }
  }
}
interfaces {
  fxp0 {
    disable;
    unit 0 {
      family inet {
        address 10.0.0.1/24;
      }
    }
  }
}
```

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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

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

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

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

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

Documentation Conventions

Table 1 on page xiii 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 xiii defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

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

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

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

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. 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 post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at <http://www.juniper.net/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/>
- Search for known bugs: <http://www2.juniper.net/kb/>
- Find product documentation: <http://www.juniper.net/techpubs/>
- Find solutions and answer questions using our Knowledge Base: <http://kb.juniper.net/>
- Download the latest versions of software and review release notes: <http://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <https://www.juniper.net/alerts/>

- Join and participate in the Juniper Networks Community Forum:
<http://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <http://www.juniper.net/cm/>

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

Opening a Case with JTAC

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

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

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

PART 1

Overview

- [Introduction to RSVP on page 3](#)

CHAPTER 1

Introduction to RSVP

- [RSVP Introduction on page 4](#)
- [Junos OS RSVP Protocol Implementation on page 4](#)
- [RSVP Operation Overview on page 5](#)
- [RSVP Authentication on page 5](#)
- [RSVP and IGP Hello Packets and Timers on page 5](#)
- [RSVP Message Types on page 6](#)
- [Path Messages on page 6](#)
- [Resv Messages on page 7](#)
- [PathTear Messages on page 7](#)
- [ResvTear Messages on page 7](#)
- [PathErr Messages on page 7](#)
- [ResvErr Messages on page 7](#)
- [ResvConfirm Messages on page 8](#)
- [RSVP Reservation Styles on page 8](#)
- [RSVP Refresh Reduction on page 9](#)
- [MTU Signaling in RSVP on page 9](#)
- [How the Correct MTU Is Signaled in RSVP on page 10](#)
- [Determining an Outgoing MTU Value on page 11](#)
- [MTU Signaling in RSVP Limitations on page 12](#)
- [Link Protection on page 12](#)
- [Fast Reroute, Node Protection, and Link Protection on page 13](#)
- [Multiple Bypass LSPs on page 14](#)
- [Node Protection on page 14](#)
- [RSVP Graceful Restart on page 16](#)
- [RSVP Graceful Restart Operation on page 16](#)
- [Processing the Restart Cap Object on page 17](#)

RSVP Introduction

RSVP is a resource reservation setup protocol that is used by both network hosts and routers. Hosts use RSVP to request a specific class of service (CoS) from the network for particular application flows. Routers use RSVP to deliver CoS requests to all routers along the data path. RSVP also can maintain and refresh states for a requested CoS application flow.

RSVP treats an application flow as a simplex connection. That is, the CoS request travels only in one direction—from the sender to the receiver. RSVP is a transport layer protocol that uses IP as its network layer. However, RSVP does not transport application flows. Rather, it is more of an Internet control protocol, similar to the Internet Control Message Protocol (ICMP) and Internet Group Management Protocol (IGMP). RSVP runs as a separate software process in the Junos OS and is not in the packet forwarding path.

RSVP is not a routing protocol, but rather is designed to operate with current and future unicast and multicast routing protocols. The routing protocols are responsible for choosing the routes to use to forward packets, and RSVP consults local routing tables to obtain routes. RSVP only ensures the CoS of packets traveling along a data path.

The receiver in an application flow requests the preferred CoS from the sender. To do this, the receiver issues an RSVP CoS request on behalf of the local application. The request propagates to all routers in reverse direction of the data paths toward the sender. In this process, RSVP requests might be merged, resulting in a protocol that scales well when there are a large number of receivers.

Because the number of receivers in an application flow is likely to change and the flow of delivery paths might change during the life of an application flow, RSVP takes a soft-state approach in its design, creating and removing the protocol states in routers and hosts incrementally over time. RSVP sends periodic refresh messages to maintain its state and to recover from occasional lost messages. In the absence of refresh messages, RSVP states automatically time out and are deleted.

Junos OS RSVP Protocol Implementation

The Junos implementation of RSVP supports RSVP version 1. The software includes support for all mandatory objects and RSVP message types, and supports message integrity and node authentications through the Integrity object.

The primary purpose of the Junos RSVP software is to support dynamic signaling within MPLS label-switched paths (LSPs). Supporting resource reservations over the Internet is only a secondary purpose of the Junos OS implementation. Since supporting resource reservations is secondary, the Junos RSVP software does not support the following features:

- IP multicasting sessions.
- Traffic control. The software cannot make resource reservations for real-time video or audio sessions.

With regard to the protocol mechanism, packet processing, and RSVP objects supported, the Junos OS implementation of the software is interoperable with other RSVP implementations.

RSVP Operation Overview

RSVP creates independent sessions to handle each data flow. A session is identified by a combination of the destination address, an optional destination port, and a protocol. Within a session, there can be one or more senders. Each sender is identified by a combination of its source address and source port. An out-of-band mechanism, such as a session announcement protocol or human communication, is used to communicate the session identifier to all senders and receivers.

A typical RSVP session involves the following sequence of events:

1. A potential sender starts sending RSVP path messages to the session address.
2. A receiver, wanting to join the session, registers itself if necessary. For example, a receiver in a multicast application would register itself with IGMP.
3. The receiver receives the path messages.
4. The receiver sends appropriate Resv messages toward the sender. These messages carry a flow descriptor, which is used by routers along the path to make reservations in their link-layer media.
5. The sender receives the Resv message and then starts sending application data.

This sequence of events is not necessarily strictly synchronized. For example, receivers can register themselves before receiving path messages from the sender, and application data can flow before the sender receives Resv messages. Application data that is delivered before the actual reservation contained in the Resv message typically is treated as best-effort, non-real-time traffic with no CoS guarantee.

RSVP Authentication

The Junos OS supports both the RSVP authentication style described in RFC 2747 (allowing for multivendor compatibility) and the RSVP authentication style described in Internet draft draft-ietf-rsvp-md5-03.txt. The Junos OS uses the authentication style described in Internet draft draft-ietf-rsvp-md5-08.txt by default. If the router receives an RFC 2747-style RSVP authentication from a neighbor, it switches to this style of authentication for that neighbor. The RSVP authentication style for each neighboring router is determined separately.

RSVP and IGP Hello Packets and Timers

RSVP monitors the status of the interior gateway protocol (IGP) (IS-IS or OSPF) neighbors and relies on the IGP protocols to detect when a node fails. If an IGP protocol declares a neighbor down (because hello packets are no longer being received), RSVP also brings down that neighbor. However, the IGP protocols and RSVP still act independently when bringing a neighbor up.

In the Junos OS, RSVP typically relies on IGP hello packet detection to check for node failures. RSVP sessions are kept up even if RSVP hello packets are no longer being received, so long as the router continues to receive IGP hello packets. RSVP sessions are maintained until either the router stops receiving IGP hello packets or the RSVP Path and Resv messages time out. Configuring a short time for the IS-IS or OSPF hello timers allows these protocols to detect node failures quickly.

RSVP hellos can be relied on when the IGP does not recognize a particular neighbor (for example, if IGP is not enabled on the interface) or if the IGP is RIP (not IS-IS or OSPF). Also, the equipment of other vendors might be configured to monitor RSVP sessions based on RSVP hello packets. This equipment might also take an RSVP session down due to a loss of RSVP hello packets.

We do not recommend configuring a short RSVP hello timer. If quick discovery of a failed neighbor is needed, configure short IGP (OSPF or IS-IS) hello timers.

OSPF and IS-IS have infrastructure to manage rapid hello message sending and receiving reliably, even if the routing protocols or some other process are straining the processing capability of the router. Under the same circumstances, RSVP hellos might time out prematurely even though the neighbor is functioning normally.

RSVP Message Types

RSVP uses the following types of messages to establish and remove paths for data flows, establish and remove reservation information, confirm the establishment of reservations, and report errors:

- [Path Messages on page 6](#)
- [Resv Messages on page 7](#)
- [PathTear Messages on page 7](#)
- [ResvTear Messages on page 7](#)
- [PathErr Messages on page 7](#)
- [ResvErr Messages on page 7](#)
- [ResvConfirm Messages on page 8](#)

Path Messages

Each sender host transmits path messages downstream along the routes provided by the unicast and multicast routing protocols. Path messages follow the exact paths of application data, creating path states in the routers along the way, thus enabling routers to learn the previous-hop and next-hop node for the session. Path messages are sent periodically to refresh path states.

The refresh interval is controlled by a variable called the *refresh-time*, which is the periodical refresh timer expressed in seconds. A path state times out if a router does not receive a specified number of consecutive path messages. This number is specified by

a variable called *keep-multiplier*. Path states are kept for $(\text{keep-multiplier} + 0.5) \times 1.5 \times \text{refresh-time}$ seconds.

Resv Messages

Each receiver host sends reservation request (Resv) messages upstream toward senders and sender applications. Resv messages must follow exactly the reverse path of path messages. Resv messages create and maintain a reservation state in each router along the way.

Resv messages are sent periodically to refresh reservation states. The refresh interval is controlled by the same refresh time variable, and reservation states are kept for $(\text{keep-multiplier} + 0.5) \times 1.5 \times \text{refresh-time}$ seconds.

PathTear Messages

PathTear messages remove (tear down) path states as well as dependent reservation states in any routers along a path. PathTear messages follow the same path as path messages. A PathTear typically is initiated by a sender application or by a router when its path state times out.

PathTear messages are not required, but they enhance network performance because they release network resources quickly. If PathTear messages are lost or not generated, path states eventually time out when they are not refreshed, and the resources associated with the path are released.

ResvTear Messages

ResvTear messages remove reservation states along a path. These messages travel upstream toward senders of the session. In a sense, ResvTear messages are the reverse of Resv messages. ResvTear messages typically are initiated by a receiver application or by a router when its reservation state times out.

ResvTear messages are not required, but they enhance network performance because they release network resources quickly. If ResvTear messages are lost or not generated, reservation states eventually time out when they are not refreshed, and the resources associated with the reservation are released.

PathErr Messages

When path errors occur (usually because of parameter problems in a path message), the router sends a unicast PathErr message to the sender that issued the path message. PathErr messages are advisory; these messages do not alter any path state along the way.

ResvErr Messages

When a reservation request fails, a ResvErr error message is delivered to all the receivers involved. ResvErr messages are advisory; these messages do not alter any reservation state along the way.

ResvConfirm Messages

Receivers can request confirmation of a reservation request, and this confirmation is sent with a ResvConfirm message. Because of the complex RSVP flow-merging rules, a confirmation message does not necessarily provide end-to-end confirmation of the entire path. Therefore, ResvConfirm messages are an indication, not a guarantee, of potential success.

Juniper Networks routers never request confirmation using the ResvConfirm message; however, a Juniper Networks router can send a ResvConfirm message if it receives a request from another vendor's equipment.

RSVP Reservation Styles

A reservation request includes options for specifying the reservation style. The reservation styles define how reservations for different senders within the same session are treated and how senders are selected.

Two options specify how reservations for different senders within the same session are treated:

- Distinct reservation—Each receiver establishes its own reservation with each upstream sender.
- Shared reservation—All receivers make a single reservation that is shared among many senders.

Two options specify how senders are selected:

- Explicit sender—List all selected senders.
- Wildcard sender—Select all senders, which then participate in the session.

The following reservation styles, formed by a combination of these four options, currently are defined:

- Fixed filter (FF)—This reservation style consists of distinct reservations among explicit senders. Examples of applications that use fixed-filter-style reservations are video applications and unicast applications, which both require flows that have a separate reservation for each sender. The fixed filter reservation style is enabled on RSVP LSPs by default.
- Wildcard filter (WF)—This reservation style consists of shared reservations among wildcard senders. This type of reservation reserves bandwidth for any and all senders, and propagates upstream toward all senders, automatically extending to new senders as they appear. A sample application for wildcard filter reservations is an audio application in which each sender transmits a distinct data stream. Typically, only a few senders are transmitting at any one time. Such a flow does not require a separate reservation for each sender; a single reservation is sufficient.
- Shared explicit (SE)—This reservation style consists of shared reservations among explicit senders. This type of reservation reserves bandwidth for a limited group of

senders. A sample application is an audio application similar to that described for wildcard filter reservations.

RSVP Refresh Reduction

RSVP relies on soft-state to maintain the path and reservation states on each router. If the corresponding refresh messages are not sent periodically, the states eventually time out and reservations are deleted. RSVP also sends its control messages as IP datagrams with no reliability guarantee. It relies on periodic refresh messages to handle the occasional loss of Path or Resv messages.

The RSVP refresh reduction extensions, based on RFC 2961, addresses the following problems that result from relying on periodic refresh messages to handle message loss:

- Scalability—The scaling problem arises from the periodic transmission and processing overhead of refresh messages, which increases as the number of RSVP sessions increases.
- Reliability and latency—The reliability and latency problem stems from the loss of nonrefresh RSVP messages or one-time RSVP messages such as PathTear or PathErr. The time to recover from such a loss is usually tied to refresh interval and the keepalive timer.

The RSVP refresh reduction capability is advertised by enabling the refresh reduction (RR) capable bit in the RSVP common header. This bit is only significant between RSVP neighbors.

RSVP refresh reduction includes the following features:

- RSVP message bundling using the bundle message
- RSVP Message ID to reduce message processing overhead
- Reliable delivery of RSVP messages using Message ID, Message Ack, and Message Nack
- Summary refresh to reduce the amount of information transmitted every refresh interval

The RSVP refresh reduction specification (RFC 2961) allows you to enable some or all of the above capabilities on a router. It also describes various procedures that a router can use to detect the refresh reduction capabilities of its neighbor.

The Junos OS supports all of the refresh reduction extensions, some of which can be selectively enabled or disabled. The Junos OS supports Message ID and therefore can perform reliable message delivery only for Path and Resv messages.

For information about how to configure RSVP refresh reduction, see [“Configuring RSVP Refresh Reduction” on page 23](#).

MTU Signaling in RSVP

The maximum transmission unit (MTU) is the largest size packet or frame, in bytes, that can be sent in a network. An MTU that is too large might cause retransmissions. Too

small an MTU might cause the router to send and handle relatively more header overhead and acknowledgments. There are default values for MTUs associated with various protocols. You can also explicitly configure an MTU on an interface.

When an LSP is created across a set of links with different MTU sizes, the ingress router does not know what the smallest MTU is on the LSP path. By default, the MTU for an LSP is 1,500 bytes.

If this MTU is larger than the MTU of one of the intermediate links, traffic might be dropped, because MPLS packets cannot be fragmented. Also, the ingress router is not aware of this type of traffic loss, because the control plane for the LSP would still function normally.

To prevent this type of packet loss in MPLS LSPs, you can configure MTU signaling in RSVP. This feature is described in RFC 3209. Juniper Networks supports the Integrated Services object for MTU signaling in RSVP. The Integrated Services object is described in RFCs 2210 and 2215. MTU signaling in RSVP is disabled by default.

To avoid packet loss due to MTU mismatches, the ingress router needs to do the following:

- Signal the MTU on the RSVP LSP—To prevent packet loss from an MTU mismatch, the ingress router needs to know what the smallest MTU value is along the path taken by the LSP. Once this MTU value is obtained, the ingress router can assign it to the LSP.
- Fragment packets—Using the assigned MTU value, packets that exceed the size of the MTU can be fragmented into smaller packets on the ingress router before they are encapsulated in MPLS and sent over the RSVP-signaled LSP.

Once both MTU signaling and packet fragmentation have been enabled on an ingress router, any route resolving to an RSVP LSP on this router uses the signaled MTU value. For information about how to configure this feature, see [“Configuring MTU Signaling in RSVP” on page 45](#).

The following sections describe how MTU signaling in RSVP works:

- [How the Correct MTU Is Signaled in RSVP on page 10](#)
- [Determining an Outgoing MTU Value on page 11](#)
- [MTU Signaling in RSVP Limitations on page 12](#)

How the Correct MTU Is Signaled in RSVP

How the correct MTU is signaled in RSVP varies depending on whether the network devices (for example, routers) explicitly support MTU signaling in RSVP or not.

If the network devices support MTU signaling in RSVP, the following occur when you enable MTU signaling:

- The MTU is signaled from the ingress router to the egress router by means of the Adspec object. Before forwarding this object, the ingress router enters the MTU value associated with the interface over which the path message is sent. At each hop in the path, the MTU value in the Adspec object is updated to the minimum of the received value and the value of the outgoing interface.
- The ingress router uses the traffic specification (Tspec) object to specify the parameters for the traffic it is going to send. The MTU value signaled for the Tspec object at the ingress router is the maximum MTU value (9192 bytes for M Series and T Series routers, 9500 bytes for PTX Series Packet Transport Switches). This value does not change en route to the egress router.
- When the Adspec object arrives at the egress router, the MTU value is correct for the path (meaning it is the smallest MTU value discovered). The egress router compares the MTU value in the Adspec object to the MTU value in the Tspec object. It signals the smaller MTU using the Flowspec object in the Resv message.
- When the Resv object arrives at the ingress router, the MTU value in this object is used as the MTU for the next hops that use the LSP.

In a network where there are devices that do not support MTU signaling in RSVP, you might have the following behaviors:

- If the egress router does not support MTU signaling in RSVP, the MTU is set to 1,500 bytes by default.
- A Juniper Networks transit router that does not support MTU signaling in RSVP sets an MTU value of 1,500 bytes in the Adspec object by default.

**Related
Documentation**

- [Configuring the Media MTU](#)

Determining an Outgoing MTU Value

The outgoing MTU value is the smaller of the values received in the Adspec object compared to the MTU value of the outgoing interface. The MTU value of the outgoing interface is determined as follows:

- If you configure an MTU value under the **[family mpls]** hierarchy level, this value is signaled.
- If you do not configure an MTU, the **inet** MTU is signaled.

MTU Signaling in RSVP Limitations

The following are limitations to MTU signaling in RSVP:

- Changes in the MTU value might cause a temporary loss of traffic in the following situations:
 - For link protection and node protection, the MTU of the bypass is only signaled at the time the bypass becomes active. During the time it takes for the new path MTU to be propagated, packet loss might occur because of an MTU mismatch.
 - For fast reroute, the MTU of the path is updated only after the detour becomes active, causing a delay in an update to the MTU at the ingress router. Until the MTU is updated, packet loss might occur if there is an MTU mismatch.

In both cases, only packets that are larger than the detour or bypass MTU are lost.


- When an MTU is updated, it triggers a change in the next hop. Any change in the next hop causes the route statistics to be lost.
- The minimum MTU supported for MTU signaling in RSVP is 1,488 bytes. This value prevents a false or incorrectly configured value from being used.
- For single-hop LSPs, the MTU value displayed by the **show** commands is the RSVP-signaled value. However, this MPLS value is ignored and the correct IP value is used.

Link Protection

Link protection helps to ensure that traffic going over a specific interface to a neighboring router can continue to reach this router if that interface fails. When link protection is configured for an interface and an LSP that traverses this interface, a bypass LSP is created that will handle this traffic if the interface fails. The bypass LSP uses a different interface and path to reach the same destination. The path used can be configured explicitly, or you can rely on CSPF. The RSVP metric for the bypass LSP is set in the range of 20,000 through 29,999 (this value is not user configurable).

If a link-protected interface fails, traffic is quickly switched to the bypass LSP. Note that a bypass LSP cannot share the same egress interface with the LSPs it monitors.

In [Figure 1 on page 13](#), link protection is enabled on Interface B between Router 1 and Router 2. It is also enabled on LSP A, an LSP that traverses the link between Router 1 and Router 2. If the link between Router 1 and Router 2 fails, traffic from LSP A is quickly switched to the bypass LSP generated by link protection.

Figure 1: Link Protection Creating a Bypass LSP for the Protected Interface

Although LSPs traversing an interface can be configured to take advantage of link protection, it is important to note that it is specifically the interface that benefits from link protection. If link protection is enabled on an interface but not on a particular LSP traversing that interface, then if the interface fails, that LSP will also fail.



NOTE: Link protection does not work on unnumbered interfaces.

To protect traffic over the entire route taken by an LSP, you should configure fast reroute. For more information, see [Configuring Fast Reroute](#).

The following sections provide more information on link protection:

- [Fast Reroute, Node Protection, and Link Protection on page 13](#)
- [Multiple Bypass LSPs on page 14](#)

Fast Reroute, Node Protection, and Link Protection

RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*, describes two different types of traffic protection for RSVP-signaled LSPs:

- **One-to-one backup**—In the Junos OS this type of traffic protection is provided by fast reroute. Each LSP requires a protecting LSP to be signaled at each hop except the egress router. This protecting LSP cannot be shared.
- **Facility backup**—This is sometimes called many-to-one backup. In the Junos OS this type of traffic protection is provided by node and link protection. Each LSP requires a protecting LSP to be signaled at each hop except the egress router. Unlike fast reroute, this protecting LSP can be shared by other LSPs.

[Table 3 on page 13](#) summarizes the traffic protection types.

Table 3: One-to-One Backup Compared with Facility Backup

Comparison	One-to-One Backup	Facility Backup
Name of the protecting LSP	Detour LSP	Bypass LSP
Sharing of the protecting LSP	Cannot be shared	Can be shared by multiple LSPs

Table 3: One-to-One Backup Compared with Facility Backup (*continued*)

Comparison	One-to-One Backup	Facility Backup
Junos configuration statements	fast-reroute	node-link-protection and link-protection

Multiple Bypass LSPs

By default, link protection relies on a single bypass LSP to provide path protection for an interface. However, you can also specify multiple bypass LSPs to provide link protection for an interface. You can individually configure each of these bypass LSPs or create a single configuration for all of the bypass LSPs. If you do not configure the bypass LSPs individually, they all share the same path and bandwidth constraints.

The following algorithm describes how and when an additional bypass LSP is activated for an LSP:

1. If any currently active bypass can satisfy the requirements of the LSP (bandwidth, link protection, or node-link protection), the traffic is directed to that bypass.
2. If no active bypass LSP is available, scan through the manual bypass LSPs in first-in, first-out (FIFO) order, skipping those that are already active (each manual bypass can only be activated once). The first inactive manual bypass that can satisfy the requirements is activated and traffic is directed to that bypass.
3. If no manual bypass LSPs are available and if the **max-bypasses** statement activates multiple bypass LSPs for link protection, determine whether an automatically configured bypass LSP can satisfy the requirements. If an automatically configured bypass LSP is available and if the total number of active automatically configured bypass LSPs does not exceed the maximum bypass LSP limit (configured with the **max-bypasses** statement), activate another bypass LSP.

For information about how to configure multiple bypass LSPs for link protection, see [“Configuring Bypass LSPs” on page 33](#).

Node Protection

Node protection extends the capabilities of link protection. Link protection helps to ensure that traffic going over a specific interface to a neighboring router can continue to reach this router if that interface fails. Node protection ensures that traffic from an LSP traversing a neighboring router can continue to reach its destination even if the neighboring router fails.


When you enable node protection for an LSP, you must also enable link protection. Once enabled, node protection and link protection establish the following types of bypass LSPs:

- Next-hop bypass LSP—Provides an alternate route for an LSP to reach a neighboring router. This type of bypass LSP is established when you enable either node protection or link protection.

- Next-next-hop bypass LSP—Provides an alternate route for an LSP to get around a neighboring router en route to the destination router. This type of bypass LSP is established exclusively when node protection is configured. If a next-next-hop bypass LSP cannot be created, an attempt is made to signal a next-hop bypass LSP.

In [Figure 2 on page 15](#), node protection is enabled on Interface B on Router 1. Node protection is also enabled on LSP A, an LSP that traverses the link transiting Router 1, Router 2, and Router 3. If Router 2 suffers a hardware or software failure, traffic from LSP A is switched to the next-next-hop bypass LSP generated by node protection.

Figure 2: Node Protection Creating a Next-Next-Hop Bypass LSP



The time needed by node protection to switch traffic to a next-next-hop bypass LSP can be significantly longer than the time needed by link protection to switch traffic to a next-hop bypass LSP. Link protection relies on a hardware mechanism to detect a link failure, allowing it to quickly switch traffic to a next-hop bypass LSP.

Node failures are often due to software problems on the node router. Node protection relies on the receipt of hello messages from a neighboring router to determine whether it is still functioning. The time it takes node protection to divert traffic partly depends on how often the node router sends hello messages and how long it takes the node-protected router to react to having not received a hello message. However, once the failure is detected, traffic can be quickly diverted to the next-next-hop bypass LSP.



NOTE:

Node protection provides traffic protection in the event of an error or interruption of the physical link between two routers. It does not provide protection in the event of control plane errors. The following provides an example of a control plane error:

- A transit router changes the label of a packet due to a control plane error.
- When the ingress router receives the packet, it considers the label change to be a catastrophic event and deletes both the primary LSP and the associated bypass LSP.

Related Documentation

- [Configuring Node Protection or Link Protection for LSPs on page 29](#)

RSVP Graceful Restart

RSVP graceful restart allows a router undergoing a restart to inform its adjacent neighbors of its condition. The restarting router requests a grace period from the neighbor or peer, which can then cooperate with the restarting router. The restarting router can still forward MPLS traffic during the restart period; convergence in the network is not disrupted. The restart is not visible to the rest of the network, and the restarting router is not removed from the network topology. RSVP graceful restart can be enabled on both transit routers and ingress routers. It is available for both point-to-point LSPs and point-to-multipoint LSPs.

RSVP graceful restart is described in the following sections:

- [RSVP Graceful Restart Standard on page 58](#)
- [RSVP Graceful Restart Terminology on page 58](#)
- [RSVP Graceful Restart Operation on page 16](#)
- [Processing the Restart Cap Object on page 17](#)

RSVP Graceful Restart Operation

For RSVP graceful restart to function, the feature must be enabled on the global routing instance. RSVP graceful restart can be disabled at the protocol level (for RSVP alone) or at the global level for all protocols.

RSVP graceful restart requires the following of a restarting router and the router's neighbors:

- For the restarting router, RSVP graceful restart attempts to maintain the routes installed by RSVP and the allocated labels, so that traffic continues to be forwarded without disruption. RSVP graceful restart is done quickly enough to reduce or eliminate the impact on neighboring nodes.
- The neighboring routers must have RSVP graceful restart helper mode enabled, thus allowing them to assist a router attempting to restart RSVP.

An object called Restart Cap that is sent in RSVP hello messages advertises a node's restart capability. The neighboring node sends a Recover Label object to the restarting node to recover its forwarding state. This object is essentially the old label that the restarting node advertised before the node went down.

The following lists the RSVP graceful restart behaviors, which vary depending on the configuration and on which features are enabled:

- If you disable helper mode, the Junos OS does not attempt to help a neighbor restart RSVP. Any information that arrives with a Restart Cap object from a neighbor is ignored.
- When you enable graceful restart under the routing instance configuration, the router can restart gracefully with the help of its neighbors. RSVP advertises a Restart Cap

object (RSVP RESTART) in hello messages in which restart and recovery times are specified (neither value is 0).

- If you explicitly disable RSVP graceful restart under the **[protocols rsdp]** hierarchy level, the Restart Cap object is advertised with restart and recovery times specified as 0. The restart of neighboring routers is supported (unless helper mode is disabled), but the router itself does not preserve the RSVP forwarding state and cannot recover its control state.
- If after a restart RSVP realizes that no forwarding state has been preserved, the Restart Cap object is advertised with restart and recovery times specified as 0.
- If graceful restart and helper mode are disabled, RSVP graceful restart is completely disabled. The router neither recognizes nor advertises the RSVP graceful restart objects.

You cannot explicitly configure values for the restart and recovery times.

Unlike other protocols, there is no way for RSVP to determine that it has completed a restart procedure, other than a fixed timeout. All RSVP graceful restart procedures are timer-based. A **show rsdp version** command might indicate that the restart is still in progress even if all RSVP sessions are back up and the routes are restored.

Processing the Restart Cap Object

The following assumptions are made about a neighbor based on the Restart Cap object (assuming that a control channel failure can be distinguished unambiguously from a node restart):

- A neighbor that does not advertise the Restart Cap object in its hello messages cannot assist a router with state or label recovery, nor can it perform an RSVP graceful restart.
- After a restart, a neighbor advertising a Restart Cap object with a restart time equal to any value and a recovery time equal to 0 has not preserved its forwarding state. When a recovery time equals 0, the neighbor is considered dead and any states related to this neighbor are purged, regardless of the value of the restart time.
- After a restart, a neighbor advertising its recovery time with a value other than 0 can keep or has kept the forwarding state. If the local router is helping its neighbor with restart or recovery procedures, it sends a Recover Label object to this neighbor.

PART 2

Configuration

- [RSVP Configuration Guidelines on page 21](#)

CHAPTER 2

RSVP Configuration Guidelines

- [Minimum RSVP Configuration on page 21](#)
- [Configuring RSVP and MPLS on page 22](#)
- [Configuring RSVP Interfaces on page 23](#)
- [Configuring RSVP Node ID Hellos on page 28](#)
- [Configuring Hello Acknowledgments for Nonsession RSVP Neighbors on page 29](#)
- [Configuring Node Protection or Link Protection for LSPs on page 29](#)
- [Switching LSPs Away from a Network Node on page 30](#)
- [Configuring Inter-AS Node and Link Protection on page 31](#)
- [Configuring Link Protection on Interfaces Used by LSPs on page 32](#)
- [Configuring RSVP Setup Protection on page 39](#)
- [Configuring RSVP Graceful Restart on page 40](#)
- [Configuring Load Balancing Across RSVP LSPs on page 41](#)
- [Configuring RSVP Automatic Mesh on page 42](#)
- [Configuring Timers for RSVP Refresh Messages on page 44](#)
- [Preempting RSVP Sessions on page 45](#)
- [Configuring MTU Signaling in RSVP on page 45](#)
- [Configuring Ultimate-Hop Popping for LSPs on page 46](#)
- [Configuring RSVP to Pop the Label on the Ultimate-Hop Router on page 50](#)
- [Enabling Ultimate-Hop Popping on Point-to-Multipoint LSPs on page 50](#)
- [Tracing RSVP Protocol Traffic on page 52](#)

Minimum RSVP Configuration

To enable RSVP on a single interface, include the **rsvp** statement and specify the interface using the **interface** statement. This is the minimum RSVP configuration. All other RSVP configuration statements are optional.

```
rsvp {  
  interface interface-name;  
}
```

You can include these statements at the following hierarchy levels:

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

To enable RSVP on all interfaces, substitute **all** for the *interface-name* variable.

If you have configured interface properties on a group of interfaces and want to disable RSVP on one of the interfaces, include the **disable** statement:

```
interface interface-name {  
  disable;  
}
```

You can include this statement at the following hierarchy levels:

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

Configuring RSVP and MPLS

The primary purpose of the Junos RSVP software is to support dynamic signaling within label-switched paths (LSPs). When you enable both MPLS and RSVP on a router, MPLS becomes a client of RSVP. No additional configuration is required to bind MPLS and RSVP.

You can configure MPLS to set up signaled paths by using the **label-switched-path** statement at the **[edit protocols mpls]** hierarchy level. Each LSP translates into a request for RSVP to initiate an RSVP session. This request is passed through the internal interface between label switching and RSVP. After examining the request information, checking RSVP states, and checking the local routing tables, RSVP initiates one session for each LSP. The session is sourced from the local router and is destined for the target of the LSP.

When an RSVP session is successfully created, the LSP is set up along the paths created by the RSVP session. If the RSVP session is unsuccessful, RSVP notifies MPLS of its status. It is up to MPLS to initiate backup paths or continue retrying the initial path.

To pass label-switching signaling information, RSVP supports four additional objects: Label Request object, Label object, Explicit Route object, and Record Route object. For an LSP to be set up successfully, all routers along the path must support MPLS, RSVP, and the four objects. Of the four objects, the Record Route object is not mandatory.

To configure MPLS and make it a client of RSVP, do the following:

- Enable MPLS on all routers that will participate in the label switching (this is, on all routers that might be part of a label-switching path).
- Enable RSVP on all routers and on all router interfaces that form the LSP.
- Configure the routers at the beginning of the LSP.

Example: Configuring RSVP and MPLS

The following shows a sample configuration for a router at the beginning of an LSP:

```
[edit]
protocols {
  mpls {
    label-switched-path sf-to-london {
      to 192.168.1.4;
    }
  }
  rsvp {
    interface so-0/0/0;
  }
}
```

The following shows a sample configuration for all the other routers that form the LSP:

```
[edit]
protocols {
  mpls {
    interface so-0/0/0;
  }
  rsvp {
    interface so-0/0/0;
  }
}
```

Configuring RSVP Interfaces

The following sections describe how to configure RSVP interfaces:

- [Configuring RSVP Refresh Reduction on page 23](#)
- [Configuring the RSVP Hello Interval on page 25](#)
- [Configuring RSVP Authentication on page 26](#)
- [Configuring the Bandwidth Subscription for Class Types on page 26](#)
- [Configuring the RSVP Update Threshold on an Interface on page 27](#)
- [Configuring RSVP for Unnumbered Interfaces on page 28](#)

Configuring RSVP Refresh Reduction

You can configure RSVP refresh reduction on each interface by including the following statements in the interface configuration:

- **aggregate** and **reliable**—Enable all RSVP refresh reduction features: RSVP message bundling, RSVP message ID, reliable message delivery, and summary refresh.

In order to have refresh reduction and reliable delivery, you must include the **aggregate** and **reliable** statements.

- **no-aggregate**—Disable RSVP message bundling and summary refresh.
- **no-reliable**—Disable RSVP message ID, reliable message delivery, and summary refresh.

For more information on RSVP refresh reduction, see [“RSVP Refresh Reduction” on page 9](#).

If the **no-reliable** statement is configured on the router (reliable message delivery is disabled), the router accepts RSVP messages that include the Message ID object but ignores the Message ID object and continues performing standard message processing. No error is generated in this case, and RSVP operates normally.

However, not all combinations between two neighbors with different refresh reduction capabilities function correctly. For example, a router is configured with either the **aggregate** statement and **no-reliable** statement or with the **reliable** and **no-aggregate** statements. If an RSVP neighbor sends a Summary Refresh object to this router, no error is generated, but the Summary Refresh object cannot be processed. Consequently, RSVP states can time out on this router if the neighbor is relying only on Summary Refresh to refresh those RSVP states.

We recommend, unless there are specific requirements, that you configure RSVP refresh reduction in a similar manner on each RSVP neighbor.

To enable all RSVP refresh reduction features on an interface, include the **aggregate** statement:

```
aggregate;
```

You can include this statement at the following hierarchy levels:

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

To disable RSVP message bundling and summary refresh, include the **no-aggregate** statement:

```
no-aggregate;
```

You can include this statement at the following hierarchy levels:

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

To enable RSVP message ID and reliable message delivery on an interface, include the **reliable** statement:

```
reliable;
```

You can include this statement at the following hierarchy levels:

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

To disable RSVP message ID, reliable message delivery, and summary refresh, include the **no-reliable** statement:

```
no-reliable;
```


You can include this statement at the following hierarchy levels:

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

Determining the Refresh Reduction Capability of RSVP Neighbors

To determine the RSVP refresh reduction capability of an RSVP neighbor, you need the following information:

- The RR bit advertised by the neighbor
- The local configuration of RSVP refresh reduction
- The actual RSVP messages received from the neighbor

To obtain this information, you can issue a **show rsvp neighbor detail** command. Sample output follows:

```
user@host> show rsvp neighbor detail
RSVP neighbor: 6 learned
  Address: 192.168.224.178 via: fxp1.0 status: Up
    Last changed time: 10:06, Idle: 5 sec, Up cnt: 1, Down cnt: 0
    Message received: 36
    Hello: sent 69, received: 69, interval: 9 sec
    Remote instance: 0x60b8feba, Local instance: 0x74bc7a8d
    Refresh reduction: not operational

  Address: 192.168.224.186 via: fxp2.0 status: Down
    Last changed time: 10:17, Idle: 40 sec, Up cnt: 0, Down cnt: 0
    Message received: 6
    Hello: sent 20, received: 0, interval: 9 sec
    Remote instance: 0x0, Local instance: 0x2ae1b339
    Refresh reduction: incomplete
    Remote end: disabled, Ack-extension: enabled

  Address: 192.168.224.188 via: fxp2.0 status: Up
    Last changed time: 4:15, Idle: 0 sec, Up cnt: 1, Down cnt: 0
    Message received: 55
    Hello: sent 47, received: 31, interval: 9 sec
    Remote instance: 0x6436a35b, Local instance: 0x663849f0
    Refresh reduction: operational
    Remote end: enabled, Ack-extension: enabled
```

For more information on the **show rsvp neighbor detail** command, see the Junos OS Operational Mode Commands.

Configuring the RSVP Hello Interval

RSVP monitors the status of the interior gateway protocol (IGP) (IS-IS or OSPF) neighbors and relies on the IGP protocols to detect when a node fails. If an IGP protocol declares a neighbor down (because hello packets are no longer being received), RSVP also brings down that neighbor. However, the IGP protocols and RSVP still act independently when bringing a neighbor up.

For Juniper Networks routers, configuring a shorter or longer RSVP hello interval has no impact on whether or not an RSVP session is brought down. RSVP sessions are kept up

even if RSVP hello packets are no longer being received. RSVP sessions are maintained until either the router stops receiving IGP hello packets or the RSVP Path and Resv messages time out.

However, the RSVP hello interval might impact when another vendor's equipment brings down an RSVP session. For example, a neighboring non-Juniper Networks router might be configured to monitor RSVP hello packets.

To modify how often RSVP sends hello packets, include the **hello-interval** statement:

```
hello-interval seconds;
```

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

Configuring RSVP Authentication

All RSVP protocol exchanges can be authenticated to guarantee that only trusted neighbors participate in setting up reservations. By default, RSVP authentication is disabled.

RSVP authentication uses a Hashed Message Authentication Code (HMAC)-MD5 message-based digest. This scheme produces a message digest based on a secret authentication key and the message contents. (The message contents also include a sequence number.) The computed digest is transmitted with RSVP messages. Once you have configured authentication, all received and transmitted RSVP messages with all neighbors are authenticated on this interface.

MD5 authentication provides protection against forgery and message modification. It also can prevent replay attacks. However, it does not provide confidentiality, because all messages are sent in clear text.

By default, authentication is disabled. To enable authentication, configure a key on each interface by including the **authentication-key** statement:

```
authentication-key key;
```

You can include this statement at the following hierarchy levels:

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

Configuring the Bandwidth Subscription for Class Types

By default, RSVP allows 100 percent of the bandwidth for a class type to be used for RSVP reservations. When you oversubscribe a class type for a multiclass LSP, the aggregate demand of all RSVP sessions is allowed to exceed the actual capacity of the class type.

For detailed instructions on how to configure the bandwidth subscription for class types, see Configuring the Bandwidth Subscription Percentage for LSPs.

Configuring the RSVP Update Threshold on an Interface

The interior gateway protocols (IGPs) maintain the traffic engineering database, but the current available bandwidth on the traffic engineering database links originates from RSVP. When a link's bandwidth changes, RSVP informs the IGPs, which can then update the traffic engineering database and forward the new bandwidth information to all network nodes. The network nodes then know how much bandwidth is available on the traffic engineering database link (local or remote), and CSPF can correctly compute the paths.

However, IGP updates can consume excessive system resources. Depending on the number of nodes in a network, it might not be desirable to perform an IGP update for small changes in bandwidth. By configuring the **update-threshold** statement at the **[edit protocols rsvp]** hierarchy level, you can adjust the threshold at which a change in the reserved bandwidth triggers an IGP update.

You can configure a value of from 1 percent through 20 percent (the default is 10 percent) for when to trigger an IGP update. If the change in the reserved bandwidth is greater than or equal to the configured threshold percentage of the static bandwidth on that interface, then an IGP update occurs. For example, if you have configured the **update-threshold** statement to be 15 percent and the router discovers that the reserved bandwidth on a link has changed by 10 percent of the link bandwidth, RSVP does not trigger an IGP update. However, if the reserved bandwidth on a link changes by 20 percent of the link bandwidth, RSVP triggers an IGP update.

To adjust the threshold at which a change in the reserved bandwidth triggers an IGP update, include the **update-threshold** statement:

update-threshold *percentage*;

You can include this statement at the following hierarchy levels:

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

Because of the update threshold, it is possible for Constrained Shortest Path First (CSPF) to compute a path using outdated traffic engineering database bandwidth information on a link. If RSVP attempts to establish an LSP over that path, it might find that there is insufficient bandwidth on that link. When this happens, RSVP triggers an IGP traffic engineering database update, flooding the updated bandwidth information on the network. CSPF can then recompute the path by using the updated bandwidth information, and attempt to find a different path, avoiding the congested link. Note that this functionality is the default and does not need any additional configuration.

You can configure the **rsvp-error-hold-time** statement at the **[edit protocols mpls]** hierarchy level or the **[edit logical-systems logical-system-name protocols mpls]** hierarchy level to improve the accuracy of the traffic engineering database (including the accuracy of bandwidth estimates for LSPs) using information provided by PathErr messages. See Improving Traffic Engineering Database Accuracy with RSVP PathErr Messages.

Configuring RSVP for Unnumbered Interfaces

The Junos OS supports RSVP traffic engineering over unnumbered interfaces. Traffic engineering information about unnumbered links is carried in the IGP traffic engineering extensions for OSPF and IS-IS as described in RFC 4203, *OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)*, and RFC 4205, *Intermediate System to Intermediate System (IS-IS) Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)*. Unnumbered links can also be specified in the MPLS traffic engineering signaling as described in RFC 3477, *Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)*. This feature allows you avoid having to configure IP addresses for each interface participating in the RSVP-signaled network.

To configure RSVP for unnumbered interfaces, you must configure the router with a router ID using the **router-id** statement specified at the **[edit routing-options]** hierarchy level. The router ID must be available for routing (you can typically use the loopback address). The RSVP control messages for the unnumbered links are sent using the router ID address (rather than a randomly selected address).

To configure link protection and fast reroute on a router with unnumbered interfaces enabled, you must configure at least two addresses. We recommend that you configure a secondary interface on the loopback in addition to configuring the router ID.

Configuring RSVP Node ID Hellos

You can configure node-ID based RSVP hellos to ensure that Juniper Networks routers can interoperate with the equipment of other vendors. By default, the Junos OS uses interface-based RSVP hellos. Node-ID based RSVP hellos are specified in RFC 4558, *Node-ID Based Resource Reservation Protocol (RSVP) Hello: A Clarification Statement*. RSVP node-ID hellos are useful if you have configured BFD to detect problems over RSVP interfaces, allowing you to disable interface-hellos for these interfaces. You can also use node-ID hellos for graceful-restart procedures.

Node-ID hellos can be enabled globally for all RSVP neighbors. By default, node-ID hello support is disabled. If you have not enabled RSVP node IDs on the router, the Junos OS does not accept any node-ID hello packets.

To enable RSVP node-ID hellos globally on the router, include the **node-hello** statement:

```
node-hello;
```

You can include this statement at the following hierarchy levels:

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

You can also explicitly disable RSVP interface hellos globally. This type of configuration might be necessary in networks where the Juniper Networks router has numerous RSVP connections with equipment from other vendors. However, if you disable RSVP interface hellos globally, you can also configure a hello interval on an RSVP interface using the **hello-interval** statement. This configuration disables RSVP interface hellos globally, but

enables RSVP interface hellos on the specified interface (the RSVP interface you configure the **hello-interval** statement on). This configuration might be necessary in a heterogeneous network in which some devices support RSVP node ID hellos and other devices support RSVP interface hellos.

To disable RSVP interface hellos globally on the router, include the **no-interface-hello** statement:

```
no-interface-hello;
```

You can include this statement at the following hierarchy levels:

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

Configuring Hello Acknowledgments for Nonsession RSVP Neighbors

The **hello-acknowledgements** statement controls the hello acknowledgment behavior between RSVP neighbors regardless of whether or not they are in the same session.

Hello messages received from RSVP neighbors that are not part of a common RSVP session are discarded. If you configure the **hello-acknowledgements** statement at the **[edit protocols rsvp]** hierarchy level, hello messages from nonsession neighbors are acknowledged with a hello acknowledgment message. When hellos are received from nonsession neighbors, an RSVP neighbor relationship is created and periodic hello messages can now be received from the nonsession neighbor. The **hello-acknowledgements** statement is disabled by default. Configuring this statement allows RSVP-capable routers to be discovered using hello packets and verifies that the interface is able to receive RSVP packets before sending any MPLS LSP setup messages.

Once you enable hello acknowledgments for nonsession RSVP neighbors, the router continues to acknowledge hello messages from any nonsession RSVP neighbors unless the interface itself goes down or you change the configuration. Interface-based neighbors are not automatically aged out.

```
hello-acknowledgements;
```

You can include this statement at the following hierarchy levels:

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

Configuring Node Protection or Link Protection for LSPs

When you configure node protection or link protection on a router, bypass LSPs are created to the next-hop or next-next-hop routers for the LSPs traversing the router. You must configure node protection or link protection for each LSP that you want protected. To extend protection along the entire path used by an LSP, you must configure protection on each router that the LSP traverses.

You can configure node protection or link protection for both static and dynamic LSPs.

To configure node protection on a router for a specified LSP, include the **node-link-protection** statement:

```
node-link-protection;
```

You can include this statement at the following hierarchy levels:

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

To configure link protection on a router for a specified LSP, include the link-protection statement:

```
link-protection;
```

You can include this statement at the following hierarchy levels:

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



NOTE: To complete the configuration of node or link protection, you must also configure link protection on all unidirectional RSVP interfaces that the LSPs traverse, as described in [“Configuring Link Protection on Interfaces Used by LSPs” on page 32](#).

Switching LSPs Away from a Network Node

You can configure the router to switch active LSPs away from a network node using a bypass LSP enabled for an interface. This feature might be used to maintain active networks when a device needs to be replaced without interrupting traffic transiting the network. The LSPs can be either static or dynamic.

1. You first need to configure either link or node protection for the traffic that needs to pass around the network device you intend to disable. To function properly, the bypass LSP must use a different logical interface than the protected LSP.
2. To prepare the router to begin switching traffic away from a network node, configure the **always-mark-connection-protection-tlv** statement:

```
always-mark-connection-protection-tlv;
```

The router then marks all OAM traffic transiting this interface in preparation for switching the traffic to an alternate path based on the OAM functionality.

You can configure this statement at the following hierarchy levels:

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

3. You then need to configure the **switch-away-lsps** statement to switch the traffic from the protected LSP to the bypass LSP, effectively bypassing the default downstream network device. The actual link itself is not brought down by this configuration.

To configure the router to switch traffic away from a network node, configure the **switch-away-lsps** statement:

```
switch-away-lsps;
```

You can configure this statement at the following hierarchy levels:

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

Note the following limitations related to switching active LSPs away from a network node:

- The switch-away feature is supported on MX Series routers only.
- The switch-away feature is not supported for switching traffic from primary point-to-multipoint LSPs to bypass point-to-multipoint LSPs. If you configure the **switch-away-lsps** statement for a point-to-multipoint LSP, traffic is not switched to the bypass point-to-multipoint LSP.
- If you configure the switch-away feature on an interface along the path of a dynamic LSP, new dynamic LSPs cannot be established over that path. The switch-away feature prevents the make-before-break behavior of RSVP-signaled LSPs. The make-before-break behavior normally causes the router to first attempt to re-signal a dynamic LSP before tearing down the original.

Related Documentation

- [Configuring Node Protection or Link Protection for LSPs on page 29](#)

Configuring Inter-AS Node and Link Protection

To interoperate with other vendors' equipment, the Junos OS supports the record route object (RRO) node ID subobject for use in inter-AS link and node protection configurations. The RRO node ID subobject is defined in RFC 4561, *Definition of a Record Route Object (RRO) Node-Id Sub-Object*. This functionality is enabled by default in Junos OS Release 9.4 and later.

If you have Juniper Networks routers running Junos OS Release 9.4 and later releases in the same MPLS-TE network as routers running Junos OS Release 8.4 and earlier releases, you might need to disable the RRO node ID subobject by configuring the **no-node-id-subobject** statement:

```
no-node-id-subobject;
```

You can include this statement at the following hierarchy levels:

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

Configuring Link Protection on Interfaces Used by LSPs

When you configure node protection or link protection on a router for LSPs as described in “[Configuring Node Protection or Link Protection for LSPs](#)” on page 29, you also must configure the **link-protection** statement on the RSVP interfaces used by the LSPs.

To configure link protection on the interfaces used by the LSPs, include the link-protection statement:

```
link-protection {  
  disable;  
  admin-group  
    exclude group-names;  
    include-all group-names;  
    include-any group-names;  
}  
bandwidth bps;  
bypass bypass-name {  
  bandwidth bps;  
  description text;  
  hop-limit number;  
  no-cspf;  
  path address <strict | loose>;  
  priority setup-priority reservation-priority;  
  to address;  
}  
class-of-service cos-value;  
hop-limit number;  
max-bypasses number;  
no-cspf;  
no-node-protection;  
optimize-timer seconds;  
path address <strict | loose>;  
priority setup-priority reservation-priority;  
subscription percent {  
  ct0 percent;  
  ct1 percent;  
  ct2 percent;  
  ct3 percent;  
}  
}
```

You can include this statement at the following hierarchy levels:

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

All the statements under **link-protection** are optional.

The following sections describe how to configure link protection:

- [Configuring Bypass LSPs on page 33](#)
- [Configuring Administrative Groups for Bypass LSPs on page 34](#)

- [Configuring the Bandwidth for Bypass LSPs on page 34](#)
- [Configuring Class of Service for Bypass LSPs on page 35](#)
- [Configuring the Hop Limit for Bypass LSPs on page 35](#)
- [Configuring the Maximum Number of Bypass LSPs on page 36](#)
- [Disabling CSPF for Bypass LSPs on page 36](#)
- [Disabling Node Protection for Bypass LSPs on page 36](#)
- [Configuring the Optimization Interval for Bypass LSPs on page 37](#)
- [Configuring an Explicit Path for Bypass LSPs on page 37](#)
- [Configuring the Amount of Bandwidth Subscribed for Bypass LSPs on page 38](#)
- [Configuring Priority and Preemption for Bypass LSPs on page 38](#)

Configuring Bypass LSPs

You can configure specific bandwidth and path constraints for a bypass LSP. You can also individually configure each bypass LSP generated when you enable multiple bypass LSPs. If you do not configure the bypass LSPs individually, they all share the same path and bandwidth constraints (if any).

If you specify the **bandwidth**, **hop-limit**, and **path** statements for the bypass LSP, these values take precedence over the values configured at the **[edit protocols rsvp interface interface-name link-protection]** hierarchy level. The other attributes (**subscription**, **no-node-protection**, and **optimize-timer**) are inherited from the general constraints.

To configure a bypass LSP, specify a name for the bypass LSP using the **bypass** statement. The name can be up to 64 characters in length.

```
bypass bypass-name {
  bandwidth bps;
  description text;
  class-of-service cos-value;
  hop-limit number;
  no-cspf;
  path address <strict | loose>;
  priority setup-priority reservation-priority;
  to address;
}
```

You can include this statement at the following hierarchy levels:

- **[edit protocols rsvp interface interface-name link-protection]**
- **[edit logical-systems logical-system-name protocols rsvp interface interface-name link-protection]**

Configuring the Next-Hop or Next-Next-Hop Node Address for Bypass LSPs

If you configure a bypass LSP, you must also configure the **to** statement. The **to** statement specifies the address for the interface of the immediate next-hop node (for link protection) or the next-next-hop node (for node-link protection). The address specified determines whether this is a link protection bypass or a node-link protection bypass. On multiaccess

networks (for example, a LAN), this address is also used to specify which next-hop node is being protected.

Configuring Administrative Groups for Bypass LSPs

Administrative groups, also known as link coloring or resource class, are manually assigned attributes that describe the “color” of links, such that links with the same color conceptually belong to the same class. You can use administrative groups to implement a variety of policy-based LSP setups. You can configure administrative groups for bypass LSPs. For more information about configuring administrative groups, see *Configuring Administrative Groups*.

To configure administrative groups for bypass LSPs, include the **admin-group** statement:

```
admin-group {  
  exclude group-names;  
  include-all group-names;  
  include-any group-names;  
}
```

To configure an administrative group for all of the bypass LSPs, include the **admin-group** statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection]

To configure an administrative groups for a specific bypass LSP, include the **admin-group** statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection **bypass** *bypass-name*]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection **bypass** *bypass-name*]

Configuring the Bandwidth for Bypass LSPs

You can specify the amount of bandwidth allocated for automatically generated bypass LSPs or you can individually specify the amount of bandwidth allocated for each LSP.

If you have enabled multiple bypass LSPs, this statement is required.

To specify the bandwidth allocation, include the **bandwidth** statement:

```
bandwidth bps;
```

For automatically generated bypass LSPs, include the **bandwidth** statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection]

For individually configured bypass LSPs, include the **bandwidth** statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection **bypass** *bypass-name*]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection **bypass** *bypass-name*]

Configuring Class of Service for Bypass LSPs

You can specify the class-of-service value for bypass LSPs by including the **class-of-service** statement:

class-of-service *cos-value*;

To apply a class-of-service value to all the automatically generated bypass LSPs, include the **class-of-service** statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection]

To configure a class-of-service value for a specific bypass LSPs, include the **class-of-service** statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection **bypass** *bypass-name*]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection **bypass** *bypass-name*]

Configuring the Hop Limit for Bypass LSPs

You can specify the maximum number of hops a bypass can traverse. By default, each bypass can traverse a maximum of 255 hops (the ingress and egress routers count as one hop each, so the minimum hop limit is two).

To configure the hop limit for bypass LSPs, include the **hop-limit** statement:

hop-limit *number*;

For automatically generated bypass LSPs, include the **hop-limit** statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection]

For individually configured bypass LSPs, include the **hop-limit** statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection **bypass** *bypass-name*]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection **bypass** *bypass-name*]

Configuring the Maximum Number of Bypass LSPs

You can specify the maximum number of dynamic bypass LSPs permitted for protecting an interface using the **max-bypasses** statement at the **[edit protocols rsvp interface interface-name link-protection]** hierarchy level. When this statement is configured, multiple bypasses for link protection are enabled. Call admission control (CAC) is also enabled.

By default, this option is disabled and only one bypass is enabled for each interface. You can configure a value of between **0** through **99** for the **max-bypasses** statement. Configuring a value of **0** prevents the creation of any dynamic bypass LSPs for the interface. If you configure a value of **0** for the **max-bypasses** statement, you need to configure one or more static bypass LSPs to enable link protection on the interface.

If you configure the **max-bypasses** statement, you must also configure the **bandwidth** statement (discussed in [“Configuring the Bandwidth for Bypass LSPs” on page 34](#)).

To configure the maximum number of bypass LSPs for a protected interface, include the **max-bypasses** statement:

```
max-bypasses number;
```

You can include this statement at the following hierarchy levels:

- **[edit protocols rsvp interface interface-name link-protection]**
- **[edit logical-systems logical-system-name protocols rsvp interface interface-name link-protection]**

Disabling CSPF for Bypass LSPs

Under certain circumstances, you might need to disable CSPF computation for bypass LSPs and use the configured Explicit Route Object (ERO) if available. For example, a bypass LSP might need to traverse multiple OSPF areas or IS-IS levels, preventing the CSPF computation from working. To ensure that link and node protection function properly in this case, you have to disable CSPF computation for the bypass LSP.

You can disable CSPF computation for all bypass LSPs or for specific bypass LSPs.

To disable CSPF computation for bypass LSPs, include the **no-cspf** statement:

```
no-cspf;
```

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

Disabling Node Protection for Bypass LSPs

You can disable node protection on the RSVP interface. Link protection remains active. When this option is configured, the router can only initiate a next-hop bypass, not a next-next-hop bypass.

To disable node protection for bypass LSPs, include the **no-node-protection** statement:

```
no-node-protection;
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection]

Configuring the Optimization Interval for Bypass LSPs

You can configure an optimization interval for bypass LSPs using the **optimize-timer** statement. At the end of this interval, an optimization process is initiated that attempts to either minimize the number of bypasses currently in use, minimize the total amount of bandwidth reserved for all of the bypasses, or both. You can configure an optimization interval from 1 through 65,535 seconds. A default value of 0 disables bypass LSP optimization.

When you configure the **optimize-timer** statement, bypass LSPs are reoptimized automatically when you configure or change the configuration of any of the following:

- Administrative group for a bypass LSP—The configuration for an administrative group has been changed on a link along the path used by the bypass LSP. Configure an administrative group using the **admin-group** statement at the [edit protocols rsvp **interface** *interface-name* link-protection] hierarchy level.
- Fate sharing group—The configuration for a fate sharing group has been changed. Configure a fate sharing group using the **group** statement at the [edit routing-options **fate-sharing**] hierarchy level.
- IS-IS overload—The configuration for IS-IS overload has been changed on a router along the path used by the bypass LSP. Configure IS-IS overload using the **overload** statement at the [edit protocols isis] hierarchy level.
- IGP metric—The IGP metric has been changed on a link along the path used by the bypass LSP.

To configure the optimization interval for bypass LSPs, include the **optimize-timer** statement:

```
optimize-timer seconds;
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection]

Configuring an Explicit Path for Bypass LSPs

By default, when you establish a bypass LSP to an adjacent neighbor, CSPF is used to discover the least-cost path. The **path** statement allows you to configure an explicit path (a sequence of strict or loose routes), giving you control over where and how the bypass LSP is established. To configure an explicit path, include the **path** statement:

path *address* <strict | loose>;

For automatically generated bypass LSPs, include the **path** statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection]

For individually configured bypass LSPs, include the **path** statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection **bypass** *bypass-name*]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection **bypass** *bypass-name*]

Configuring the Amount of Bandwidth Subscribed for Bypass LSPs

You can configure the amount of bandwidth subscribed to bypass LSPs. You can configure the bandwidth subscription for the whole bypass LSP or for each class type that might traverse the bypass LSP. You can configure any value between 1 percent and 65,535 percent. By configuring a value less than 100 percent, you are undersubscribing the bypass LSPs. By configuring a value greater than 100 percent, you are oversubscribing the bypass LSPs.

The ability to oversubscribe the bandwidth for the bypass LSPs makes it possible to more efficiently use network resources. You can configure the bandwidth for the bypass LSPs based on the average network load as opposed to the peak load.

To configure the amount of bandwidth subscribed for bypass LSPs, include the **subscription** statement:

```
subscription percentage {  
  ct0 percentage;  
  ct1 percentage;  
  ct2 percentage;  
  ct3 percentage;  
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection]

Configuring Priority and Preemption for Bypass LSPs

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

For more detailed information on configuring setup priority and reservation priority for LSPs, see Configuring Priority and Preemption for LSPs.

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

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

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

Configuring RSVP Setup Protection

You can configure the facility-backup fast reroute mechanism to provide setup protection for LSPs which are in the process of being signaled. Both point-to-point LSPs and point-to-multipoint LSPs are supported. This feature is applicable in the following scenario:

1. A failed link or node is present on the strict explicit path of an LSP before the LSP is signaled.
2. There is also a bypass LSP protecting the link or node.
3. RSVP signals the LSP through the bypass LSP. The LSP appears as if it was originally set up along its primary path and then failed over to the bypass LSP because of the link or node failure.
4. When the link or node has recovered, the LSP can be automatically reverted to the primary path.

You should configure the **setup-protection** statement at the **[edit protocols rsvp]** on each of the routers along the LSP path on which you want to enable LSP setup protection. You should also configure IGP traffic engineering on all of the routers on the LSP path. You can issue a **show rsvp session** command to determine whether or not the LSP has setup protection enabled on a router acting as a point of local repair (PLR) or a merge point.

To enable RSVP setup protection, include the **setup-protection** statement

```
setup-protection;
```

You can include this statement at the following hierarchy levels:

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

Configuring RSVP Graceful Restart

The following RSVP graceful restart configurations are possible:

- Graceful restart and helper mode are both enabled (the default).
- Graceful restart is enabled but helper mode is disabled. A router configured in this way can restart gracefully, but cannot help a neighbor with its restart and recovery procedures.
- Graceful restart is disabled but helper mode is enabled. A router configured in this way cannot restart gracefully, but can help a restarting neighbor.
- Graceful restart and helper mode both are disabled. This configuration completely disables RSVP graceful restart (including restart and recovery procedures and helper mode). The router behaves like a router that does not support RSVP graceful restart.



NOTE: In order to turn on RSVP graceful restart, you must set the global graceful restart timer to at least 180 seconds.

The following sections describe how to configure RSVP graceful restart:

- [Enabling Graceful Restart for All Routing Protocols on page 40](#)
- [Disabling Graceful Restart for RSVP on page 40](#)
- [Disabling RSVP Helper Mode on page 41](#)
- [Configuring the Maximum Helper Recovery Time on page 41](#)
- [Configuring the Maximum Helper Restart Time on page 41](#)

Enabling Graceful Restart for All Routing Protocols

To enable graceful restart for RSVP, you need to enable graceful restart for all the protocols that support graceful restart on the router. For more information about graceful restart, see the Junos OS Routing Protocols Configuration Guide.

To enable graceful restart on the router, include the **graceful-restart** statement:

```
graceful-restart;
```

You can include this statement at the following hierarchy levels:

- **[edit routing-options]**
- **[edit logical-systems *logical-system-name* routing-options]**

Disabling Graceful Restart for RSVP

By default, RSVP graceful restart and RSVP helper mode are enabled when you enable graceful restart. However, you can disable one or both of these capabilities.

To disable RSVP graceful restart and recovery, include the **disable** statement at the **[edit protocols rsvp graceful-restart]** hierarchy level:

```
disable;
```

Disabling RSVP Helper Mode

To disable RSVP helper mode, include the **helper-disable** statement at the **[edit protocols rsvp graceful-restart]** hierarchy level:

```
helper-disable;
```

Configuring the Maximum Helper Recovery Time

To configure the amount of time the router retains the state of its RSVP neighbors while they undergo a graceful restart, include the **maximum-helper-recovery-time** statement at the **[edit protocols rsvp graceful-restart]** hierarchy level. This value is applied to all neighboring routers, so it should be based on the time required by the slowest RSVP neighbor to recover.

```
maximum-helper-recovery-time seconds;
```

Configuring the Maximum Helper Restart Time

To configure the delay between when the router discovers that a neighboring router has gone down and when it declares the neighbor down, include the **maximum-helper-restart-time** statement at the **[edit protocols rsvp graceful-restart]** hierarchy level. This value is applied to all neighboring routers, so it should be based on the time required by the slowest RSVP neighbor to restart.

```
maximum-helper-restart-time seconds;
```

Configuring Load Balancing Across RSVP LSPs

By default, when you have configured several RSVP LSPs to the same egress router, the LSP with the lowest metric is selected and carries all traffic. If all of the LSPs have the same metric, one of the LSPs is selected at random and all traffic is forwarded over it.

Alternatively, you can load-balance traffic across all of the LSPs by enabling per-packet load balancing.

To enable per-packet load balancing on an ingress LSP, configure the **policy-statement** statement as follows:

```
[edit policy-options]
policy-statement policy-name {
  then {
    load-balance per-packet;
  }
  accept;
}
```

You then need to apply this statement as an export policy to the forwarding table. For more information on how to configure the **policy-statement** statement, see the Routing Policy Configuration Guide.

Once per-packet load balancing is applied, traffic is distributed equally between the LSPs (by default).

You need to configure per-packet load balancing if you want to enable PFE fast reroute. To enable PFE fast reroute, include the **policy-statement** statement for per-packet load balancing shown in this section in the configuration of each of the routers where a reroute might take place. See also Configuring Fast Reroute.

You can also load-balance the traffic between the LSPs in proportion to the amount of bandwidth configured for each LSP. This capability can better distribute traffic in networks with asymmetric bandwidth capabilities across external links, since the configured bandwidth of an LSP typically reflects the traffic capacity of that LSP.

To configure RSVP LSP load balancing, include the **load-balance** statement with the **bandwidth** option:

```
load-balance {  
    bandwidth;  
}
```

You can configure this statement at the following hierarchy levels:

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

Keep the following information in mind when you use the **load-balance** statement:

- If you configure the **load-balance** statement, the behavior of currently running LSPs is not altered. To force currently running LSPs to use the new behavior, you can issue a **clear mpls lsp** command.
- The **load-balance** statement only applies to ingress LSPs that have per-packet load balancing enabled.
- For Differentiated Services–aware traffic engineered LSPs, the bandwidth of an LSP is calculated by summing the bandwidth of all of the class types.

Configuring RSVP Automatic Mesh

BGP and MPLS VPNs are based on a peer model. To add a new site to an existing VPN, you need to configure the CE router at the new site and the PE router connected to the CE router. You do not have to modify the configuration of all of the other PE routers participating in the VPN. The PE routers automatically learn about the routes associated with the new site (a process called automatic discovery).

The requirements are a bit different if you need to add a new PE router (as opposed to a CE router) to the network. A BGP and MPLS VPN requires that the BGP session be fully meshed and that there also be a full mesh of PE router-to-PE router MPLS LSPs between all of the PE routers in the network. When you add a new PE router to the network, all of the existing PE routers must be reconfigured to peer with the new PE router. Much of the configuration effort can be reduced if you configure BGP route reflectors (mitigating the

full mesh requirement for BGP) and if you configure LDP as the signaling protocol for MPLS.

However, if you need to add a new PE router to a network configured with a full mesh of RSVP-signaled LSPs, you need to reconfigure each of the PE routers to have a peer relationship with the new PE router. You can configure RSVP automatic mesh to address this particular operational scenario. When you enable RSVP automatic mesh, RSVP LSPs are dynamically created between a new PE router and the existing PE routers, eliminating the need to reconfigure all of the PE routers manually. For dynamic LSP creation to function properly, BGP must be configured to exchange routes between all of the participating PE routers. If two BGP peers did not exchange routes, it would not be possible to configure a dynamic LSP between them. The local router's inet.0 routing table must include a labeled route to each potential IBGP next-hop (future potential PE routers or LSP destinations).

RSVP includes numerous capabilities that are not available in LDP, including fast reroute. RSVP automatic mesh helps to reduce the operation and maintenance requirements for RSVP, making it possible to deploy RSVP in larger and more complicated networks.

Every PE router can reach every other PE router in the network because this information is distributed by the IGP. A PE router can set up an RSVP LSP to any other PE router in the network so long as it knows that such an LSP is required. To build a full mesh of LSPs between the PE routers requires that each PE router know which of the other PE routers make up the full mesh.

You can configure RSVP to establish LSPs automatically for any new PE router added to a full mesh of LSPs. To enable this feature, you must configure the **rsvp-te** statement on all of the PE routers in the full mesh.



NOTE: You cannot configure RSVP automatic mesh in conjunction with CCC. CCC cannot use the dynamically generated LSPs.

To configure RSVP automatic mesh, include the **rsvp-te** statement:

```
rsvp-te {
  destination-networks network-prefix;
  label-switched-path-template {
    default-template;
    template-name;
  }
}
```

You can configure these statements at the following hierarchy levels:

- [edit routing-options dynamic-tunnels *tunnel-name*]
- [edit logical-systems *logical-system-name* routing-options dynamic-tunnels *tunnel-name*]

You must also configure the following statements to enable RSVP automatic mesh:

- **destination-networks**—Specify the IP version 4 (IPv4) prefix range for the destination network. Dynamic tunnels within the specified IPv4 prefix range can be initiated.

- **label-switched-path-template**—You can configure either the default template explicitly using the **default-template** option, or you can configure an LSP template of your own using the **template-name** option. The LSP template acts as a model configuration for the dynamically generated LSPs.

**Related
Documentation**

- Flooding Unknown Traffic Using Point-to-Multipoint LSPs

Configuring Timers for RSVP Refresh Messages

RSVP uses two related timing parameters:

- **refresh-time**—The refresh time controls the interval between the generation of successive refresh messages. The default value for the refresh time is 45 seconds. This number is derived from the **refresh-time** statement's default value of 30, multiplied by a fixed value of 1.5. This computation differs from RFC 2205, which states that the refresh time should be multiplied by a random value in the range from 0.5 through 1.5.

Refresh messages include path and Resv messages. Refresh messages are sent periodically so that reservation states in neighboring nodes do not time out. Each path and Resv message carries the refresh timer value, and the receiving node extracts this value from the messages.

- **keep-multiplier**—The keep multiplier is a small, locally configured integer from 1 through 255. The default value is 3. It indicates the number of messages that can be lost before a particular state is declared stale and must be deleted. The keep multiplier directly affects the lifetime of an RSVP state.

To determine the lifetime of a reservation state, use the following formula:

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

In the worst case, (**keep-multiplier** – 1) successive refresh messages must be lost before a reservation state is deleted.

We do not recommend configuring a short RSVP hello timer. If quick discovery of a failed neighbor is needed, configure short IGP (OSPF or IS-IS) hello timers.

By default, the refresh timer value is 30 seconds. To modify this value, include the **refresh-time** statement:

```
refresh-time seconds;
```

You can include this statement at the following hierarchy levels:

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

The default value of the keep multiplier is 3. To modify this value, include the **keep-multiplier** statement:

```
keep-multiplier number;
```

You can include this statement at the following hierarchy levels:

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

Preempting RSVP Sessions

Whenever bandwidth is insufficient to handle all RSVP sessions, you can control the preemption of RSVP sessions. By default, an RSVP session is preempted only by a new higher-priority session.

To always preempt a session when the bandwidth is insufficient, include the **preemption** statement with the **aggressive** option:

```
preemption aggressive;
```

You can include this statement at the following hierarchy levels:

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

To disable RSVP session preemption, include the **preemption** statement with the **disabled** option:

```
preemption disabled;
```

To return to the default (that is, preempt a session only for a new higher-priority session), include the **preemption** statement with the **normal** option:

```
preemption normal;
```

You can include this statement at the following hierarchy levels:

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

Configuring MTU Signaling in RSVP

To configure maximum transmission unit (MTU) signaling in RSVP, you need to configure MPLS to allow IP packets to be fragmented before they are encapsulated in MPLS. You also need to configure MTU signaling in RSVP. For troubleshooting purposes, you can configure MTU signaling alone without enabling packet fragmentation.

To configure MTU signaling in RSVP, include the **path-mtu** statement:

```
path-mtu {  
  allow-fragmentation;  
  rsvp {  
    mtu-signaling;  
  }  
}
```

You can include this statement at the following hierarchy levels:

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

The following sections describe how to enable packet fragmentation and MTU signaling in RSVP:

- [Enabling MTU Signaling in RSVP on page 46](#)
- [Enabling Packet Fragmentation on page 46](#)

Enabling MTU Signaling in RSVP

To enable MTU signaling in RSVP, include the **rsvp mtu-signaling** statement:

```
rsvp mtu-signaling;
```

You can include this statement at the following hierarchy levels:

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

Once you have committed the configuration, changes in the MTU signaling behavior for RSVP take effect the next time the path is refreshed.

You can configure the **mtu-signaling** statement by itself at the **[edit protocols mpls path-mtu rsvp]** hierarchy level. This can be useful for troubleshooting. If you configure just the **mtu-signaling** statement, you can use the **show rsvp session detail** command to determine what the smallest MTU is on an LSP. The **show rsvp session detail** command displays the MTU value received and sent in the Adspec object.

Enabling Packet Fragmentation

To allow IP packets to be fragmented before they are encapsulated in MPLS, include the **allow-fragmentation** statement:

```
allow-fragmentation;
```

You can include this statement at the following hierarchy levels:

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




NOTE: Do not configure the **allow-fragmentation** statement alone. Always configure it in conjunction with the **mtu-signaling** statement.

Configuring Ultimate-Hop Popping for LSPs

By default, RSVP-signaled LSPs use penultimate-hop popping (PHP). [Figure 3 on page 47](#) illustrates a penultimate-hop popping LSP between Router PE1 and Router PE2. Router


CE1 forwards a packet to its next hop (Router PE1), which is also the LSP ingress. Router PE1 pushes label 1 on the packet and forwards the labeled packet to Router P1. Router P1 completes the standard MPLS label swapping operation, swapping label 1 for label 2, and forwards the packet to Router P2. Since Router P2 is the penultimate-hop router for the LSP to Router PE2, it first pops the label and then forwards the packet to Router PE2. When Router PE2 receives it, the packet can have a service label, an explicit-null label, or just be a plain IP or VPLS packet. Router PE2 forwards the unlabeled packet to Router CE2.

Figure 3: Penultimate-Hop Popping for an LSP



You can also configure ultimate-hop popping (UHP) (as shown in [Figure 4 on page 47](#)) for RSVP-signaled LSPs. Some network applications can require that packets arrive at the egress router (Router PE2) with a non-null outer label. For an ultimate-hop popping LSP, the penultimate router (Router P2 in [Figure 4 on page 47](#)) performs the standard MPLS label swapping operation (in this example, label 2 for label 3) before forwarding the packet to egress Router PE2. Router PE2 pops the outer label and performs a second lookup of the packet address to determine the end destination. It then forwards the packet to the appropriate destination (either Router CE2 or Router CE4).

Figure 4: Ultimate-Hop Popping for an LSP



The following network applications require that you configure UHP LSPs:

- MPLS-TP for performance monitoring and in-band OAM
- Edge protection virtual circuits

The following features do not support the UHP behavior:

- LDP-signaled LSPs
- Static LSPs
- Point-to-multipoint LSPs
- **traceroute** command

For more information about UHP behavior, see Internet draft draft-ietf-mpls-rsvp-te-no-php-oob-mapping-01.txt, *Non PHP behavior and Out-of-Band Mapping for RSVP-TE LSPs*.

For point-to-point RSVP-signaled LSPs, UHP behavior is signaled from the LSP ingress. Based on the ingress router configuration, RSVP can signal the UHP LSP with the non-PHP flag set. RSVP PATH messages carry the two flags in the LSP-ATTRIBUTES object. When the egress router receives the PATH message, it assigns a non-null label to the LSP. RSVP also creates and installs two routes in the mpls.0 routing table. S refers to the S bit of the MPLS label, which indicates whether or not the bottom of the label stack has been reached.

- Route S=0—Indicates that there are more labels in the stack. The next hop for this route points to the mpls.0 routing table, triggering a chained MPLS label lookup to discover the remaining MPLS labels in the stack.
- Route S=1—Indicates that there are no more labels. The next hop points to the inet.0 routing table if the platform supports chained and multi-family lookup. Alternatively, the label route can point to a VT interface to initiate IP forwarding.

If you enable UHP LSPs, MPLS applications such as Layer 3 VPNs, VPLS, Layer 2 VPNs, and Layer 2 circuits can use the UHP LSPs. The following explains how UHP LSPs affect the different types of MPLS applications:

- Layer 2 VPNs and Layer 2 circuits—A packet arrives at the PE router (egress of the UHP LSP) with two labels. The outer label (S=0) is the UHP label, and the inner label (S=1) is the VC label. A lookup based on the transport label results in a table handle for the mpls.0 routing table. There is an additional route in the mpls.0 routing table corresponding to the inner label. A lookup based on the inner label results in the CE router next hop.
- Layer 3 VPN—A packet arrives at the PE router (egress of the UHP LSP) with two labels. The outer label (S=0) is the UHP label, and the inner label is the VPN label (S=1). A lookup based on the transport label results in the table handle for the mpls.0 routing table. There are two cases in this scenario. By default, Layer 3 VPNs advertise the per-next hop label. A lookup based on the inner label results in the next hop toward the CE router. However, if you have configured the **vrf-table-label** statement for the Layer 3 VPN routing instance, the inner LSI label points to the VRF routing table. An IP lookup is also completed for the VRF routing table.



NOTE: UHP for Layer 3 VPNs configured with the **vrf-table-label** statement is supported on MX 3D Universal Edge Routers only.

- VPLS—A packet arrives at the PE router (egress of the UHP LSP) with two labels. The outer label is the transport label (S=0) and the inner label is the VPLS label (S=1). A lookup based on the transport label results in the table handle for the mpls.0 routing table. A lookup based on the inner label in mpls.0 routing table results in the LSI tunnel interface of the VPLS routing instance if tunnel-services is not configured (or a VT interface not available). MX 3D Series routers support chained lookup and multi-family lookup.



NOTE: UHP for VPLS configured with the `no-tunnel-service` statement is supported on MX 3D Series routers only.

- **IPv4 over MPLS**—A packet arrives at the PE router (egress of the UHP LSP) with one label ($S=1$). A lookup based on this label returns a VT tunnel interface. Another IP lookup is completed on the VT interface to determine where to forward the packet. If the routing platform supports multi-family and chained lookups (for example, MX 3D routers), lookup based on label route ($S=1$) points to the `inet.0` routing table.
- **IPv6 over MPLS**—For IPv6 tunneling over MPLS, PE routers advertise IPv6 routes to each other with a label value of 2. This is the explicit null label for IPv6. As a result, the forwarding next hops for IPv6 routes that are learned from remote PE routers normally push two labels. The inner label is 2 (it could be different if the advertising PE router is from another vendor), and the router label is the LSP label. Packets arrive at the PE router (egress of the UHP LSP) with two labels. The outer label is the transport label ($S=0$), and the inner label is the IPv6 explicit-null label (label 2). Lookup based on the inner label in the `mpls.0` routing table redirects back to the `mpls.0` routing table. On MX 3D Series routers, the inner label (label 2) is stripped off and an IPv6 lookup is done using the `inet6.0` routing table.
- **Enabling both PHP and UHP LSPs**—You can configure both PHP and UHP LSPs over the same network paths. You can separate PHP and UHP traffic by selecting forwarding LSP next hops using a regular expression with the `install-nexthop` statement. You can also separate traffic by simply naming the LSPs appropriately.

The following statements enable ultimate-hop popping for an LSP. You can enable this feature on a specific LSP or for all of the ingress LSPs configured on the router. Configure these statements on the router at the LSP ingress.

1. To enable ultimate-hop popping, include the `ultimate-hop-popping` statement:

`ultimate-hop-popping;`

Include this statement at the `[edit protocols mpls label-switched-path label-switched-path-name]` hierarchy level to enable ultimate-hop popping on a specific LSP. Include this statement at the `[edit protocols mpls]` hierarchy level to enable ultimate-hop popping on all of the ingress LSPs configured on the router. You can also configure the `ultimate-hop-popping` statement under the equivalent `[edit logical-routers]` hierarchy levels.



NOTE: When you enable ultimate-hop popping, RSVP attempts to resignal existing LSPs as ultimate-hop popping LSPs in a make-before-break fashion. If an egress router does not support ultimate-hop popping, the existing LSP is torn down (RSVP sends a PathTear message along an LSP's path, removing the path state and dependent reservation state and releasing the associated networking resources).

If you disable ultimate-hop popping, RSVP resignals existing LSPs as penultimate-hop popping LSPs in a make-before-break fashion.

2. If you want to enable both ultimate-hop-popping and chained next hops on MX 3D Series routers only, you also need to configure the **enhanced-ip** option for the **network-services** statement:

```
network-services enhanced-ip;
```

You configure this statement at the **[edit chassis]** hierarchy level. Once you have configured the **network-services** statement, you need to reboot the router to enable UHP behavior.

**Related
Documentation**

- Label Allocation
- Configuring Corouted Bidirectional LSPs
- network-services
- [ultimate-hop-popping on page 102](#)

Configuring RSVP to Pop the Label on the Ultimate-Hop Router

You can control the label value advertised on the egress router of an LSP. The default advertised label is label 3 (Implicit Null label). If label 3 is advertised, the penultimate-hop router removes the label and sends the packet to the egress router. When ultimate-hop popping is enabled, label 0 (IP version 4 [IPv4] Explicit Null label) is advertised. Ultimate-hop popping ensures that any packets traversing an MPLS network include a label.



NOTE: Juniper Networks routers queue packets based on the incoming label. Routers from other vendors might queue packets differently. Keep this in mind when working with networks containing routers from multiple vendors.

To configure ultimate-hop popping for RSVP, include the **explicit-null** statement:

```
explicit-null;
```

You can include this statement at the following hierarchy levels:

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

**Related
Documentation**

- Label Description
- Label Allocation

Enabling Ultimate-Hop Popping on Point-to-Multipoint LSPs

By default, for both point-to-point and point-to-multipoint LSPs, penultimate-hop popping is used for MPLS traffic. MPLS labels are removed from packets on the router just before the egress router of the LSP. The plain IP packets are then forwarded to the

egress router. For ultimate-hop popping, the egress router is responsible for both removing the MPLS label and processing the plain IP packet.

It can be beneficial to enable ultimate-hop popping on point-to-multipoint LSPs, particularly when transit traffic is traversing the same egress device. If you enable ultimate-hop popping, a single copy of traffic can be sent over the incoming link, saving significant bandwidth. By default, ultimate-hop popping is disabled.

You enable ultimate-hop popping for point-to-multipoint LSPs by configuring the **tunnel-services** statement. When you enable ultimate-hop popping, the Junos OS selects one of the available virtual loopback tunnel (VT) interfaces to loop back the packets to the PFE for IP forwarding. By default, the VT interface selection process is performed automatically. Bandwidth admission control is used to limit the number of LSPs that can be used on one VT interface. Once all the bandwidth is consumed on one interface, the Junos OS selects another VT interface with sufficient bandwidth for admission control.

If an LSP requires more bandwidth than is available from any of the VT interfaces, ultimate-hop popping cannot be enabled and penultimate-hop popping is enabled instead.

For ultimate-hop popping on point-to-multipoint LSPs to function properly, the egress router must have a PIC that provides tunnel services, such as the tunnel services PIC or the adaptive services PIC. Tunnel services are needed for popping the final MPLS label and for returning packets for IP address lookups.

You can explicitly configure which VT interfaces handle the RSVP traffic by including the **devices** option for the **tunnel-services** statement. The **devices** option allows you to specify which VT interfaces are to be used by RSVP. If you do not configure this option, all of the VT interfaces available to the router can be used.

To enable ultimate-hop popping for the egress point-to-multipoint LSPs on a router, configure the **tunnel-services** statement:

```
tunnel-services {
  devices device-names;
}
```

You can configure this statement at the following hierarchy levels:

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

To enable ultimate-hop popping for egress point-to-multipoint LSPs, you must also configure the **interface** statement with the **all** option:

```
interface all;
```

You must configure this statement at the **[edit protocols rsvp]** hierarchy level.

Tracing RSVP Protocol Traffic

To trace RSVP protocol traffic, include the **traceoptions** statement:

```
traceoptions {  
  file filename <files number> <size size> <world-readable | no-world-readable>;  
  flag flag <flag-modifier> <disable>;  
}
```

You can include this statement at the following hierarchy levels:

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

You can specify the following RSVP-specific flags in the RSVP **traceoptions** statement:

Use the **file** statement to specify the name of the file that receives the output of the tracing operation. All files are placed in the directory **/var/log**. We recommend that you place RSVP tracing output in the file **rsvp-log**.

- **all**—All tracing operations.
- **error**—All detected error conditions
- **event**—RSVP-related events (helps to trace events related to RSVP graceful restart)
- **lmp**—RSVP-Link Management Protocol (LMP) interactions
- **packets**—All RSVP packets
- **path**—All path messages
- **pathtear**—PathTear messages
- **resv**—Resv messages
- **resvtear**—ResvTear messages
- **route**—Routing information
- **state**—Session state transitions

For general information about tracing and global tracing options, see the Junos OS Routing Protocols Configuration Guide.

Examples: Tracing RSVP Protocol Traffic

Trace RSVP path messages in detail:

```
[edit]  
protocols {  
  rsvp {  
    traceoptions {  
      file rsvp size 10m files 5;  
      flag path;  
    }  
  }  
}
```

```
}
```

Trace all RSVP messages:

```
[edit]
protocols {
  rsvp {
    traceoptions {
      file rsvp size 10m files 5;
      flag packets;
    }
  }
}
```

Trace all RSVP error conditions:

```
[edit]
protocols {
  rsvp {
    traceoptions {
      file rsvp size 10m files 5;
      flag error;
    }
  }
}
```


PART 3

Administration

- [RSVP Standards and Terminology on page 57](#)
- [Summary of RSVP Configuration Statements on page 61](#)

CHAPTER 3

RSVP Standards and Terminology

- [Supported RSVP Standards on page 57](#)
- [RSVP Graceful Restart Standard on page 58](#)
- [RSVP Graceful Restart Terminology on page 58](#)

Supported RSVP Standards

Junos OS substantially supports the following RFCs and Internet drafts, which define standards for RSVP.

- RFC 2205, *Resource ReSerVation Protocol (RSVP)—Version 1 Functional Specification*
- RFC 2210, *The Use of RSVP with IETF Integrated Services*
- RFC 2211, *Specification of the Controlled-Load Network Element Service*
- RFC 2212, *Specification of Guaranteed Quality of Service*
- RFC 2215, *General Characterization Parameters for Integrated Service Network Elements*
- RFC 2745, *RSVP Diagnostic Messages*
- RFC 2747, *RSVP Cryptographic Authentication* (updated by RFC 3097)
- RFC 2961, *RSVP Refresh Overhead Reduction Extensions*
- RFC 3097, *RSVP Cryptographic Authentication—Updated Message Type Value*
- RFC 3209, *RSVP-TE: Extensions to RSVP for LSP Tunnels*

The Null Service Object for maximum transmission unit (MTU) signaling in RSVP is not supported.

- RFC 3473, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReSerVation Protocol-Traffic Engineering (RSVP-TE) Extensions*

Only Section 9, “Fault Handling,” is supported.

- RFC 3477, *Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)*
- RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*
- RFC 4203, *OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)*

(OSPF extensions can carry traffic engineering information over unnumbered links.)

- RFC 4558, *Node-ID Based Resource Reservation Protocol (RSVP) Hello: A Clarification Statement*
- RFC 4561, *Definition of a Record Route Object (RRO) Node-Id Sub-Object*

The RRO node ID subobject is for use in inter-AS link and node protection configurations.

- RFC 4875, *Extensions to RSVP-TE for Point-to-Multipoint TE LSPs*

The following RFCs do not define standards, but provide information about RSVP and related technologies. The IETF classifies them variously as “Experimental” or “Informational.”

- RFC 2209, *Resource ReSerVation Protocol (RSVP)—Version 1 Message Processing Rules*
- RFC 2216, *Network Element Service Specification Template*
- RFC 4125, *Maximum Allocation Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering*
- RFC 4127, *Russian Dolls Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering*

Related Documentation

- Supported GMPLS Standards
- Supported LDP Standards
- Supported MPLS Standards
- Accessing Standards Documents on the Internet

RSVP Graceful Restart Standard

RSVP graceful restart is described in RFC 3473, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions* (only Section 9, “Fault Handling”).

RSVP Graceful Restart Terminology

R

Recovery time (in milliseconds)

Applies only when the control channel is up (the hello exchange is complete) before the restart time. Applies only to nodal faults.

When a graceful restart is in progress, the time left to complete a recovery is advertised. At other times, this value is zero. The maximum advertised recovery time is 2 minutes (120,000 milliseconds).

During the recovery time, a restarting node attempts to recover its lost states with assistance from its neighbors. The neighbor of the restarting node must send the path messages with the recovery labels to the restarting node within a period of one-half the recovery time. The restarting node considers its graceful restart complete after its advertised recovery time.

Restart time (in milliseconds) The default value is 60,000 milliseconds (1 minute). The restart time is advertised in the hello message. The time indicates how long a neighbor should wait to receive a hello message from a restarting router before declaring that router dead and purging states.

The Junos OS can override a neighbor's advertised restart time if the time is greater than one-third the local restart time. For example, given the default restart time of 60 seconds, a router would wait 20 seconds or less to receive a hello message from a restarting neighbor. If the restart time is zero, the restarting neighbor can immediately be declared dead.

CHAPTER 4

Summary of RSVP Configuration Statements

admin-group

Syntax	<pre>admin-group { exclude [<i>group-names</i>]; include-all [<i>group-names</i>]; include-any [<i>group-names</i>]; }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection], [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>], [edit protocols rsvp interface <i>interface-name</i> link-protection], [edit protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>]</pre>
Release Information	Statement introduced in Junos OS Release 9.2.
Description	Enable you to configure administrative groups for bypass label-switched paths (LSPs). You can configure administrative groups either globally for all bypass LSPs traversing an interface or for just a specific bypass LSP.
Options	<p>exclude <i>group-names</i>—Specify the administrative groups to exclude for a bypass LSP.</p> <p>include-all <i>group-names</i>—Specify the administrative groups whose links the bypass LSP must traverse.</p> <p>include-any <i>group-names</i>—Specify the administrative groups whose links the bypass LSP can traverse.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Administrative Groups for Bypass LSPs on page 34

aggregate (Protocols RSVP)

Syntax	(aggregate no-aggregate);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols rsvp peer-interface <i>peer-interface-name</i>], [edit protocols rsvp interface <i>interface-name</i>], [edit protocols rsvp peer-interface <i>peer-interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>Control the use of RSVP aggregate messages on an interface or peer interface:</p> <ul style="list-style-type: none">• aggregate—Use RSVP aggregate messages.• no-aggregate—Do not use RSVP aggregate messages. <p>Aggregate messages can pack multiple RSVP messages into a single transmission, thereby reducing network overhead and enhancing efficiency. The number of supportable sessions and processing overhead are significantly improved when aggregation is enabled.</p> <p>Not all routers connected to a subnet need to support aggregation simultaneously. Each RSVP router negotiates its intention to use aggregate messages on a per-neighbor basis. Only when both routers agree are aggregate messages sent.</p> <p>To have refresh reduction and reliable delivery, you must include the aggregate and reliable statements.</p>
Default	Aggregation is disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring RSVP Refresh Reduction on page 23• reliable on page 92

authentication-key (Protocols RSVP)

Syntax	authentication-key <i>key</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols rsvp peer-interface <i>peer-interface-name</i>], [edit protocols rsvp interface <i>interface-name</i>], [edit protocols rsvp peer-interface <i>peer-interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>Authentication key (password). Neighboring routers use the password to verify the authenticity of packets sent from this interface or peer interface.</p> <p>RSVP uses HMAC-MD5 authentication, which is defined in RFC 2104, <i>HMAC: Keyed-Hashing for Message Authentication</i>.</p> <p>All routers that are connected to the same IP subnet must use the same authentication scheme and password.</p>
Options	key —Authentication password. It can be 1 through 16 contiguous digits or letters. Separate decimal digits with periods. Separate hexadecimal digits with periods and precede the string with 0x. If you include spaces in the password, enclose the entire password in quotation marks (" ").
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring RSVP Authentication on page 26

bandwidth (Protocols RSVP)

Syntax	<code>bandwidth <i>bps</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i></code> <code>link-protection],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i></code> <code>link-protection bypass <i>bypass-name</i>],</code> <code>[edit protocols rsvp interface <i>interface-name</i>],</code> <code>[edit protocols rsvp interface <i>interface-name</i> link-protection],</code> <code>[edit protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>For certain logical interfaces (such as Asynchronous Transfer Mode [ATM], Permanent Virtual Circuit [PVC], or Frame Relay), you cannot determine the correct bandwidth from the hardware. This statement enables you to specify the actual available bandwidth.</p> <p>This statement also enables you to specify the bandwidth for a bypass label switched path (LSP). If you have configured multiple bypasses, this statement is mandatory and is applied to all of the bypass LSPs.</p>
Default	The hardware raw bandwidth is used.
Options	<p><i>bps</i>—Bandwidth in bits per second. You can specify this as an integer value. If you do so, count your zeros carefully, or you can use the abbreviations k (for a thousand), m (for a million), or g (for a billion [also called a thousand million]).</p> <p>Range: Any positive integer</p> <p>Default: 0 (no bandwidth is reserved)</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring the Bandwidth for Bypass LSPs on page 34• Configuring Link Protection on Interfaces Used by LSPs on page 32• Configuring Bypass LSPs on page 33

bypass (Signaled LSP)

Syntax	<pre> bypass <i>bypass-name</i> { bandwidth <i>bps</i>; description <i>text</i>; hop-limit <i>number</i>; no-cspf; path <i>address</i> <strict loose>; priority <i>setup-priority reservation-priority</i>; to address; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection],</p> <p>[edit protocols rsvp interface <i>interface-name</i> link-protection]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>The description option was added in Junos OS Release 10.4.</p>
Description	<p>Enables you to configure specific bandwidth and path constraints for a bypass LSP. It is possible to individually configure multiple bypass LSPs. If you do not configure the bypass LSPs individually, they all share the same path and bandwidth constraints.</p> <p>If you specify the bandwidth, hop-limit, and path statements for the bypass LSP, these values take precedence over the values configured at the [edit protocols rsvp interface <i>interface-name</i> link-protection] hierarchy level. The other attributes (subscription, no-node-protection, and optimize-timer) are inherited from the general constraints.</p>
Options	<p>bypass-name—(Required) Specify a name for the bypass LSP. The name can be up to 64 characters.</p> <p>description—Provides a textual description of the bypass LSP. Enclose any descriptive text that includes spaces in quotation marks (" "). Any descriptive text you include is displayed in the output of the show mpls lsp bypass detail command and has no effect on the operation of the bypass LSP. The description text can be no more than 80 characters in length.</p> <p>to address—(Required) Specify the address for the interface of the immediate next-hop node (for link protection) or the next-next-hop node (for node-link protection). The address specified determines whether this is a link protection bypass or a node-link protection bypass. On multiaccess networks (for example, a LAN), this address is also used to specify which next-hop node is being protected.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Bypass LSPs on page 33

bypass (Static LSP)

Syntax	<pre>bypass <i>bypass-name</i> { bandwidth <i>bps</i>; description <i>string</i>; next-hop (<i>address</i> <i>interface-name</i> <i>address/interface-name</i>); push <i>out-label</i>; to <i>address</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i>], [edit protocols mpls static-label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced before Junos OS Release 10.1.
Description	<p>Configure specific bandwidth and path constraints for a bypass ingress LSP. It is possible to configure multiple bypass LSPs individually. If you do not, they all share the same path and bandwidth constraints.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Configuring Static LSPs

class-of-service (Protocols RSVP)

Syntax	<code>class-of-service <i>cos-value</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>],</p> <p>[edit protocols rsvp interface <i>interface-name</i> link-protection],</p> <p>[edit protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>]</p>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>Class-of-service (CoS) value given to all packets in the bypass LSP. You can specify a single CoS value for all the bypass LSPs traversing an interface. You can also configure CoS values for specific bypass LSPs traversing an interface.</p> <p>The CoS value might affect the scheduling or queuing algorithm of traffic traveling along an LSP.</p>
Options	<p><i>cos-value</i>—CoS value. A higher value typically corresponds to a higher level of service.</p> <p>Range: 0 through 7</p> <p>Default: If you do not specify a CoS value, the IP precedence bits from the packet's IP header are used as the packet's CoS value.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Class of Service for Bypass LSPs on page 35

destination-networks

Syntax	<code>destination-networks <i>prefix</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options dynamic-tunnels <i>tunnel-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-options dynamic-tunnels <i>tunnel-name</i> rsvp-te <i>entry</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-options dynamic-tunnels <i>tunnel-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> routing-options dynamic-tunnels <i>tunnel-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> routing-options dynamic-tunnels <i>tunnel-name</i> rsvp-te <i>entry</i>],</code> <code>[edit routing-options dynamic-tunnels <i>tunnel-name</i>],</code> <code>[edit routing-options dynamic-tunnels <i>tunnel-name</i> rsvp-te <i>entry</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Specify the IPv4 prefix range for the destination network. Only tunnels within the specified IPv4 prefix range can be created.
Options	<i>prefix</i> —Destination prefix of the network.
Required Privilege Level	<code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Configuring GRE Tunnels for Layer 3 VPNsConfiguring Dynamic TunnelsConfiguring RSVP Automatic Mesh on page 42

devices

Syntax	<code>devices <i>device-names</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
Release Information	Statement introduced in Junos OS Release 8.1.
Description	Specifies one of the virtual tunnel (VT) interfaces to de-encapsulate the egress traffic for ultimate-hop popping on point-to-multipoint LSPs. If no device is specified, the selection process is performed automatically.
Default	The device selection process is performed automatically if no device is configured. Junos OS selects one of the available VT interfaces to de-encapsulate the egress traffic.
Options	<i>device-names</i> —Specify which VT interfaces are used to handle the RSVP traffic. Range: 0 to 8 devices
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Enabling Ultimate-Hop Popping on Point-to-Multipoint LSPs on page 50 • Understanding Redundant Virtual Tunnel Interfaces in MBGP MVPNs

disable (Protocols RSVP)

Syntax	disable;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit logical-systems <i>logical-system-name</i> protocols rsvp graceful-restart], [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection], [edit logical-systems <i>logical-system-name</i> protocols rsvp peer-interface <i>peer-interface-name</i>], [edit protocols rsvp], [edit protocols rsvp graceful-restart], [edit protocols rsvp interface <i>interface-name</i>], [edit protocols rsvp interface <i>interface-name</i> link-protection], [edit protocols rsvp peer-interface <i>peer-interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Explicitly disable RSVP or RSVP graceful restart. Explicitly disable link protection on the specified interface.
Default	RSVP is enabled on interfaces and peer interfaces configured with the RSVP interface statement. RSVP graceful restart is enabled on the router. Link protection is disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Minimum RSVP Configuration on page 21• Configuring RSVP Graceful Restart on page 40• Configuring Link Protection on Interfaces Used by LSPs on page 32

fast-reroute (Protocols RSVP)

Syntax	<code>fast-reroute optimize-timer <i>seconds</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
Release Information	Statement added in Junos OS Release 7.5.
Description	Configure the optimize timer for fast reroute. The optimize timer triggers a periodic optimization process that recomputes the fast reroute detour LSPs to use network resources more efficiently.
Options	<i>seconds</i> —Specify the number of seconds between fast reroute detour LSP optimizations. Range: 0 through 65,535 seconds Default: 0 (disabled)
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> Configuring the Optimization Interval for Fast Reroute Paths

graceful-deletion-timeout

Syntax	<code>graceful-deletion-timeout <i>seconds</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the time, in seconds, before completing graceful deletion of signaling.
Options	<i>seconds</i> —Time before completing graceful deletion of signaling. Range: 1 through 300 seconds Default: 30 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> Gracefully Tearing Down GMPLS LSPs

graceful-restart (Protocols RSVP)

Syntax	<pre>graceful-restart { disable; helper-disable; maximum-helper-recovery-time <i>seconds</i>; maximum-helper-restart-time <i>seconds</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options], [edit protocols rsvp], [edit routing-options]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Enable graceful restart on the router. You must configure the graceful-restart statement at the [edit routing-options] hierarchy level to enable graceful restart on the router.
Options	<p>disable—Disable graceful restart on the router or for RSVP.</p> <p>helper-disable—Disable RSVP graceful restart helper mode (this option is only available at the [edit protocols rsvp] hierarchy level).</p> <p>Default: Helper mode is enabled by default.</p> <p>maximum-helper-recovery-time <i>seconds</i>—The maximum length of time the router stores the state of neighboring routers when they undergo a graceful restart. The value applies to all neighboring routers, so it should be based on the time that the slowest RSVP neighbor requires for restart.</p> <p>Default: 180 seconds</p> <p>Range: 1 through 3600 seconds</p> <p>maximum-helper-restart-time <i>seconds</i>—The maximum length of time the router waits between when it discovers that a neighboring router has gone down and when it declares the neighbor down. This value is applied to all neighboring routers, so it should be based on the time that the slowest RSVP neighbor requires for restart.</p> <p>Default: 20 seconds</p> <p>Range: 1 through 1800 seconds</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring RSVP Graceful Restart on page 40

hello-acknowledgements

Syntax	hello-acknowledgements;
Hierarchy Level	[edit logical-systems <i>logical-systems-name</i> protocols rsvp], [edit protocols rsvp]
Release Information	Statement introduced in Junos OS Release 10.2.
Description	Enable hello messages from nonsession neighbors to be acknowledged with a hello acknowledgment message. Once hello acknowledgments are enabled, the router continues to acknowledge hello messages from any nonsession RSVP neighbors unless the interface itself goes down or the configuration is changed by an administrator.
Default	Hello acknowledgments are disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Hello Acknowledgments for Nonsession RSVP Neighbors on page 29

hello-interval (Protocols RSVP)

Syntax	hello-interval <i>seconds</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols rsvp peer-interface <i>peer-interface-name</i>], [edit protocols rsvp interface <i>interface-name</i>], [edit protocols rsvp peer-interface <i>peer-interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Enable the sending of hello packets on the interface.
Options	<p><i>seconds</i>—Length of time between hello packets. A value of 0 disables the sending of hello packets on the interface.</p> <p>Range: 1 through 60 seconds</p> <p>Default: 9 seconds</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the RSVP Hello Interval on page 25

hop-limit

Syntax	<code>hop-limit <i>number</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> fast-reroute], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>], [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection], [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>], [edit protocols mpls], [edit protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i> fast-reroute], [edit protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>], [edit protocols rsvp interface <i>interface-name</i> link-protection], [edit protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>]</p>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>Specify the maximum number of routers that an LSP can traverse. This limit can be applied to any of the following:</p> <ul style="list-style-type: none"> • LSPs—The configured hop limit includes the ingress and egress routers. You can specify a hop limit for an LSP and for both primary and secondary paths. • Fast reroute detour—Specify the number of additional routers a fast reroute detour can traverse relative to the protected LSP. For example, if an LSP traverses 4 routers, any detour for the LSP can be no more than 10 router hops, including the ingress and egress routers. • Link protection bypass—Specify the maximum number of routers that a link protection bypass can traverse.
Options	<p><i>number</i>—Maximum number of hops.</p> <p>Range: 2 through 255 (for an LSP or for a link protection bypass); 0 through 255 (for fast reroute)</p> <p>Default: 255 (for an LSP or for a link protection bypass); 6 (for fast reroute)</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Fast Reroute • Limiting the Number of Hops in LSPs • Configuring the Hop Limit for Bypass LSPs on page 35

interface (Protocols RSVP)

```

Syntax  interface interface-name {
        disable;
        (aggregate | no-aggregate);
        authentication-key key;
        bandwidth bps;
        hello-interval seconds;
        link-protection {
            disable;
            admin-group {
                exclude [ group-names ];
                include-all [ group-names ];
                include-any [ group-names ];
            }
            bandwidth bps;
            bypass bypass-name {
                bandwidth bps {
                    ct0 bps;
                    ct1 bps;
                    ct2 bps;
                    ct3 bps;
                }
                description text;
                class-of-service cos-value;
                hop-limit number;
                no-cspf;
                path address <strict | loose>;
                priority setup-priority reservation-priority;
                to address;
            }
            class-of-service cos-value;
            hop-limit number;
            max-bypasses number;
            no-cspf;
            no-node-protection;
            optimize-timer seconds;
            path address <strict | loose>;
            priority setup-priority reservation-priority;
            subscription percentage;
        }
        (reliable | no-reliable);
        subscription percentage {
            ct0 percentage;
            ct1 percentage;
            ct2 percentage;
            ct3 percentage;
        }
        update-threshold threshold;
    }

```

Hierarchy Level [edit logical-systems *logical-system-name* protocols rsvp],
[edit protocols rsvp]

Release Information	Statement introduced before Junos OS Release 7.4.
Description	Enable RSVP on one or more router interfaces.
Default	RSVP is disabled on all interfaces.
Options	<i>interface-name</i> —Name of an interface. To configure all interfaces, specify all . For details about specifying interfaces, see the Junos® OS Network Interfaces. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Minimum RSVP Configuration on page 21

keep-multiplier

Syntax	keep-multiplier <i>number</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Set the keep multiplier value.
Options	<i>number</i> —Multiplier value. Range: 1 through 255 Default: 3
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Timers for RSVP Refresh Messages on page 44

label-switched-path-template

Syntax	label-switched-path-template { (default-template <i>lsp-template-name</i>); }
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel rsvp-te],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel ingress-replication label-switched-path],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>address</i> source <i>source-address</i> rsvp-te],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options dynamic-tunnels <i>tunnel-name</i> rsvp-te <i>entry-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> provider-tunnel ingress-replication label-switched-path],</p> <p>[edit routing-instances <i>routing-instance-name</i> provider-tunnel rsvp-te],</p> <p>[edit routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>address</i> source <i>source-address</i> rsvp-te],</p> <p>[edit routing-options dynamic-tunnels <i>tunnel-name</i> rsvp-te <i>entry-name</i>]</p>
Release Information	Statement introduced in Junos OS Release 8.5.
Description	Specify the LSP template. An LSP template is used as the basis for other dynamically generated LSPs. This feature can be used for a number of applications, including point-to-multipoint LSPs, flooding VPLS traffic, configuring ingress replication for IP multicast using MBGP MVPNs, and to enable RSVP automatic mesh. There is no default setting for the label-switched-path-template statement, so you must configure either the default-template using the default-template option, or you must specify the name of your preconfigured LSP template.
Options	<p>default-template—Specify that the default LSP template be used for the dynamically generated LSPs.</p> <p><i>lsp-template-name</i>—Specify the name of an LSP to be used as a template for the dynamically generated LSPs.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Ingress Replication for IP Multicast Using MBGP MVPNs • Configuring Point-to-Multipoint LSPs for an MBGP MVPN • Flooding Unknown Traffic Using Point-to-Multipoint LSPs • Configuring RSVP Automatic Mesh on page 42

link-protection (RSVP)

Syntax	<pre> link-protection { disable; admin-group { exclude [group-names]; include-all [group-names]; include-any [group-names]; } bandwidth bps; bypass bypass-name { bandwidth bps { ct0 bps; ct1 bps; ct2 bps; ct3 bps; } description text; class-of-service cos-value; hop-limit number; no-cspf; path address <strict loose>; priority setup-priority reservation-priority; to address; } class-of-service cos-value; hop-limit number; max-bypasses number; no-cspf; no-node-protection; optimize-timer seconds; path address <strict loose>; priority setup-priority reservation-priority; subscription percentage; } </pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i>], [edit protocols rsvp interface <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Enable link protection on the specified interface. Using link protection, you can configure a network to reroute traffic quickly around broken links. To fully enable link protection, you also need to configure the link-protection statement at the [edit protocols mpls label-switched-path <i>lsp-name</i>] hierarchy level. You can configure single or multiple bypasses for protected interface.
Default	Link protection is disabled.
Options	no-node-protection —Disable node-link protection on the RSVP interface. Link protection remains active. When this option is configured, the router can only initiate a next-hop bypass, not a next-next-hop bypass.

The remaining statements are explained separately.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Link Protection on Interfaces Used by LSPs on page 32• link-protection (Dynamic LSPs)

load-balance (Protocols RSVP)

Syntax	load-balance { bandwidth; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Load-balance traffic between RSVP LSPs.
Options	bandwidth —Load-balance traffic between RSVP LSPs based on the bandwidth configured for each LSP.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Load Balancing Across RSVP LSPs on page 41

max-bypasses

Syntax	<code>max-bypasses <i>number</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i>], [edit protocols rsvp interface <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Range modified in Junos OS Release 9.3.
Description	Specify the maximum number of dynamic bypass LSPs permitted for protecting this interface. When this option is configured, multiple bypasses for link protection are enabled. Call admission control (CAC) is also enabled. The limit on bypasses configured applies only to dynamically generated bypass LSPs. By default, this option is disabled and only one dynamic bypass LSP is enabled for each interface. If you configure max-bypasses , you must also configure the bandwidth statement.
Options	number —Configure the maximum number of bypass LSPs. If you configure a value of 0, no dynamic bypass LSPs are allowed to be established for the interface. Only static bypass LSPs can be configured. Range: 0 through 99 Default: 1
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Maximum Number of Bypass LSPs on page 36

no-local-reversion

Syntax	no-local-reversion;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
Release Information	Statement introduced in Junos OS Release 10.4.
Description	<p>Disables RSVP local revertive mode as specified in RFC 4090, <i>Fast Reroute Extensions to RSVP-TE for LSP Tunnels</i>. RSVP local revertive mode is supported on all Juniper Networks routers running the Junos OS. It is the default behavior. If you include this statement, the Juniper Networks router uses global revertive mode instead. You might need to disable RSVP local revertive mode on Juniper Networks routers if your network includes equipment that does not support this mode.</p> <p>The following information can also be found in RFC 4090. Refer to the full RFC for additional information. When an LSP fails, the connection can be repaired locally using a traffic protection mechanism such as fast reroute. To restore the LSP to a full working path, RFC 4090 specifies the following strategies:</p> <ul style="list-style-type: none"> • Local revertive mode—Upon detecting that the path is restored, the point of local repair (PLR) resignals each of the LSPs that were formerly routed over the restored path. Every LSP successfully resingaled along the restored path is switched back. • Global revertive mode—The ingress router of each tunnel is responsible for reoptimizing the LSPs that used the failed path. There are several potential reoptimization triggers: RSVP error messages, inspection of OSPF LSAs or IS-IS LSPs, and timers. This re-optimization process can proceed as soon as the failure is detected. It is not tied to the restoration of the failed path.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

node-hello

Syntax	node-hello;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
Release Information	Statement introduced in JUNOS Release 10.0.
Description	Enables node-ID based RSVP hellos globally on all of the RSVP interfaces on the router to allow Juniper Networks routers to interoperate with the equipment of other vendors. By default, the JUNOS Software uses interface-based RSVP hellos and node-ID based RSVP hellos are disabled. If you have not enabled RSVP node IDs on the router, the JUNOS software does not accept any node-ID hello packets.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring RSVP Node ID Hellos on page 28

no-adjacency-down-notification (Protocols IS-IS)

Syntax	no-adjacency-down-notification;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i>], [edit protocols isis interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 8.0.
Description	<p>Disable adjacency down notification for IS-IS to allow for migration from IS-IS to OSPF without disruption of the RSVP neighbors and associated RSVP-signaled label-switched paths (LSPs).</p> <p>Whenever IS-IS is deactivated, the IS-IS adjacencies are brought down. IS-IS signals to RSVP to bring down any RSVP neighbors associated with the IS-IS adjacencies, and this further causes the associated LSPs signaled by RSVP to go down as well.</p> <p>A similar process occurs whenever OSPF is deactivated. The OSPF neighbors are brought down. OSPF signals to RSVP to bring down any of the RSVP neighbors associated with the OSPF neighbors, and this further causes the associated LSPs signaled by RSVP to go down as well.</p> <p>If you need to migrate from IS-IS to OSPF or from OSPF to IS-IS, the internal gateway protocol (IGP) notification to RSVP for an adjacency or neighbor down event needs to be ignored. Using the no-adjacency-down-notification or no-neighbor-down-notification statements, you can disable IS-IS adjacency down notification or OSPF neighbor down notification, respectively, until the migration is complete. The network administrator is responsible for configuring the statements before the migration, and then removing them from the configuration afterward, so that IGP notification can function normally.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • no-neighbor-down-notification on page 85

no-cspf (Protocols RSVP)

Syntax	no-cspf;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection], [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>], [edit protocols rsvp interface <i>interface-name</i> link-protection], [edit protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>]
Release Information	Statement introduced in Junos OS Release 7.5.
Description	Disable CSPF computation on all bypass LSPs or on a specific bypass LSP. You need to disable CSPF for link protection to function properly on interarea paths.
Default	CSPF is enabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Disabling CSPF for Bypass LSPs on page 36

no-interface-hello

Syntax	no-interface-hello;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
Release Information	Statement introduced in JUNOS Release 10.0.
Description	Allows you to explicitly disable RSVP interface hellos globally on the router. This type of configuration might be necessary in networks where the Juniper Networks router has numerous RSVP connections with equipment from other vendors. However, if you disable RSVP interface hellos globally, you can also configure a hello interval on an RSVP interface using the hello-interval (Protocols RSVP) statement. This configuration disables RSVP interface hellos globally but enables RSVP interface hellos on the specified interface. This configuration might be necessary in a heterogeneous network where some devices support RSVP node ID hellos and other devices support RSVP interface hellos.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring RSVP Node ID Hellos on page 28• hello-interval (Protocols RSVP) on page 73

no-neighbor-down-notification

Syntax	no-neighbor-down-notification;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i>], [edit protocols ospf area <i>area-id</i> interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 8.0.
Description	Disable neighbor down notification for OSPF to allow for migration from OSPF to IS-IS without disruption of the RSVP neighbors and associated RSVP-signaled LSPs.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

no-node-id-subobject

Syntax	no-node-id-subobject;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
Release Information	Statement introduced in Junos OS Release 9.4.
Description	Disable the record route object (RRO) node ID subobject for compatibility with earlier versions of the Junos OS. To interoperate with other vendors' equipment, the Junos OS supports the RRO node ID subobject for use in inter-AS link and node protection configurations.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Inter-AS Node and Link Protection on page 31

no-p2mp-sublsp

Syntax	no-p2mp-sublsp;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
Release Information	Statement introduced in Junos OS Release 9.2.
Description	Reject Resv messages that include the S2L_SUB_LSP object. By default, Resv messages that include the S2L_SUB_LSP object are accepted. However, in a network which includes Juniper Networks devices running both Junos OS Release 9.2 and later and Junos OS Release 9.1 and earlier, it is necessary to configure the no-p2mp-sublsp statement on devices running Junos OS Release 9.2 and later to ensure that point-to-multipoint LSPs function properly.
Default	Resv messages that include the S2L_SUB_LSP object are accepted.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Preserving Point-to-Multipoint LSP Functioning with Different Junos OS Releases

node-link-protection (Protocols MPLS)

Syntax	node-link-protection;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Enable node and link protection on the specified LSP. To fully enable node and link protection, you also need to include the link-protection statement at the [edit protocols rsvp interface <i>interface-name</i>] hierarchy level.
Default	Node and link protection is disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Configuring Node Protection or Link Protection for LSPs on page 29

optimize-timer (Protocols RSVP)

Syntax	<code>optimize-timer <i>seconds</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection], [edit protocols rsvp interface <i>interface-name</i> link-protection]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure an optimize timer for a bypass LSP. The optimize timer initiates a periodic optimization process that reshuffles data LSPs among bypass LSPs to achieve the most efficient use of network resources. The optimization process attempts to either minimize the number of bypasses currently in use, minimize the total amount of bandwidth reserved for all bypasses, or both.
Options	<i>seconds</i> —Specify the number of seconds between optimizations. Range: 0 through 65,535 seconds Default: 0 (disabled)
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Optimization Interval for Bypass LSPs on page 37

path (Protocols RSVP)

Syntax	<code>path address <strict loose>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection], [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>], [edit protocols rsvp interface <i>interface-name</i> link-protection], [edit protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure an explicit path (a sequence of strict or loose routes) to control where and how a bypass LSP is established. If multiple bypasses are configured, they all will use the same explicit path.
Default	No path is configured. CSPF automatically calculates the path the bypass LSP takes.
Options	<p>address—IP address of each transit router in the LSP. You must specify the address or hostname of each transit router, although you do not need to list each transit router if its type is loose. As an option, you can include the ingress and egress routers in the path. Specify the addresses in order, starting with the ingress router (optional) or the first transit router, and continuing sequentially along the path until reaching the egress router (optional) or the router immediately before the egress router.</p> <p>Default: If you do not specify any routers explicitly, no routing limitations are imposed on the bypass LSP.</p> <p>loose—(Optional) The next address in the path statement is loose. The LSP can traverse other routers before reaching this router.</p> <p>Default: strict</p> <p>strict—(Optional) The LSP must go to the next address specified in the path statement without traversing other nodes. This is the default.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring an Explicit Path for Bypass LSPs on page 37

peer-interface (Protocols RSVP)

Syntax	<pre>peer-interface <i>peer-interface-name</i> { disable; (aggregate no-aggregate); authentication-key <i>key</i>; hello-interval <i>seconds</i>; (reliable no-reliable); }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure the name of the LMP peer device. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Configuring RSVP and OSPF for LMP Peer Interfaces

preemption

Syntax	<pre>preemption { (aggressive disabled normal); soft-preemption { cleanup-timer <i>seconds</i>; } }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Control RSVP session preemption.
Default	normal
Options	<p>aggressive—Preempt RSVP sessions whenever bandwidth is insufficient to handle all sessions. A session is preempted whenever bandwidth is lowered or a new higher-priority session is established.</p> <p>disabled—Do not preempt RSVP sessions.</p> <p>normal—Preempt RSVP sessions to accommodate new higher-priority sessions when bandwidth is insufficient to handle all sessions.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Preempting RSVP Sessions on page 45

priority (Protocols RSVP)

Syntax	<code>priority setup-priority reservation-priority;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>],</p> <p>[edit protocols rsvp interface <i>interface-name</i> link-protection],</p> <p>[edit protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>],</p>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure the setup priority and reservation priority for a bypass LSP. If insufficient link bandwidth is available during session establishment, the setup priority is compared with other setup priorities for established sessions on the link to determine whether some of them should be preempted to accommodate the new session. The session with the lower-hold priority is preempted.
Options	<p>reservation-priority—Reservation priority, used to keep a reservation after it has been set up. A smaller number has a higher priority. The priority must be greater than or equal to the setup priority to prevent preemption loops.</p> <p>Range: 0 through 7, where 0 is the highest and 7 is the lowest priority.</p> <p>Default: 0 (Once the session is set up, no other session can preempt it.)</p> <p>setup-priority—Setup priority.</p> <p>Range: 0 through 7, where 0 is the highest and 7 is the lowest priority.</p> <p>Default: 7 (The session cannot preempt any existing sessions.)</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Priority and Preemption for Bypass LSPs on page 38 • Configuring Priority and Preemption for LSPs

refresh-time

Syntax	refresh-time <i>seconds</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Set the refresh time.
Options	<i>seconds</i> —Refresh time. Range: 1 through 65,535 Default: 30 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Timers for RSVP Refresh Messages on page 44

reliable

Syntax	(reliable no-reliable);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols rsvp peer-interface <i>peer-interface-name</i>], [edit protocols rsvp interface <i>interface-name</i>], [edit protocols rsvp peer-interface <i>peer-interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Enable reliable message delivery on the interface. In order to have refresh reduction and reliable delivery, you must include the aggregate and reliable statements.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring RSVP Refresh Reduction on page 23• aggregate on page 62

rsvp

Syntax	rsvp;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols], [edit protocols]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.5 for EX Series switches.
Description	<p>Enable Resource Reservation Protocol (RSVP) signaling.</p> <p>You must include the rsvp statement in the configuration to enable RSVP on the router.</p> <p>The primary purpose of RSVP in Junos OS for EX Series switches is to support dynamic signaling within label switched paths (LSPs).</p>
Default	RSVP is disabled.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Minimum RSVP Configuration on page 21 • Example: Configuring MPLS on EX Series Switches • Configuring MPLS on Provider Edge Switches Using Circuit Cross-Connect (CLI Procedure) • Configuring MPLS on Provider Edge Switches Using IP Over MPLS (CLI Procedure) • Configuring MPLS on Provider Switches (CLI Procedure) • Junos OS MPLS Applications Configuration Guide

rsvp-te (Routing Options)

Syntax	<pre>rsvp-te entry-name { destination-networks network-prefix; label-switched-path-template { default-template; template-name; } }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options dynamic-tunnels <i>tunnel-name</i>], [edit routing-options dynamic-tunnels <i>tunnel-name</i>]
Release Information	Statement introduced in Junos OS Release 10.1. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Enable RSVP to automatically establish LSPs for any new PE router added to a full mesh of LSPs. To enable this feature, you must configure the rsvp-te statement on all of the PE routers in the full mesh.
Options	entry-name —Specify the entry for the RSVP tunnel. The other options are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring RSVP Automatic Mesh on page 42• Flooding Unknown Traffic Using Point-to-Multipoint LSPs

setup-protection

Syntax	setup-protection;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
Description	The facility-backup fast reroute mechanism can provide setup protection for LSPs which are in the process of being signaled. Both point-to-point LSPs and point-to-multipoint LSPs are supported. You should configure the setup-protection statement on each of the routers along the LSP path on which you want to enable LSP setup protection. You should also configure IGP traffic engineering on all of the routers on the LSP path.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring RSVP Setup Protection on page 39

soft-preemption (Protocols RSVP)

Syntax	<code>soft-preemption { cleanup-timer <i>seconds</i>; }</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp preemption], [edit protocols rsvp preemption]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Enable soft preemption to attempt to establish a new path for a preempted LSP before tearing it down.
Options	cleanup-timer —A value of 0 disables soft preemption. Range: 0 through 180 seconds Default: 30 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Configuring MPLS Soft Preemption

static-label-switched-path

Syntax static-label-switched-path *lsp-name* {
 bypass *bypass-name* {
 bandwidth *bps*;
 description *string*;
 next-hop (*address* | *interface-name* | *address/interface-name*);
 push *out-label*;
 to *address*;
 }
 ingress {
 bandwidth *bps*;
 class-of-service *cos-value*;
 description *string*;
 install {
 destination-prefix <active>;
 }
 link-protection *bypass-name name*;
 metric *metric*;
 next-hop (*address* | *interface-name* | *address/interface-name*);
 node-protection *bypass-name name* next-next-label *label*;
 no-install-to-address;
 policing {
 filter *filter-name*;
 no-auto-policing;
 }
 preference *preference*;
 push *out-label*;
 to *address*;
 }
 transit *incoming-label* {
 bandwidth *bps*;
 description *string*;
 link-protection *bypass-name name*;
 next-hop (*address* | *interface-name* | *address/interface-name*);
 node-protection *bypass-name name* next-next-label *label*;
 pop;
 swap *out-label*;
 }
 }

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls],
 [edit protocols mpls]

Release Information Statement introduced in Junos OS Release 10.1.

Description Configure a static LSP.

Options *lsp-name*—Name of the path.

The remaining statements are explained separately.

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

Related Documentation

- Configuring Static LSPs

subscription

Syntax	<pre>subscription <i>percentage</i> { ct0 <i>percentage</i>; ct1 <i>percentage</i>; ct2 <i>percentage</i>; ct3 <i>percentage</i>; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection], [edit protocols rsvp interface <i>interface-name</i>], [edit protocols rsvp interface <i>interface-name</i> link-protection]</p>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>Configure the amount of bandwidth subscribed to a class type (when you have enabled Differentiated Services) or bypass LSP (when you have enabled link protection). subscription is the percentage of the link bandwidth that can be used for the RSVP reservation process.</p>
Options	<p>ctnumber percentage—Percentage of the class-type bandwidth allowed for reservations. If you specify a value greater than 100, you are oversubscribing the class type. You can specify bandwidth subscriptions for class types 0 through 3. This option is not available for bypass LSPs.</p> <p>Range: 0 through 65,000 Default: 100 percent</p> <p>percentage—Percentage of the class-type or bypass LSP bandwidth allowed for reservations. If you specify a value greater than 100, you are oversubscribing the class type or bypass LSP.</p> <p>Range: 0 through 65,000 Default: 100 percent</p>
Required Privilege Level	<p>routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring the Bandwidth Subscription Percentage for LSPs • Configuring the Amount of Bandwidth Subscribed for Bypass LSPs on page 38

traceoptions (Protocols RSVP)

Syntax	<pre>traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier> <disable>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Enable RSVP-level trace options.
Default	The default RSVP-level trace options are those inherited from the routing protocols traceoptions statement included at the [edit routing-options] hierarchy level.
Options	<p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>filename—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. We recommend that you place RSVP tracing output in the file rsvp-log.</p> <p>files number—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>Range: 2 through 1000</p> <p>Default: 2 files</p> <p>If you specify a maximum number of files, you must also include the size statement to specify the maximum file size.</p> <p>flag—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements.</p> <ul style="list-style-type: none">• all—All tracing operations• error—All detected error conditions• event—RSVP-related events• lmp—RSVP-LMP interactions• packets—All RSVP packets• path—All path messages• pathtear—PathTear messages• resv—Resv messages

- **resvtear**—ResvTear messages
- **route**—Routing information
- **state**—Session state transitions

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

no-world-readable—(Optional) Enable only certain users to read the log file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches this size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Syntax: *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 1 MB

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of files.

world-readable—(Optional) Enable any user to read the log file.

Required Privilege Level	routing and trace—To view this statement in the configuration. routing-control and trace-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Tracing RSVP Protocol Traffic on page 52

transit

Syntax	<pre>transit <i>incoming-label</i> { bandwidth <i>bps</i>; description <i>string</i>; link-protection bypass-name <i>name</i>; next-hop (<i>address</i> <i>interface-name</i> <i>address/interface-name</i>); node-protection bypass-name <i>name</i> next-next-label <i>label</i>; pop; swap <i>out-label</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i>], [edit protocols mpls static-label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced in Junos OS Release 10.1.
Description	Configure a transit static LSP. The remaining statements are explained separately.
Options	<i>incoming-label</i> —Incoming label value. Range: 1000000 through 1048575
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Configuring Static LSPs

tunnel-services (RSVP)

Syntax	tunnel-services { devices <i>device-names</i> ; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
Release Information	Statement introduced in Junos OS Release 8.1.
Description	Enable ultimate-hop popping on point-to-multipoint LSPs. The Junos OS selects one of the available virtual tunnel (VT) interfaces to de-encapsulate the egress traffic. By default, the selection process is performed automatically.
Default	Ultimate-hop popping is disabled.
Options	devices <i>device-names</i> —Specify which VT interfaces are used to handle the RSVP traffic. Range: 0 to 8 devices
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Enabling Ultimate-Hop Popping on Point-to-Multipoint LSPs on page 50

ultimate-hop-popping

Syntax	ultimate-hop-popping;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>label-switched-path-name</i>], [edit protocols mpls], [edit protocols mpls label-switched-path <i>label-switched-path-name</i>]
Release Information	Statement introduced in Junos OS Release 12.3.
Description	<p>Enable ultimate-hop popping on LSPs. Configure this statement on the device at the LSP ingress. In ultimate-hop popping, the MPLS label is popped from the IP packet at the PE router. The IP address is checked in a second address lookup (also at the PE router), and then the packet is forwarded to its destination.</p> <p>Be aware of the following platform requirements and restrictions:</p> <ul style="list-style-type: none">• UHP LSPs using VT interfaces—Supported on all M Series, MX Series, T Series, and TX Matrix routers.• UHP LSPs using LSI interfaces—Supported on MX 3D Series routers only.• UHP LSP requirements for the egress PE device—For M Series and T Series routers, a VT interface is needed.• UHP LSPs and Layer 3 VPNs—UHP LSPs are supported for Layer 3 VPNs configured on MX 3D Series routers only.• UHP LSPs and VPLS—UHP LSPs are supported for VPLS configured on MX 3D Series routers only. You must configure the no-tunnel-services statement at the [edit routing-instances <i>routing-instance-name</i> protocols vpls] hierarchy level.
Default	Ultimate-hop popping is disabled by default on LSPs. Penultimate-hop popping is the default behavior. In penultimate-hop popping, the final MPLS label is popped from the IP packet at the last provider router in the network before being forwarded to the PE router. The PE router receives the packet and checks the IP address, and then the packet is forwarded to its destination.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Ultimate-Hop Popping for LSPs on page 46• explicit-null

update-threshold

Syntax	update-threshold <i>threshold</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i>], [edit protocols rsvp interface <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Adjust the threshold at which a change in bandwidth triggers an interior gateway protocol (IGP) update.
Options	threshold —Specify the percentage change in bandwidth to trigger an IGP update. Range: 1 through 20 percent Default: 10 percent
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the RSVP Update Threshold on an Interface on page 27

PART 4

Index

- [Index on page 107](#)

Index

Symbols

#, comments in configuration statements.....	xiv
(), in syntax descriptions.....	xiv
< >, in syntax descriptions.....	xiv
[], in configuration statements.....	xiv
{ }, in configuration statements.....	xiv
(pipe), in syntax descriptions.....	xiv

A

admin-group statement	
bypass LSPs.....	61
configuration	
bypass LSPs.....	34
aggregate statement	
RSVP.....	62
usage guidelines.....	23
aggregation, RSVP.....	62
all (tracing flag)	
RSVP.....	98
allow-fragmentation statement	
usage guidelines.....	45
always-mark-connection-protection-tlv statement	
usage guidelines.....	30
authentication	
RSVP.....	26, 63
authentication-key statement	
RSVP.....	63
usage guidelines.....	26
automatic mesh, RSVP.....	42
automatic reoptimization, bypass LSPs.....	37

B

bandwidth	
RSVP reservations.....	97
bandwidth statement	
link protection.....	64
usage guidelines.....	34
RSVP.....	64
bandwidth update threshold.....	27
braces, in configuration statements.....	xiv

brackets

angle, in syntax descriptions.....	xiv
square, in configuration statements.....	xiv
bypass LSPs.....	33
administrative groups.....	34
bandwidth.....	34
bandwidth subscription.....	38
class-of-service.....	35
CSPF, disabling.....	36
explicit paths.....	37
hop limit.....	35
maximum number.....	36
multiple.....	14
node protection, disabling.....	36
optimization interval.....	37
priority and preemption.....	38
switching away from a network node.....	30
types.....	14
bypass statement	
RSVP.....	65
static LSP.....	66
usage guidelines.....	33

C

class-of-service statement	
bypass LSPs.....	67
usage guidelines.....	35
comments, in configuration statements.....	xiv
conventions	
text and syntax.....	xiii
CoS requests using RSVP.....	4
curly braces, in configuration statements.....	xiv
customer support.....	xv
contacting JTAC.....	xv

D

destination-networks statement	
tunnel.....	68
detail (tracing flag modifier)	
RSVP.....	99
devices statement.....	69
disable option to traceoptions statement	
RSVP.....	98

disable statement		
link protection.....	70	
usage guidelines.....	29	
RSVP.....	70	
usage guidelines.....	21	
RSVP graceful restart.....	70	
usage guidelines.....	40	
distinct reservations.....	8	
documentation		
comments on.....	xv	
dynamic tunnels		
destination.....	68	
E		
error (tracing flag)		
RSVP.....	98	
event (tracing flag)		
RSVP.....	98	
explicit senders, RSVP.....	8	
explicit-null statement		
RSVP		
usage guidelines.....	50	
F		
facility backup.....	13	
fast reroute.....	13	
PFE fast reroute.....	41	
fast-reroute statement		
RSVP.....	71	
FF (reservation style).....	8	
fixed-filter reservation style.....	8	
font conventions.....	xiii	
G		
graceful restart		
RSVP.....	72	
graceful-deletion-timeout statement.....	71	
graceful-restart statement		
RSVP.....	72	
usage guidelines.....	40	
H		
hello acknowledgments		
RSVP.....	29	
hello interval		
RSVP.....	25, 73	
hello messages		
interface.....	25	
node ID.....	28	
hello packets		
RSVP.....	5	
hello-acknowledgements statement.....	73	
RSVP		
usage guidelines.....	29	
hello-interval statement		
RSVP.....	73	
usage guidelines.....	25	
helper-disable statement		
RSVP		
usage guidelines.....	40	
hop-limit statement.....	74	
usage guidelines.....	35	
I		
Integrity object.....	4	
interface statement		
RSVP.....	75	
usage guidelines.....	21	
K		
keep multiplier, RSVP.....	44, 76	
keep-multiplier statement.....	76	
usage guidelines.....	44	
L		
label-switched-path statement		
MPLS with RSVP		
usage guidelines.....	22	
label-switched-path-template statement.....	77	
LDP		
Explicit Null label.....	50	
Implicit Null label.....	50	
ultimate-hop popping.....	50	
link protection.....	13, 29	
bypass LSPs		
administrative groups.....	34	
multiple bypass LSPs.....	36	
RSVP.....	12	
static LSPs.....	29	
switching away from a network node.....	30	
link-protection statement		
MPLS		
usage guidelines.....	29	
RSVP.....	78	
usage guidelines.....	32	
Imp (tracing flag).....	98	
load-balance statement.....	79	
usage guidelines.....	41	

LSPs		no-world-readable option to traceoptions	
bypass.....	29	statement	
multiple bypass.....	14	RSVP.....	99
path		node ID hellos, RSVP.....	28
length.....	74	node protection.....	13, 14
switching away from a network node.....	30	static LSPs.....	29
ultimate-hop popping.....	46	switching away from a network node.....	30
		node-hello statement.....	82
		RSVP	
		usage guidelines.....	28
		node-link-protection statement.....	86
		usage guidelines.....	29
M		O	
manuals		one-to-one backup.....	13
comments on.....	xv	optimize-timer statement	
max-bypasses statement.....	80	bypass LSPs.....	87
usage guidelines.....	36	usage guidelines.....	37
MD5 authentication.....	26	outgoing MTU value in RSVP	
messages		determining.....	11
Resv, RSVP.....	7		
ResvConfirm, RSVP.....	8	P	
ResvErr, RSVP.....	7	packets (tracing flag)	
ResvTear, RSVP.....	7	RSVP.....	98
RSVP message types.....	6	parentheses, in syntax descriptions.....	xiv
RSVP refresh.....	44	path	
MPLS		length, LSP.....	74
ultimate-hop popping.....	46	path (tracing flag)	
MTU signaling, in RSVP.....	9	RSVP.....	98
mtu-signaling statement		path messages, RSVP.....	6
usage guidelines.....	45	path statement	
multiple bypass LSPs.....	14, 33, 34, 36	RSVP.....	88
		usage guidelines.....	37
N		path-mtu statement.....	45
next-hop bypass LSP.....	14	PathErr messages.....	7
next-next-hop bypass LSP.....	14	pathtear (tracing flag).....	98
no-adjacency-down-notification statement		PathTear messages, RSVP.....	7
IS-IS.....	83	peer-interface statement	
no-aggregate statement.....	62	RSVP.....	89
usage guidelines.....	23	PFE fast reroute.....	41
no-cspf statement.....	84	point-to-multipoint LSPs	
usage guidelines.....	36	ultimate-hop popping.....	50
no-interface-hello statement.....	84	preemption	
RSVP		RSVP sessions.....	45
usage guidelines.....	28	preemption statement.....	90
no-local-reversion statement.....	81	usage guidelines.....	45
no-neighbor-down-notification statement.....	85	priority statement	
no-node-id-subobject statement.....	85	RSVP.....	91
usage guidelines.....	31	usage guidelines.....	38
no-node-protection statement			
usage guidelines.....	36		
no-p2mp-sublsp statement.....	86		
no-reliable statement.....	92		
usage guidelines.....	23		

R

receive (tracing flag modifier)	
RSVP.....	99
refresh messages, RSVP.....	44
refresh reduction, RSVP.....	9
refresh time, RSVP.....	44
refresh-time statement.....	92
usage guidelines.....	44
reliable statement.....	92
usage guidelines.....	23
requests, CoS.....	4
reservation styles.....	8
resv (tracing flag).....	98
Resv messages, RSVP.....	7
ResvConfirm messages, RSVP.....	8
ResvErr messages, RSVP.....	7
resvtcar (tracing flag).....	99
ResvTear messages, RSVP.....	7
route (tracing flag).....	99
RRO node ID sub-object, disabling.....	31
RSVP	
aggregation.....	62
authentication.....	26, 63
automatic mesh, configuration.....	42
bandwidth	
reserving.....	97
update threshold.....	27
configuration, minimum.....	21
disabling.....	21
enabling.....	21
example configurations.....	23, 52
Explicit Null label.....	50
graceful restart.....	16, 72
hello acknowledgments.....	29
hello interval.....	25, 73
hello packets.....	5
IGP hello packets.....	5
Implicit Null label.....	50
Junos implementation.....	4
keep multiplier.....	76
link protection.....	29
load balancing.....	41
message types.....	6
MPLS, configuring with RSVP.....	22
MTU signaling in.....	9
node ID hellos.....	28
overview.....	4
preemption.....	45
reservation styles.....	8

sessions.....	22
setup protection.....	39
supported software standards.....	57
timers.....	44, 92
timers, hello packets.....	5
tracing protocol traffic.....	52, 98
ultimate-hop popping.....	50
unnumbered interfaces.....	28
RSVP refresh reduction	
configuration.....	23
overview.....	9
rsvp statement.....	93
mpls.....	93
usage guidelines.....	21
rsvp-te statement.....	94
usage guidelines.....	42

S

SE (reservation style).....	8
send (tracing flag modifier)	
RSVP.....	99
sessions, RSVP.....	22
setup protection, RSVP.....	39
setup-protection statement.....	94
usage guidelines.....	39
shared explicit reservation style.....	8
shared reservations.....	8
signed LSPs	
path	
length.....	74
soft-preemption statement	
RSVP.....	95
state (tracing flag)	
RSVP.....	99
static LSPs	
link protection.....	29
node protection.....	29
static-label-switched-path statement	
static LSP.....	96
subscribing to bandwidth.....	97
subscription statement.....	97
support, technical See technical support	
switch-away-lsps statement	
usage guidelines.....	30
syntax conventions.....	xiii

T

technical support	
contacting JTAC.....	xv

timers	
RSVP.....	44, 92
traceoptions statement	
RSVP.....	98
usage guidelines.....	52
tracing flag modifiers	
detail	
RSVP.....	99
receive	
RSVP.....	99
send	
RSVP.....	99
tracing flags	
all	
RSVP.....	98
error	
RSVP.....	98
event	
RSVP.....	98
Imp.....	98
packets	
RSVP.....	98
path	
RSVP.....	98
pathtear.....	98
resv.....	98
resvtear.....	99
route.....	99
state	
RSVP.....	99
tracing operations	
RSVP.....	52, 98
traffic engineering database	
bandwidth update threshold.....	27
transit statement	
static LSP.....	100
tunnel-services statement.....	101
usage guidelines.....	50

U

ultimate-hop popping	
point-to-multipoint LSPs.....	50
ultimate-hop-popping statement.....	102
usage guidelines.....	46
unnumbered interfaces, RSVP.....	28
update-threshold statement.....	103
usage guidelines.....	27

W

wildcard filter (WF) reservation style.....	8
wildcard senders, RSVP.....	8
world-readable option to traceoptions statement	
RSVP.....	99

