



Alcatel-Lucent

Service Router | Release 12.0 R1
7750 SR OS MPLS Guide

93-0075-11-01 Edition 01
93-0075-11-01

Alcatel, Lucent, Alcatel-Lucent and the Alcatel-Lucent logo are trademarks of Alcatel-Lucent. All other trademarks are the property of their respective owners. The information presented is subject to change without notice. Alcatel-Lucent assumes no responsibility for inaccuracies contained herein.
Copyright © 2014 Alcatel-Lucent
All Rights Reserved.



This document is protected by copyright. Except as specifically permitted herein, no portion of the provided information can be reproduced in any form, or by any means, without prior written permission from Alcatel-Lucent.
Alcatel, Lucent, Alcatel-Lucent and the Alcatel-Lucent logo are trademarks of Alcatel-Lucent. All other trademarks are the property of their respective owners.
The information presented is subject to change without notice.
Alcatel-Lucent assumes no responsibility for inaccuracies contained herein.

Copyright 2014 Alcatel-Lucent. All rights reserved.

Table of Contents

Preface	13
Getting Started	
Alcatel Router Configuration Process	17
MPLS and RSVP	
MPLS	21
MPLS Label Stack	22
Label Values	24
Label Switching Routers	25
LSP Types	25
MPLS Facility Bypass Method of MPLS Fast Re-Route (FRR)	27
Manual Bypass LSP	27
Uniform FRR Failover Time	32
Automatic Bandwidth Allocation for RSVP LSPs	33
MPLS Transport Profile (MPLS-TP)	40
MPLS-TP Model	41
MPLS-TP Provider Edge and Gateway	42
VLL Services	42
Detailed Descriptions of MPLS-TP	44
MPLS-TP LSPs	44
MPLS-TP on Pseudowires	45
MPLS-TP Maintenance Identifiers	45
Generic Associated Channel	49
MPLS-TP Operations, Administration and Maintenance (OAM)	50
PW Control Channel Status Notifications (Static Pseudowire Status Signaling)	54
PW Control Channel Status Request Mechanism	55
Pseudowire Redundancy and Active / Standby Dual Homing	56
MPLS-TP LSP Protection	56
Switching Static MPLS-TP to Dynamic T-LDP Signaled PWs	60
Configuring MPLS-TP	61
Configuration Overview	61
Node-Wide MPLS-TP Parameter Configuration	62
Node-Wide MPLS-TP Identifier Configuration	62
Static LSP and pseudowire (VC) Label and Tunnel Ranges	63
Interface Configuration for MPLS-TP	64
LER Configuration for MPLS-TP	67
Intermediate LSR Configuration for MPLS-TP LSPs	73
MPLS-TP Show Commands	75
Static MPLS Labels	75
MPLS-TP Tunnel Configuration	75
MPLS-TP Path configuration	76
MPLS-TP Protection	80
MPLS TP Node Configuration	80
MPLS-TP Interfaces	82
MPLS-TP Debug Commands	83
RSVP	85

Table of Contents

Using RSVP for MPLS	87
RSVP Traffic Engineering Extensions for MPLS	87
Reservation Styles	89
RSVP Message Pacing	89
RSVP Overhead Refresh Reduction	90
RSVP Graceful Restart Helper	91
Enhancements to RSVP control plane congestion control	92
RSVP LSP Statistics	93
P2MP RSVP-TE LSP Statistics	94
Configuring RSVP P2MP LSP Egress Statistics	95
Configuring RSVP P2MP LSP Ingress Statistics	96
Configuring Implicit Null	98
Using Unnumbered Point-to-Point Interface in RSVP	99
Operation of RSVP FRR Facility Backup over Unnumbered Interface	100
Traffic Engineering	102
TE Metric (IS-IS and OSPF)	103
Admin Group Support on Facility Bypass Backup LSP	104
Procedures at Head-End Node	104
Procedures at PLR Node	105
Diff-Serv Traffic Engineering	107
Mapping of Traffic to a Diff-Serv LSP	107
Admission Control of Classes	107
RSVP Control Plane Extensions	113
IGP Extensions	113
Diff-Serv TE Configuration and Operation	114
Diff-Serv TE LSP Class Type Change under Failure	118
LSP Primary Path Retry Procedures	118
Bandwidth Sharing Across Class Types	121
Downgrading the CT of Bandwidth Sharing LSP Paths	123
Upgrading the CT of Bandwidth Sharing LSP Paths	124
Advanced MPLS/RSVP Features	125
Extending RSVP LSP to use Loopback Interfaces Other Than router-id	125
LSP Path Change	125
Manual LSP Path Switch	127
Make-Before-Break (MBB) Procedures for LSP/Path Parameter Configuration Change	128
Automatic Creation of RSVP-TE LSP Mesh	130
RSVP-TE LSP Shortcut for IGP Resolution	132
Using LSP Relative Metric with IGP Shortcut	133
ECMP Considerations	134
Handling of Control Packets	135
Forwarding Adjacency	135
LDP Forwarding over IGP Shortcut	136
Handling of Multicast Packets	137
RSVP-TE LSP Signaling using LSP Template	139
Shared Risk Link Groups	140
Enabling Disjoint Backup Paths	140
Static Configurations of SRLG Memberships	143
TE Graceful Shutdown	144
Soft Pre-emption of Diff-Serv RSVP LSP	144
Least-Fill Bandwidth Rule in CSPF ECMP Selection	145
Inter Area TE LSP (ERO Expansion Method)	145

Area Border Node FRR Protection for Inter Area LSP	145
Inter-Area LSP support of OSPF Virtual Links	149
Area Border Node FRR Protection for Inter Area LSP	149
Automatic Creation of a RSVP Mesh LSP	152
Feature Configuration	152
Feature Behavior	153
Multi-Area and Multi-Instance Support	154
Mesh LSP Name Encoding and Statistics	154
Automatic Creation of an RSVP One-Hop LSP	156
Feature Configuration	156
Feature Behavior	156
Point-to-Multipoint (P2MP) RSVP LSP	158
Application in Video Broadcast	158
P2MP LSP Data Plane	159
Procedures at Ingress LER Node	159
Procedures at LSR Node	160
Procedures at Branch LSR Node	160
Procedures at Egress LER Node	161
Procedures at BUD LSR Node	161
Ingress Path Management for P2MP LSP Packets	162
Ingress P2MP Path Management on IOM-3/IMMs	162
Ingress P2MP Path Management on IOM-2	164
RSVP Control Plane in a P2MP LSP	165
Forwarding Multicast Packets over RSVP P2MP LSP in the Base Router	168
Procedures at Ingress LER Node	168
Procedures at Egress LER Node	169
MPLS Service Usage	171
Service Distribution Paths	171
MPLS/RSVP Configuration Process Overview	172
Configuration Notes	173
Configuring MPLS and RSVP with CLI	175
MPLS Configuration Overview	176
LSPs	176
Paths	176
Router Interface	177
Choosing the Signaling Protocol	177
Basic MPLS Configuration	178
Common Configuration Tasks	179
Configuring MPLS Components	180
Configuring Global MPLS Parameters	180
Configuring an MPLS Interface	181
Configuring MPLS Paths	182
Configuring an MPLS LSP	183
Configuring a Static LSP	184
Configuring Manual Bypass Tunnels	185
Configuring RSVP Parameters	187
Configuring RSVP Message Pacing Parameters	188
Configuring Graceful Shutdown	189
MPLS Configuration Management Tasks	190
Deleting MPLS	190
Modifying MPLS Parameters	190

Table of Contents

Modifying an MPLS LSP	191
Modifying MPLS Path Parameters	192
Modifying MPLS Static LSP Parameters	193
Deleting an MPLS Interface	194
RSVP Configuration Management Tasks	195
Modifying RSVP Parameters	195
Modifying RSVP Message Pacing Parameters	196
Deleting an Interface from RSVP	196
MPLS/RSVP Command Reference	197
Command Hierarchies	197
MPLS Commands	197
MPLS-TP Commands	200
LSP Commands	201
MPLS Path Commands	204
RSVP Commands	205
Show Commands	207
Tools Commands	209
Router Commands	209
Clear Commands	210
Debug Commands	211

Label Distribution Protocol

Label Distribution Protocol	430
LDP and MPLS	430
LDP Architecture	431
Subsystem Interrelationships	432
Memory Manager and LDP	433
Label Manager	433
LDP Configuration	433
Logger	433
Service Manager	433
Execution Flow	434
Initialization	434
Session Lifetime	434
Label Exchange	436
Other Reasons for Label Actions	436
Cleanup	436
Configuring Implicit Null Label	437
Global LDP Filters	438
Per LDP Peer FEC Import and Export Policies	438
Configuring Multiple LDP LSR ID	439
T-LDP hello reduction	440
Tracking a T-LDP Peer with BFD	440
Tracking a Link LDP Peer with BFD	441
LDP LSP Statistics	442
TTL Security for <i>BGP and LDP</i>	443
ECMP Support for LDP	444
Label Operations	444
Unnumbered Interface Support in LDP	446
Feature Configuration	446

Operation of LDP over an Unnumbered IP Interface	446
Link LDP	447
Targeted LDP	448
FEC Resolution	448
LDP over RSVP Tunnels	449
LDP over RSVP Without Area Boundary	454
<i>LDP over RSVP and ECMP</i>	455
LDP ECMP Uniform Failover	456
LDP Fast-Reroute for IS-IS and OSPF Prefixes	458
LDP FRR Configuration	458
LDP FRR Procedures	459
IS-IS and OSPF Support for Loop-Free Alternate Calculation	463
mLDP Fast Upstream Switchover	467
LDP FEC to BGP Label Route Stitching	469
Configuration	470
Detailed LDP FEC Resolution	471
Detailed BGP Labeled Route Resolution	472
Data Plane Forwarding	472
Automatic Creation of a Targeted Hello Adjacency and LDP Session	473
Feature Configuration	473
Feature Behavior	474
Multicast P2MP LDP for GRT	478
LDP P2MP Support	479
LDP P2MP Configuration	479
LDP P2MP Protocol	479
Make Before Break (MBB)	479
ECMP Support	479
Multicast LDP Fast Upstream Switchover	480
Feature Configuration	480
Feature Behavior	481
Uniform Failover from Primary to Backup ILM	483
<i>Multi-Area and Multi-Instance Extensions to LDP</i>	485
LDP Shortcut for BGP Next-Hop Resolution	486
LDP Shortcut for IGP Routes	487
ECMP Considerations	489
Disabling TTL Propagation in an LSP Shortcut	489
LDP Graceful Handling of Resource Exhaustion	490
LDP Base Graceful Handling of Resources	490
LDP Enhanced Graceful Handling of Resources	491
LSR Overload Notification	492
LSR Overload Protection Capability	493
Procedures for LSR overload protection	493
User Guidelines and Troubleshooting Procedures	495
Common Procedures	495
Base Resource Handling Procedures	496
Enhanced Resource Handling Procedures	499
LDP Process Overview	503
Configuring LDP with CLI	505
LDP Configuration Overview	506
Basic LDP Configuration	507
Common Configuration Tasks	508

Table of Contents

Enabling LDP	508
Configuring FEC Originate Parameters	509
Configuring Graceful-Restart Helper Parameters	510
Applying Export and Import Policies	511
Targeted Session Parameters	512
Interface Parameters	513
Peer Parameters	514
LDP Signaling and Services	515
LDP Configuration Management Tasks	518
Disabling LDP	518
Modifying Targeted Session Parameters	519
Modifying Interface Parameters	520
LDP Command Reference	521
Command Hierarchies	521
LDP Commands	521
Show Commands	524
Clear Commands	525
Debug Commands	525
Tools Commands	525
Index	629

List of Tables

Getting Started

Table 1:	Configuration Process.....	17
----------	----------------------------	----

MPLS and RSVP

Table 2:	Packet/Label Field Description	22
Table 3:	Mapping from RSVP-TE to MPLS-TP Maintenance Identifiers	47
Table 4:	Bypass LSP Admin-Group Constraint Behavior	105
Table 5:	Internal TE Class Definition when Diff-Serv TE is Disabled	115
Table 6:	Default Mapping of Forwarding Class to TE Class	115
Table 7:	RSVP LSP Role As Outcome of LSP level and IGP level configuration options.	133
Table 8:	Impact of LSP level configuration on IGP shortcut and forwarding adjacency features	136
Table 9:	Impact of IGP Shortcut and Forwarding Adjacency on Unicast and Multicast RTM	137

Label Distribution Protocol

Table 10:	Triggering Events and the Associated Priority	476
-----------	---	-----

LIST OF FIGURES

MPLS and RSVP

Figure 1:	Label Placement	22
Figure 2:	Label Packet Placement	23
Figure 3:	Bypass Tunnel Nodes	28
Figure 4:	FRR Node-Protection Example	31
Figure 5:	MPLS-TP Model	41
Figure 6:	MPLS-TP Provider Edge and Gateway, VLL Services	42
Figure 7:	MPLS-TP Provider Edge and Gateway, spoke-SDP Termination on VPLS	43
Figure 8:	MPLS-TP Provider Edge and Gateway, spoke-SDP Termination on IES/VP RN	43
Figure 9:	MPLS-TP LSR	44
Figure 10:	MPLS-TP Maintenance Architecture	46
Figure 11:	MPLS-TP LSP and Tunnel Information Model	47
Figure 12:	MPLS-TP PW Information Model	48
Figure 13:	Example usage of PW Identifiers	49
Figure 14:	Label for LSP and PW G-ACh Packets	50
Figure 15:	BFD used for proactive CC on MPLS-TP LSP	52
Figure 16:	BFD used for proactive CV on MPLS-TP LSP	53
Figure 17:	Normal Operation	57
Figure 18:	Failed Condition	58
Figure 19:	Failed Condition - Switching at A	58
Figure 20:	Failed Condition - Switching at Z	59
Figure 21:	Static - Dynamic PW Switching with MPLS-TP	60
Figure 22:	Establishing LSPs	85
Figure 23:	LSP Using RSVP Path Set Up	86
Figure 24:	RDM with Two Class Types	110
Figure 25:	First LSP Reservation	110
Figure 26:	Second LSP Reservation	111
Figure 27:	RDM Admission Control Policy Example	121
Figure 28:	Sharing bandwidth when an LSP primary path is downgraded to backup CT	123
Figure 29:	Sharing Bandwidth When an LSP Primary Path is Upgraded to Main CT	124
Figure 30:	Shared Risk Link Groups	142
Figure 31:	Automatic ABR Node Selection for Inter-Area LSP	146
Figure 32:	ABR Node Protection Using Dynamic Bypass LSP	150
Figure 33:	Application of P2MP LSP in Video Broadcast	158
Figure 34:	MPLS and RSVP Configuration and Implementation Flow	172

Label Distribution Protocol

Figure 35:	Subsystem Interrelationships	432
Figure 36:	LDP Adjacency and Session over Unnumbered Interface	446
Figure 37:	LDP over RSVP Application	449
Figure 38:	LDP over RSVP Application Variant	450
Figure 39:	LDP over RSVP Without ABR Stitching Point	454
Figure 40:	Topology with Primary and LFA Routes	463
Figure 41:	Example Topology with Broadcast Interfaces	464
Figure 42:	Application of LDP to BGP FEC Stitching	469
Figure 43:	Video Distribution using P2MP LDP	478
Figure 44:	mLDP LSP with Backup Upstream LSR Nodes	482

List of Figures

Figure 45:	LSR Overload Status TLV (Type = 0x3E02)	492
Figure 46:	LSR Overload Protection Capability TLV (Type== 0x3E03)	493
Figure 47:	LDP Configuration and Implementation.	504

About This Guide

This guide describes the services and protocol support provided by the router and presents examples to configure and implement MPLS, RSVP, and LDP protocols.

This document is organized into functional chapters and provides concepts and descriptions of the implementation flow, as well as Command Line Interface (CLI) syntax and command usage.

Audience

This manual is intended for network administrators who are responsible for configuring routers. It is assumed that the network administrators have an understanding of networking principles and configurations. Protocols and concepts described in this manual include the following:

- Multiprotocol Label Switching (MPLS)
- Resource Reservation Protocol (RSVP)
- Label Distribution Protocol (LDP)

List of Technical Publications

- **7750 SR OS Basic System Configuration Guide**
This guide describes basic system configurations and operations.
- **7750 SR OS System Management Guide**
This guide describes system security and access configurations as well as event logging and accounting logs.
- **7750 SR OS Interface Configuration Guide**
- **7750 SR OS Router Configuration Guide**
This guide describes logical IP routing interfaces and associated attributes such as an IP address, as well as IP and MAC-based filtering, and VRRP and Cflowd.
- **7750 SR OS Routing Protocols Guide**
This guide provides an overview of routing concepts and provides configuration examples for RIP, OSPF, IS-IS, BGP, and route policies.
- **7750 SR OS MPLS Guide**
This guide describes how to configure Multiprotocol Label Switching (MPLS) and Label Distribution Protocol (LDP).
- **7750 SR OS Services Guide**
This guide describes how to configure service parameters such as service distribution points (SDPs), customer information, and user services.
- **7750 SR OAM and Diagnostic Guide**
This guide describes how to configure features such as service mirroring and Operations, Administration and Management (OAM) tools.
- **7750 SR OS Triple Play Guide**
This guide describes Triple Play services and support provided by the 7750 SR and presents examples to configure and implement various protocols and services.
- **7750 SR OS Quality of Service Guide**
This guide describes how to configure Quality of Service (QoS) policy management.
- **OS Multi-Service ISA Guide**
This guide describes services provided by integrated service adapters such as Application Assurance, IPSec, ad insertion (ADI) and Network Address Translation (NAT).
- **7750 SR-OS RADIUS Attributes Reference Guide**
This guide describes all supported RADIUS Authentication, Authorization and Accounting attributes.

Technical Support

If you purchased a service agreement for your router and related products from a distributor or authorized reseller, contact the technical support staff for that distributor or reseller for assistance.

If you purchased an TiMetra Systems service agreement, contact technical assistance at:

<http://www.alcatel-lucent.com/wps/portal/support>

Report documentation errors, omissions and comments to:

ipd_online_feedback@alcatel-lucent.com

Include document name, version, part number and page(s) affected.

GETTING STARTED

In This Chapter

This chapter provides process flow information to configure MPLS, RSVP, and LDP protocols.

Alcatel Router Configuration Process

[Table 1](#) lists the tasks necessary to configure MPLS applications functions.

This guide is presented in an overall logical configuration flow. Each section describes a software area and provides CLI syntax and command usage to configure parameters for a functional area.

Table 1: Configuration Process

Area	Task	Chapter
Protocol configuration	Configure MPLS protocols: <ul style="list-style-type: none">• MPLS• RSVP• LDP	MPLS on page 21 RSVP on page 85 Label Distribution Protocol on page 430
Reference	List of IEEE, IETF, and other proprietary entities.	Standards and Protocol Support on page 623

In This Chapter

This chapter provides information to configure MPLS and RSVP.

- [MPLS on page 21](#)
 - [MPLS Label Stack on page 22](#)
 - [Label Switching Routers on page 25](#)
- [MPLS Transport Profile \(MPLS-TP\) on page 40](#)
- [RSVP on page 85](#)
 - [Using RSVP for MPLS on page 87](#)
 - [Reservation Styles on page 89](#)
 - [RSVP Overhead Refresh Reduction on page 90](#)
 - [RSVP Graceful Restart Helper on page 91](#)
 - [Enhancements to RSVP control plane congestion control on page 92](#)
 - [RSVP LSP Statistics on page 93](#)
- [Traffic Engineering on page 102](#)
 - [TE Metric \(IS-IS and OSPF\) on page 103](#)
 - [Diff-Serv Traffic Engineering on page 107](#)
 - [Diff-Serv TE LSP Class Type Change under Failure on page 118](#)
- [Advanced MPLS/RSVP Features on page 125](#)
 - [LSP Path Change on page 125](#)
 - [Make-Before-Break \(MBB\) Procedures for LSP/Path Parameter Configuration Change on page 128](#)
 - [Automatic Creation of RSVP-TE LSP Mesh on page 130](#)
 - [RSVP-TE LSP Shortcut for IGP Resolution on page 132](#)
 - [Shared Risk Link Groups on page 140](#)
 - [TE Graceful Shutdown on page 144](#)

- [Soft Pre-emption of Diff-Serv RSVP LSP on page 144](#)
- [Least-Fill Bandwidth Rule in CSPF ECMP Selection on page 145](#)
- [Automatic Creation of an RSVP One-Hop LSP on page 156](#)
- [Point-to-Multipoint \(P2MP\) RSVP LSP on page 158](#)
 - [Application in Video Broadcast on page 158](#)
 - [P2MP LSP Data Plane on page 159](#)
 - [Ingress Path Management for P2MP LSP Packets on page 162](#)
 - [RSVP Control Plane in a P2MP LSP on page 165](#)
 - [Forwarding Multicast Packets over RSVP P2MP LSP in the Base Router on page 168](#)
- [MPLS Service Usage on page 171](#)
 - [Service Distribution Paths on page 171](#)
- [MPLS/RSVP Configuration Process Overview on page 172](#)
- [Configuration Notes on page 173](#)

MPLS

Multiprotocol Label Switching (MPLS) is a label switching technology that provides the ability to set up connection-oriented paths over a connectionless IP network. MPLS facilitates network traffic flow and provides a mechanism to engineer network traffic patterns independently from routing tables. MPLS sets up a specific path for a sequence of packets. The packets are identified by a label inserted into each packet. MPLS is not enabled by default and must be explicitly enabled.

MPLS is independent of any routing protocol but is considered multiprotocol because it works with the Internet Protocol (IP), Asynchronous Transport Mode (ATM), and frame relay network protocols.

MPLS Label Stack

MPLS requires a set of procedures to enhance network layer packets with label stacks which thereby turns them into labeled packets. Routers that support MPLS are known as Label Switching Routers (LSRs). In order to transmit a labeled packet on a particular data link, an LSR must support the encoding technique which, when given a label stack and a network layer packet, produces a labeled packet.

In MPLS, packets can carry not just one label, but a set of labels in a stack. An LSR can swap the label at the top of the stack, pop the stack, or swap the label and push one or more labels into the stack. The processing of a labeled packet is completely independent of the level of hierarchy. The processing is always based on the top label, without regard for the possibility that some number of other labels may have been above it in the past, or that some number of other labels may be below it at present.

As described in RFC 3032, *MPLS Label Stack Encoding*, the label stack is represented as a sequence of label stack entries. Each label stack entry is represented by 4 octets. [Figure 1](#) displays the label placement in a packet.





Figure 1: Label Placement

Table 2: Packet/Label Field Description

Field	Description
Label	This 20-bit field carries the actual value (unstructured) of the label.
Exp	This 3-bit field is reserved for experimental use. It is currently used for Class of Service (CoS).
S	This bit is set to 1 for the last entry (bottom) in the label stack, and 0 for all other label stack entries.
TTL	This 8-bit field is used to encode a TTL value.

A stack can carry several labels, organized in a last in/first out order. The top of the label stack appears first in the packet and the bottom of the stack appears last (Figure 2).



OSSG014

Figure 2: Label Packet Placement

The label value at the top of the stack is looked up when a labeled packet is received. A successful lookup reveals:

- The next hop where the packet is to be forwarded.
- The operation to be performed on the label stack before forwarding.

In addition, the lookup may reveal outgoing data link encapsulation and other information needed to properly forward the packet.

An empty label stack can be thought of as an unlabeled packet. An empty label stack has zero (0) depth. The label at the bottom of the stack is referred to as the Level 1 label. The label above it (if it exists) is the Level 2 label, and so on. The label at the top of the stack is referred to as the Level m label.

Labeled packet processing is independent of the level of hierarchy. Processing is always based on the top label in the stack which includes information about the operations to perform on the packet's label stack.

Label Values

Packets travelling along an LSP (see [Label Switching Routers on page 25](#)) are identified by its label, the 20-bit, unsigned integer. The range is 0 through 1,048,575. Label values 0-15 are reserved and are defined below as follows:

- A value of 0 represents the IPv4 Explicit NULL Label. This Label value is legal only at the bottom of the Label stack. It indicates that the Label stack must be popped, and the packet forwarding must be based on the IPv4 header.
- A value of 1 represents the router alert Label. This Label value is legal anywhere in the Label stack except at the bottom. When a received packet contains this Label value at the top of the Label stack, it is delivered to a local software module for processing. The actual packet forwarding is determined by the Label beneath it in the stack. However, if the packet is further forwarded, the router alert Label should be pushed back onto the Label stack before forwarding. The use of this Label is analogous to the use of the router alert option in IP packets. Since this Label cannot occur at the bottom of the stack, it is not associated with a particular network layer protocol.
- A value of 2 represents the IPv6 explicit NULL Label. This Label value is only legal at the bottom of the Label stack. It indicates that the Label stack must be popped, and the packet forwarding must be based on the IPv6 header.
- A value of 3 represents the Implicit NULL Label. This is a Label that a Label Switching Router (LSR) can assign and distribute, but which never actually appears in the encapsulation. When an LSR would otherwise replace the Label at the top of the stack with a new Label, but the new Label is Implicit NULL, the LSR pops the stack instead of doing the replacement. Although this value may never appear in the encapsulation, it needs to be specified in the Label Distribution Protocol (LDP), so a value is reserved.
- Values 4-15 are reserved for future use.

The router uses labels for MPLS, RSVP-TE, and LDP, as well as packet-based services such as VLL and VPLS.

Label values 16 through 1,048,575 are defined as follows:

- Label values 16 through 31 are reserved for future use.
- Label values 32 through 1,023 are available for static LSP label assignments.
- Label values 1,024 through 2,047 are reserved for future use.
- Label values 2,048 through 18,431 are available for static service label assignments
- Label values 18,432 through 262,143 (131,071 in chassis modes lower than D) are assigned dynamically by RSVP, LDP, and BGP control planes for both MPLS LSP and service labels.
- Label values 262,144 (131,072 in chassis modes lower than D) through 1,048,575 are reserved for future use.

Label Switching Routers

LSRs perform the label switching function. LSRs perform different functions based on its position in an LSP. Routers in an LSP do one of the following:

- The router at the beginning of an LSP is the ingress label edge router (ILER). The ingress router can encapsulate packets with an MPLS header and forward it to the next router along the path. An LSP can only have one ingress router.
- A Label Switching Router (LSR) can be any intermediate router in the LSP between the ingress and egress routers. An LSR swaps the incoming label with the outgoing MPLS label and forwards the MPLS packets it receives to the next router in the MPLS path (LSP). An LSP can have 0-253 transit routers.
- The router at the end of an LSP is the egress label edge router (ELER). The egress router strips the MPLS encapsulation which changes it from an MPLS packet to a data packet, and then forwards the packet to its final destination using information in the forwarding table. Each LSP can have only one egress router. The ingress and egress routers in an LSP cannot be the same router.

A router in your network can act as an ingress, egress, or transit router for one or more LSPs, depending on your network design.

An LSP is confined to one IGP area for LSPs using constrained-path. They cannot cross an autonomous system (AS) boundary.

Static LSPs can cross AS boundaries. The intermediate hops are manually configured so the LSP has no dependence on the IGP topology or a local forwarding table.

LSP Types

The following are LSP types:

- **Static LSPs** — A static LSP specifies a static path. All routers that the LSP traverses must be configured manually with labels. No signaling such as RSVP or LDP is required.
- **Signaled LSP** — LSPs are set up using a signaling protocol such as RSVP-TE or LDP. The signaling protocol allows labels to be assigned from an ingress router to the egress router. Signaling is triggered by the ingress routers. Configuration is required only on the ingress router and is not required on intermediate routers. Signaling also facilitates path selection.

There are two signaled LSP types:

- **Explicit-path LSPs** — MPLS uses RSVP-TE to set up explicit path LSPs. The hops within the LSP are configured manually. The intermediate hops must be configured as either strict or loose meaning that the LSP must take either a direct path from the

previous hop router to this router (strict) or can traverse through other routers (loose). You can control how the path is set up. They are similar to static LSPs but require less configuration. See [RSVP on page 85](#).

- Constrained-path LSPs — The intermediate hops of the LSP are dynamically assigned. A constrained path LSP relies on the Constrained Shortest Path First (CSPF) routing algorithm to find a path which satisfies the constraints for the LSP. In turn, CSPF relies on the topology database provided by the extended IGP such as OSPF or IS-IS.

Once the path is found by CSPF, RSVP uses the path to request the LSP set up. CSPF calculates the shortest path based on the constraints provided such as bandwidth, class of service, and specified hops.

If fast reroute is configured, the ingress router signals the routers downstream. Each downstream router sets up a detour for the LSP. If a downstream router does not support fast reroute, the request is ignored and the router continues to support the LSP. This can cause some of the detours to fail, but otherwise the LSP is not impacted.

No bandwidth is reserved for the rerouted path. If the user enters a value in the bandwidth parameter in the **config>router>mpls>lsp>fast-reroute** context, it will have no effect on the LSP backup LSP establishment.

Hop-limit parameter specifies the maximum number of hops that an LSP can traverse, including the ingress and egress routers. An LSP is not set up if the hop limit is exceeded. The hop count is set to 255 by default for the primary and secondary paths. It is set to 16 by default for a bypass or detour LSP path.

MPLS Facility Bypass Method of MPLS Fast Re-Route (FRR)

The MPLS facility bypass method of MPLS Fast Re-Route (FRR) functionality is extended to the ingress node.

The behavior of an LSP at an ingress LER with both fast reroute and a standby LSP path configured is as follows:

- When a down stream detour becomes active at a point of local repair (PLR):
The ingress LER switches to the standby LSP path. If the primary LSP path is repaired subsequently at the PLR, the LSP will switch back to the primary path. If the standby goes down, the LSP is switched back to the primary, even though it is still on the detour at the PLR. If the primary goes down at the ingress while the LSP is on the standby, the detour at the ingress is cleaned up and for one-to-one detours a “path tear” is sent for the detour path. In other words, the detour at the ingress does not protect the standby. If and when the primary LSP is again successfully re-signaled, the ingress detour state machine will be restarted.
 - When the primary fails at the ingress:
The LSP switches to the detour path. If a standby is available then LSP would switch to standby on expiration of **hold-timer**. If **hold-timer** is disabled then switchover to standby would happen immediately. On successful global revert of primary path, the LSP would switch back to the primary path.
 - Admin groups are not taken into account when creating detours for LSPs.
-

Manual Bypass LSP

In prior releases, the router implemented dynamic bypass tunnels as per RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*. When an LSP is signaled and the local protection flag is set in the session_attribute object and/or the FRR object in the path message indicates that facility backup is desired, the PLR will establish a bypass tunnel to provide node and link protection. If a bypass LSP which merges in a downstream node with the protected LSP exist, and if this LSP satisfies the constraints in the FRR object, then this bypass tunnel is selected.

With the manual bypass feature, an LSP can be pre-configured from a PLR which will be used exclusively for bypass protection. When a path message for a new LSP requests bypass protection, the node will first check if a manual bypass tunnel satisfying the path constraints exists. If one is found, it will be selected. If no manual bypass tunnel is found, the router will dynamically signal a bypass LSP in the default behavior. Users can disable the dynamic bypass creation on a per node basis using the CLI.

A maximum of 1000 associations of primary LSP paths can be made with a single manual bypass by default. The **max-bypass-associations** *integer* command increases the number of associations.

If dynamic bypass creation is disabled on the node, it is recommended to configure additional manual bypass LSPs to handle the required number of associations.

Refer to [Configuring Manual Bypass Tunnels on page 185](#) for configuration information.

PLR Bypass LSP Selection Rules




Figure 3: Bypass Tunnel Nodes

The PLR uses the following rules to select a bypass LSP among multiple manual and dynamic bypass LSPs at the time of establishment of the primary LSP path or when searching for a bypass for a protected LSP which does not have an association with a bypass tunnel:

1. The MPLS/RSVP task in the PLR node checks if an existing manual bypass satisfies the constraints. If the path message for the primary LSP path indicated node protection desired, which is the default LSP FRR setting at the head end node, MPLS/RSVP task searches for a node-protect' bypass LSP. If the path message for the primary LSP path indicated link protection desired, then it searches for a link-protect bypass LSP.
2. If multiple manual bypass LSPs satisfying the path constraints exist, it will prefer a manual-bypass terminating closer to the PLR over a manual bypass terminating further away. If multiple manual bypass LSPs satisfying the path constraints terminate on the same downstream node, it selects one with the lowest IGP path cost or if in a tie, picks the first one available.
3. If none satisfies the constraints and dynamic bypass tunnels have not been disabled on PLR node, then the MPLS/RSVP task in the PLR will check if any of the already established dynamic bypasses of the requested type satisfies the constraints.
4. If none do, then the MPLS/RSVP task will ask CSPF to check if a new dynamic bypass of the requested type, node-protect or link-protect, can be established.
5. If the path message for the primary LSP path indicated node protection desired, and no manual bypass was found after Step 1, and/or no dynamic bypass LSP was found after 3 attempts of performing Step 3, the MPLS/RSVP task will repeat Steps 1-3 looking for a suitable link-protect bypass LSP. If none are found, the primary LSP will have no protection and the PLR node must clear the "local protection available" flag in the IPv4

address sub-object of the RRO starting in the next Resv refresh message it sends upstream.

6. If the path message for the primary LSP path indicated link protection desired, and no manual bypass was found after step 1, and/or no dynamic bypass LSP was found after performing Step 3, the primary LSP will have no protection and the PLR node must clear the “local protection available” flag in the IPv4 address sub-object of the RRO starting in the next RESV refresh message it sends upstream. The PLR will not search for a node-protect’ bypass LSP in this case.
7. If the PLR node successfully makes an association, it must set the “local protection available” flag in the IPv4 address sub-object of the RRO starting in the next RESV refresh message it sends upstream.
8. For all primary LSP that requested FRR protection but are not currently associated with a bypass tunnel, the PLR node on reception of RESV refresh on the primary LSP path repeats Steps 1-7.

If the user disables dynamic-bypass tunnels on a node while dynamic bypass tunnels were activated and were passing traffic, traffic loss will occur on the protected LSP. Furthermore, if no manual bypass exist that satisfy the constraints of the protected LSP, the LSP will remain without protection.

If the user configures a bypass tunnel on node B and dynamic bypass tunnels have been disabled, LSPs which have been previously signaled and which were not associated with any manual bypass tunnel, for example, none existed, will be associated with the manual bypass tunnel if suitable. The node checks for the availability of a suitable bypass tunnel for each of the outstanding LSPs every time a RESV message is received for these LSPs.

If the user configures a bypass tunnel on node B and dynamic bypass tunnels have not been disabled, LSPs which have been previously signaled over dynamic bypass tunnels will not automatically be switched into the manual bypass tunnel even if the manual bypass is a more optimized path. The user will have to perform a make before break at the head end of these LSPs.

If the manual bypass goes into the down state in node B and dynamic bypass tunnels have been disabled, node B (PLR) will clear the “protection available” flag in the RRO IPv4 sub-object in the next RESV refresh message for each affected LSP. It will then try to associate each of these LSPs with one of the manual bypass tunnels that are still up. If it finds one, it will make the association and set again the “protection available” flag in the next RESV refresh message for each of these LSPs. If it could not find one, it will keep checking for one every time a RESV message is received for each of the remaining LSPs. When the manual bypass tunnel is back UP, the LSPs which did not find a match will be associated back to this tunnel and the protection available flag is set starting in the next RESV refresh message.

If the manual bypass goes into the down state in node B and dynamic bypass tunnels have not been disabled, node B will automatically signal a dynamic bypass to protect the LSPs if a suitable one does not exist. Similarly, if an LSP is signaled while the manual bypass is in the down state, the node will only signal a dynamic bypass tunnel if the user has not disabled dynamic tunnels.

When the manual bypass tunnel is back into the UP state, the node will not switch the protected LSPs from the dynamic bypass tunnel into the manual bypass tunnel.

FRR Node-Protection (Facility)

The MPLS Fast Re-Route (FRR) functionality enables PLRs to be aware of the missing node protection and lets them regularly probe for a node-bypass. The following describes an LSP scenario:




Figure 4: FRR Node-Protection Example

Where:

- LSP 1: between PE_1 to PE_2, with CSPF, FRR facility node-protect enabled.
- P_1 protects P_2 with bypass-nodes P_1 -P_3 - P_4 - PE_4 -PE_3.
- If P_4 fails, P_1 tries to establish the bypass-node three times.
- When the bypass-node creation fails, P_1 will protect link P_1-P_2.
- P_1 protects the link to P_2 through P_1 - P_5 - P_2.
- P_4 returns online.

Since LSP 1 had requested node protection, but due to lack of any available path, it could only obtain link protection. Therefore, every 60 seconds the PLR for LSP 1 will search for a new path that might be able to provide node protection. Once P_4 is back online and such a path is available, A new bypass tunnel will be signalled and LSP 1 will get associated with this new bypass tunnel.

Uniform FRR Failover Time

The failover time during FRR consists of a detection time and a switchover time. The detection time corresponds to the time it takes for the RSVP control plane protocol to detect that a network IP interface is down or that a neighbor/next-hop over a network IP interface is down. The control plane can be informed of an interface down event when event is due to a failure in a lower layer such in the physical layer. The control plane can also detect the failure of a neighbor/next-hop on its own by running a protocol such as Hello, Keep-Alive, or BFD.

The switchover time is measured from the time the control plane detected the failure of the interface or neighbor/next-hop to the time the IOM completed the reprogramming of all the impacted ILM or service records in the data path. This includes the time it takes for the control plane to send a down notification to all IOMs to request a switch to the backup NHLFE.

Uniform Fast-Reroute (FRR) failover enables the switchover of MPLS and service packets from the outgoing interface of the primary LSP path to that of the FRR backup LSP within the same amount of time regardless of the number of LSPs or service records. This is achieved by updating Ingress Label Map (ILM) records and service records to point to the backup Next-Hop Label to Forwarding Entry (NHLFE) in a single operation.

Automatic Bandwidth Allocation for RSVP LSPs

Enabling and Disabling Auto-Bandwidth Allocation on an LSP

This section discusses an auto-bandwidth hierarchy configurable in the **config>router>mpls>lsp** context.

Adding auto-bandwidth at the LSP level starts the measurement of LSP bandwidth described in [Measurement of LSP Bandwidth on page 34](#) and allows auto-bandwidth adjustments to take place based on the triggers described in [Periodic Automatic Bandwidth Adjustment on page 36](#).

When an LSP is first established, the bandwidth reserved along its primary path is controlled by the bandwidth parameter in the **config>router>mpls>lsp>primary** context, whether or not the LSP has auto-bandwidth enabled. When auto-bandwidth is enabled and a trigger occurs, the system will attempt to change the bandwidth of the LSP to a value between **min-bandwidth** and **max-bandwidth**, which are configurable values in the **lsp>auto-bandwidth** context. **min-bandwidth** is the minimum bandwidth that auto-bandwidth can signal for the LSP and **max-bandwidth** is the maximum bandwidth that can be signaled. The user can set the **min-bandwidth** to the same value as the primary path bandwidth but the system will not enforce this restriction. The system will allow:

- No **min-bandwidth** to be configured. In this case, the implicit minimum is 0 Mbps
- No **max-bandwidth** to be configured, as long as overflow-triggered auto-bandwidth is not configured. In this case, the implicit maximum is infinite (effectively 100 Gbps).
- The configured primary path bandwidth to be outside the range of min-bandwidth to max-bandwidth.
- **auto-bandwidth** parameters can be changed at any time on an operational LSP; in most cases the changes have no immediate impact but subsequent sections will describe some exceptions

All of the auto-bandwidth adjustments discussed are performed using MBB procedures.

Auto bandwidth can be added to an operational LSP at any time (without the need to shut down the LSP or path), but no bandwidth change occurs until a future trigger event. Auto bandwidth may also be removed from an operational LSP at any time and this causes an immediate MBB bandwidth change to be attempted using the configured primary path bandwidth.

Note that changing the configured bandwidth of an auto-bandwidth LSP has no immediate affect, it will only matters if the LSP/path goes down (due to failure or administrative action) and comes back up or if auto-bandwidth is removed from the LSP. The operator can force an auto-bandwidth LSP to be resized immediately to an arbitrary bandwidth using the appropriate tools commands.

Measurement of LSP Bandwidth

Automatic adjustment of RSVP LSP bandwidth based on measured traffic rate into the tunnel requires the LSP to be configured for egress statistics collection at the ingress LER. The following CLI shows an example:

```
config router mpls lsp name
    egress-statistics
    accounting-policy 99
    collect-stats
    no shutdown
exit
```

All LSPs configured for accounting, including any configured for auto-bandwidth based on traffic measurements, must reference the same accounting policy. An example configuration of such an accounting-policy is shown below: in the CLI example below.

```
config log
    accounting-policy 99
    collection-interval 5
        record combined-mpls-lsp-egress
    exit
exit
```

Note that the record **combined-mpls-lsp-egress** command in the accounting policy has the effect of recording both egress packet and byte counts and bandwidth measurements based on the byte counts if auto-bandwidth is enabled on the LSP.

When egress statistics are enabled the CPM collects stats from of all IOMs involved in forwarding traffic belonging to the LSP (whether the traffic is currently leaving the ingress LER via the primary LSP path, a secondary LSP path, an FRR detour path or an FRR bypass path). The egress statistics have counts for the number of packets and bytes forwarded per LSP on a per-forwarding class, per-priority (in-profile vs. out-of-profile) basis. When auto-bandwidth is configured for an LSP the ingress LER calculates a traffic rate for the LSP as follows:

Average data rate of LSP[x] during interval[i] = $F(x, i) - F(x, i-1) / \text{sample interval}$

$F(x, i)$ — The total number of bytes belonging to LSP[x], regardless of forwarding-class or priority, at time[i]

sample interval = time[i] — time [i-1], time[i+1] — time[i], etc.

The sample interval is the product of sample-multiplier and the collection-interval specified in the auto-bandwidth accounting policy. A default sample-multiplier for all LSPs may be configured using the **config>router>mpls>auto-bandwidth-defaults** command but this value can be overridden on a per-LSP basis at the **config>router>mpls>lsp>auto-bandwidth** context. The

default value of sample-multiplier (the value that would result from the no auto-bandwidth-defaults command) is 1, which means the default sample interval is 300 seconds.

Over a longer period of time called the adjust interval the router keeps track of the maximum average data rate recorded during any constituent sample interval. The adjust interval is the product of adjust-multiplier and the collection-interval specified in the auto-bandwidth accounting-policy. A default adjust-multiplier for all LSPs may be configured using the **config>router>mpls>auto-bandwidth-multiplier** command but this value can be overridden on a per-LSP basis at the **config>router>mpls>lsp>auto-bandwidth** context. The default value of adjust-multiplier (the value that would result from the no auto-bandwidth-multiplier command) is 288, which means the default adjust interval is 86400 seconds or 24 hours. The system enforces the restriction that adjust-multiplier is equal to or greater than sample-multiplier. It is recommended that the **adjust-multiplier** be an integer multiple of the **sample-multiplier**.

The collection-interval in the auto-bandwidth accounting policy can be changed at any time, without disabling any of the LSPs that rely on that policy for statistics collection.

The sample-multiplier (at the **mpls>auto-bandwidth** level or the **lsp>auto-bandwidth** level) can be changed at any time. This will have no effect until the beginning of the next sample interval. In this case the adjust-interval does not change and information about the current adjust interval (such as the remaining adjust-multiplier, the maximum average data rate) is not lost when the sample-multiplier change takes effect.

The system allows adjust-multiplier (at the **mpls** level or the **lsp>auto-bandwidth** level) to be changed at any time as well but in this case the new value shall have no effect until the beginning of the next adjust interval.

Byte counts collected for LSP statistics include layer 2 encapsulation (Ethernet headers and trailers) and therefore average data rates measured by this feature include Layer 2 overhead as well.

Passive Monitoring of LSP Bandwidth

The system offers the option to measure the bandwidth of an RSVP LSP (see [Measurement of LSP Bandwidth on page 34](#)) without taking any action to adjust the bandwidth reservation, regardless of how different the measured bandwidth is from the current reservation. Passive monitoring is enabled using the **config>router>mpls>lsp>auto-bandwidth>monitor-bandwidth** command.

The **show>router>mpls>lsp detail** command can be used to view the maximum average data rate in the current adjust interval and the remaining time in the current adjust interval.

Periodic Automatic Bandwidth Adjustment

Automatic bandwidth allocation is supported on any RSVP LSP that has MBB enabled. MBB is enabled in the `config>router>mpls>lsp` context using the **adaptive** command. For automatic adjustments of LSP bandwidth to occur the `monitor-bandwidth` command must not be present at `config>router>mpls>lsp>auto-bandwidth` context, otherwise only passive measurements will occur.

If an eligible RSVP LSP is configured for auto-bandwidth, by entering `auto-bandwidth` at the `config>router>mpls>lsp` context, then the ingress LER decides every adjust interval whether to attempt auto-bandwidth adjustment. The following parameters are defined:

- `current_bw` — The currently reserved bandwidth of the LSP; this is the operational bandwidth that is already maintained in the MIB.
- `measured_bw` — The maximum average data rate in the current adjust interval.
- `signaled_bw` — The bandwidth that is provided to the CSPF algorithm and signaled in the `SENDER_TSPEC` and `FLOWSPEC` objects when an auto-bandwidth adjustment is attempted.
- `min` — The configured min-bandwidth of the LSP.
- `max` — The configured max-bandwidth of the LSP.
- `up%` — The minimum difference between `measured_bw` and `current_bw`, expressed as a percentage of `current_bw`, for increasing the bandwidth of the LSP.
- `up` — The minimum difference between `measured_bw` and `current_bw`, expressed as an absolute bandwidth relative to `current_bw`, for increasing the bandwidth of the LSP. This is an optional parameter; if not defined the value is 0.
- `down%` — The minimum difference between `current_bw` and `measured_bw`, expressed as a percentage of `current_bw`, for decreasing the bandwidth of the LSP.
- `down` — The minimum difference between `current_bw` and `measured_bw`, expressed as an absolute bandwidth relative to `current_bw`, for decreasing the bandwidth of the LSP. This is an optional parameter; if not defined the value is 0.

At the end of every adjust interval the system decides if an auto-bandwidth adjustment should be attempted. The heuristics are as follows:

- If the measured bandwidth exceeds the current bandwidth by more than the percentage threshold and also by more than the absolute threshold then the bandwidth is re-signaled to the measured bandwidth (subject to min and max constraints).
- If the measured bandwidth is less than the current bandwidth by more than the percentage threshold and also by more than the absolute threshold then the bandwidth is re-signaled to the measured bandwidth (subject to min and max constraints).

- If the current bandwidth is greater than the max bandwidth then the LSP bandwidth is re-sigaled to max bandwidth, even if the thresholds have not been triggered.
- If the current bandwidth is greater than the min bandwidth then the LSP bandwidth is re-sigaled to min bandwidth, even if the thresholds have not been triggered.

Changes to min-bandwidth, max-bandwidth and any of the threshold values (up, up%, down, down%) are permitted at any time on an operational LSP but the changes have no effect until the next auto-bandwidth trigger (for example, adjust interval expiry).

If the measured bandwidth exceeds the current bandwidth by more than the percentage threshold and also by more than the absolute threshold then the bandwidth is re-sigaled to the measured bandwidth (subject to min and max constraints).

The adjust-interval and maximum average data rate are reset whether the adjustment succeeds or fails. If the bandwidth adjustment fails (for example, CSPF cannot find a path) then the existing LSP is maintained with its existing bandwidth reservation. The system does not retry the bandwidth adjustment (for example, per the configuration of the LSP retry-timer and retry-limit).

Overflow-Triggered Auto-Bandwidth Adjustment

For cases where the measured bandwidth of an LSP has increased significantly since the start of the current adjust interval it may be desirable for the system to preemptively adjust the bandwidth of the LSP and not wait until the end of the adjust interval.

The following parameters are defined:

- `current_bw` — The currently reserved bandwidth of the LSP.
- `sampled_bw` — The average data rate of the sample interval that just ended.
- `measured_bw` — The maximum average data rate in the current adjust interval.
- `signaled_bw` — The bandwidth that is provided to the CSPF algorithm and signaled in the `SENDER_TSPEC` and `FLOWSPEC` objects when an auto-bandwidth adjustment is attempted.
- `max` — The configured max-bandwidth of the LSP.
- `%_threshold` — The minimum difference between `sampled_bw` and `current_bw`, expressed as a percentage of the `current_bw`, for counting an overflow event.
- `min_threshold` — The minimum difference between `sampled_bw` and `current_bw`, expressed as an absolute bandwidth relative to `current_bw`, for counting an overflow event. This is an optional parameter; if not defined the value is 0.

When a sample interval ends it is counted as an overflow if:

- The sampled bandwidth exceeds the current bandwidth by more than the percentage threshold and by more than the absolute bandwidth threshold (if defined).
- When the number of overflow samples reaches a configured limit, an immediate attempt is made to adjust the bandwidth to the measured bandwidth (subject to the min and max constraints).

If the bandwidth adjustment is successful then the adjust-interval, maximum average data rate and overflow count are all reset. If the bandwidth adjustment fails then the overflow count is reset but the adjust-interval and maximum average data rate continue with current values. It is possible that the overflow count will once again reach the configured limit before the end of adjust-interval is reached and this will once again trigger an immediate auto-bandwidth adjustment attempt.

The overflow configuration command fails if the max-bandwidth of the LSP has not been defined.

The threshold limit can be changed on an operational auto-bandwidth LSP at any time and the change should take effect at the end of the current sample interval (for example, if the user decreases the overflow limit to a value lower than the current overflow count then auto-bandwidth adjustment will take place as soon as the sample interval ends). The threshold values can also be changed at any time (for example, `%_threshold` and `min_threshold`) but the new values will not take effect until the end of the current sample interval.

Manually-Triggered Auto-Bandwidth Adjustment

Manually-triggered auto-bandwidth adjustment feature is configured with the **tools>perform>router>mpls adjust-autobandwidth [lsp *lsp-name* [force [bandwidth *mbps*]]]** command to attempt immediate auto-bandwidth adjustment for either one specific LSP or all active LSPs. If the LSP is not specified then the system assumes the command applies to all LSPs. If an LSP name is provided then the command applies to that specific LSP only and the optional **force** parameter (with or without a bandwidth) can be used.

If **force** is not specified (or the command is not LSP-specific) then `measured_bw` is compared to `current_bw` and bandwidth adjustment may or may not occur

If **force** is specified and a bandwidth is not provided then the threshold checking is bypassed but the min and max bandwidth constraints are still enforced.

If **force** is specified with a bandwidth (in Mbps) then `signaled_bw` is set to this bandwidth. There is no requirement that the bandwidth entered as part of the command fall within the range of min-bandwidth to max-bandwidth.

The `adjust-interval`, maximum average data rate and overflow count are not reset by the manual auto-bandwidth command, whether or not the bandwidth adjustment succeeds or fails. The overflow count is reset only if the manual auto-bandwidth adjustment is successful.

MPLS Transport Profile (MPLS-TP)

MPLS can be used to provide a network layer to support packet transport services. In some operational environments, it is desirable that the operation and maintenance of such an MPLS based packet transport network follow operational models typical in traditional optical transport networks (e.g. SONET/SDH), while providing additional OAM, survivability and other maintenance functions targeted at that environment.

MPLS-TP defines a profile of MPLS targeted at transport applications. This profile defines the specific MPLS characteristics and extensions required to meet transport requirements, while retaining compliance to the standard IETF MPLS architecture and label switching paradigm. The basic requirements for MPLS-TP are described by the IETF in RFC 5654, RFC 5921 and RFC 5960, in order to meet two objectives:

1. To enable MPLS to be deployed in a transport network and operated in a similar manner to existing transport technologies.
2. To enable MPLS to support packet transport services with a similar degree of predictability to that found in existing transport networks.

In order to meet these objectives, MPLS-TP has a number of high level characteristics:

- It does not modify the MPLS forwarding architecture, which is based on existing pseudowire and LSP constructs. Point-to-point LSPs may be unidirectional or bi-directional. Bi-directional LSPs must be congruent (i.e. co-routed and follow the same path in each direction). The system supports bidirectional co-routed MPLS-TP LSPs.
- There is no LSP merging.
- OAM, protection and forwarding of data packets can operate without IP forwarding support. When static provisioning is used, there is no dependency on dynamic routing or signaling.
- LSP and pseudowire monitoring is only achieved through the use of OAM and does not rely on control plane or routing functions to determine the health of a path. e.g. LDP hello failures, do not trigger protection.
- MPLS-TP can operate in the absence of an IP control plane and IP forwarding of OAM traffic. In release 11.0, MPLS-TP is only supported on static LSPs and PWs.

The system supports MPLS-TP on LSPs and PWs with static labels. MPLS-TP is not supported on dynamically signalled LSPs and PWs. MPLS-TP is supported for EPIPE, APIPE and CPIPE VLLs, and EPIPE Spoke SDP termination on IES, VPRN and VPLS. Static PWs may use SDPs that use either static MPLS-TP LSPs or RSVP-TE LSPs.

The following MPLS-TP OAM and protection mechanisms, defined by the IETF, are supported:

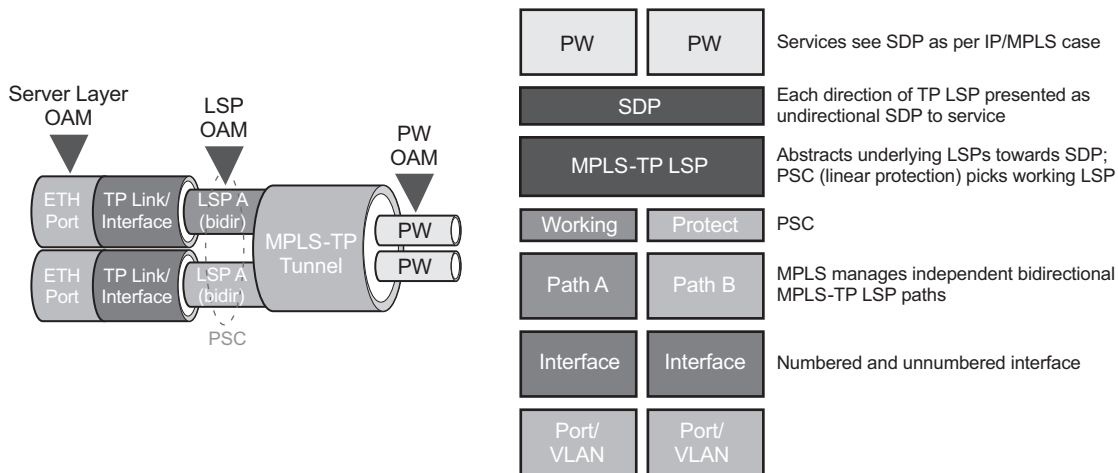
- MPLS-TP Generic Associated Channel for LSPs and PWs (RFC 5586)

- MPLS-TP Identifiers (RFC 6370)
- Proactive CC, CV, and RDI using BFD for LSPs (RFC 6428)
- On-Demand CV for LSPs and PWs using LSP Ping and LSP Trace (RFC 6426)
- 1-for-1 Linear protection for LSPs (RFC 6378)
- Static PW Status Signaling (RFC 6478)

The system can play the role of an LER and an LSR for static MPLS-TP LSPs, and a PE/T-PE and an S-PE for static MPLS-TP PWs. It can also act as a S-PE for MPLS-TP segments between an MPLS network that strictly follows the transport profile, and an MPLS network that supports both MPLS-TP and dynamic IP/MPLS.

MPLS-TP Model

Figure 5 shows a high level functional model for MPLS-TP in SROS. LSP A and LSP B are the working and protect LSPs of an LSP tunnel. These are modelled as working and protect paths of an MPLS-TP LSP in SROS. MPLS-TP OAM runs in-band on each path. 1:1 linear protection coordinates the working and protect paths, using a protection switching coordination protocol (PSC) that runs in-band on each path over a Generic Associated Channel (G-ACh) on each path. Each path can use either an IP numbered, IP unnumbered, or MPLS-TP unnumbered (i.e. non-IP) interface.



al_0221

Figure 5: MPLS-TP Model

Note that in SR OS, all MPLS-TP LSPs are bidirectional co-routed, as detailed in RFC5654. That is, the forward and backward directions follow the same route (in terms of links and nodes) across the network. Both directions are setup, monitored and protected as a single entity. Therefore, both

ingress and egress directions of the same LSP segment are associated at the LER and LSR and use the same interface (although this is not enforced by the system).

In the above model, an SDP can use one MPLS-TP LSP. This abstracts the underlying paths towards the overlying services, which are transported on pseudowires. Pseudowires are modelled as spoke SDPs and can also use MPLS-TP OAM. PWs with static labels may use SDPs that in-turn use either signaled RSVP-TE LSPs, or one static MPLS-TP LSP.

MPLS-TP Provider Edge and Gateway

This section describes some example roles for the system in an MPLS-TP network.

VLL Services

The system may use MPLS TP LSPs, and PWs, to transport point to point virtual leased line services. The 7750 may play the role of a terminating PE or switching PE for VLLs. Epipe, Apipe and Cpipe VLLs are supported.

Figure 6 illustrates the use of the system as a T-PE for services in an MPLS-TP domain, and as a S-PE for services between

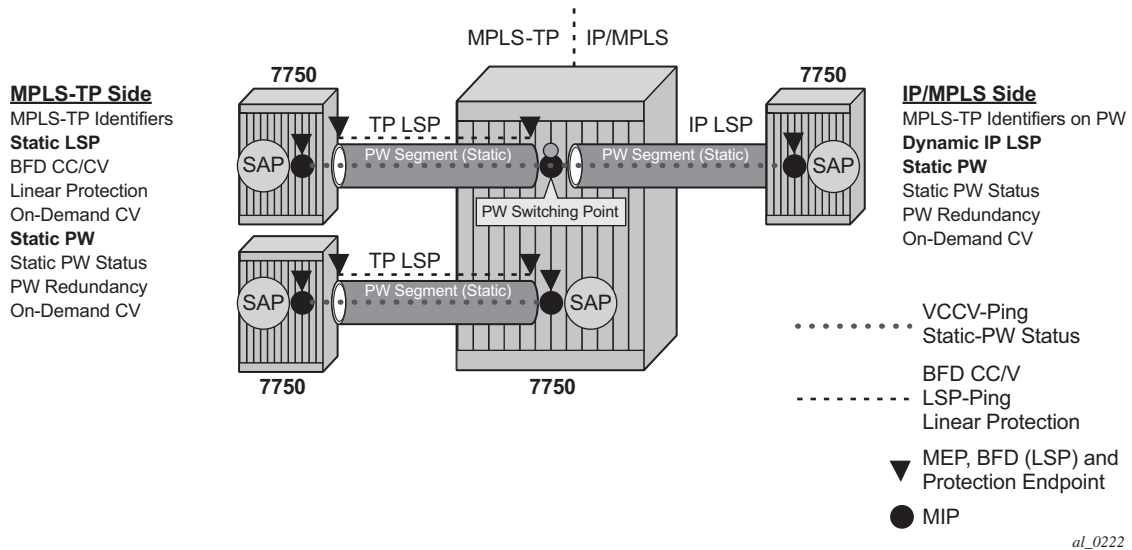


Figure 6: MPLS-TP Provider Edge and Gateway, VLL Services

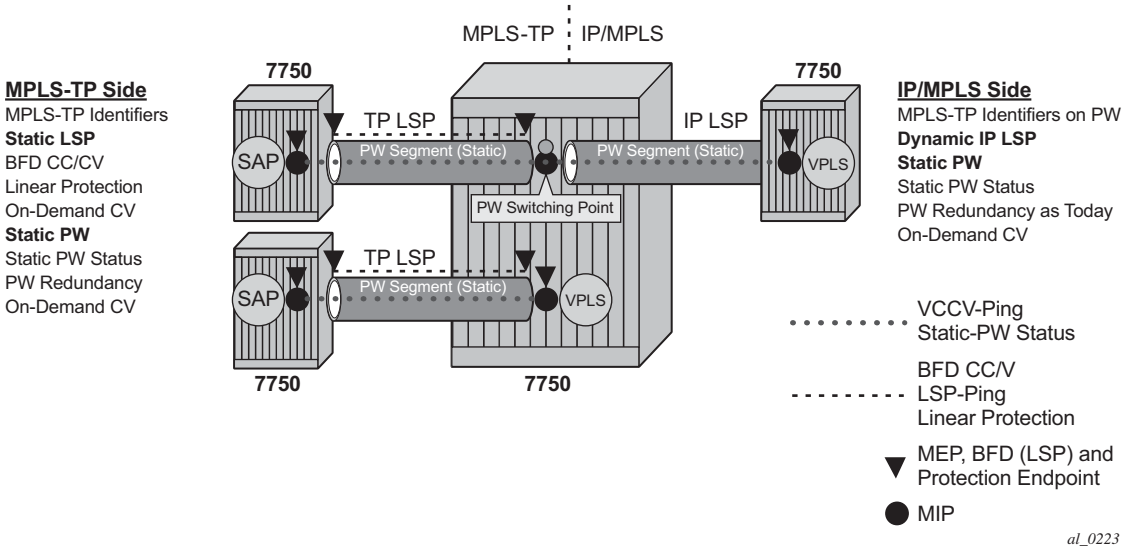


Figure 7: MPLS-TP Provider Edge and Gateway, spoke-SDP Termination on VPLS




Figure 8: MPLS-TP Provider Edge and Gateway, spoke-SDP Termination on IES/VPRN




Figure 9: MPLS-TP LSR

Detailed Descriptions of MPLS-TP

MPLS-TP LSPs

SR OS supports the configuration of MPLS-TP tunnels, which comprise a working and, optionally, a protect LSP. In SROS, a tunnel is referred to as an LSP, while an MPLS-TP LSP is referred to as a path. It is then possible to bind an MPLS-TP tunnel to an SDP.

MPLS-TP LSPs (i.e. paths) with static labels are supported. MPLS-TP is not supported for signaled LSPs.

Both bidirectional associated (where the forward and reverse directions of a bidirectional LSP are associated at a given LER, but may take different routes through the intervening network) and bidirectional co-routed (where the forward and reverse directions of the LSP are associated at each LSR, and take the same route through the network) are possible in MPLS-TP. However, only bidirectional co-routed LSPs are supported.

It is possible to configure MPLS-TP identifiers associated with the LSP, and MPLS-TP OAM parameters on each LSP of a tunnel. MPLS-TP protection is configured for a tunnel at the level of the protect path level. Both protection and OAM configuration is managed via templates, in order to simplify provisioning for large numbers of tunnels.

The 7750 may play the role of either an LER or an LSR.

MPLS-TP on Pseudowires

MPLS-TP is supported on PWs with static labels. The provisioning model supports RFC6370-style PW path identifiers for MPLS-TP PWs.

MPLS-TP PWs reuse the static PW provisioning model of previous SR OS releases. Including the use of the PW-switching key work to distinguish an S-PE. Therefore, the primary distinguishing feature for an MPLS-TP PW is the ability to configure MPLS-TP PW path identifiers, and to support MPLS-TP OAM and static PW status signaling.

The system can perform the role of a T-PE or an S-PE for a PW with MPLS-TP.

A spoke-SDP with static PW labels and MPLS-TP identifiers and OAM capabilities can use an SDP that uses either an MPLS-TP tunnel, or that uses regular RSVP-TE LSPs. The control word is supported for all MPLS-TP PWs.

MPLS-TP Maintenance Identifiers

MPLS-TP is designed for use both with, and without, a control plane. MPLS-TP therefore specifies a set of identifiers that can be used for objects in either environment. This includes a path and maintenance identifier architecture comprising Node, Interface, PW and LSP identifiers, Maintenance Entity Groups (MEGs), Maintenance End Points (MEPs) and Maintenance Intermediate Points (MIPs). These identifiers are specified in RFC6370.

MPLS-TP OAM and protection switching operates within a framework that is designed to be similar to existing transport network maintenance architectures. MPLS-TP introduces concept of maintenance domains to be managed and monitored. In these, Maintenance Entity Group End Points (MEPs) are edges of a maintenance domain. OAM of a maintenance level must not leak beyond corresponding MEP and so MEPs typically reside at the end points of LSPs and PWs. Maintenance Intermediate Points (MIPs) define intermediate nodes to be monitored. Maintenance Entity Groups (MEGs) comprise all the MEPs and MIPs on an LSP or PW.




Figure 10: MPLS-TP Maintenance Architecture

Both IP-compatible and ICC (ITU-T carrier code) based identifiers for the above objects are specified in the IETF, but only the IP-compatible identifiers defined in RFC6370 are supported.

SROS supports the configuration of the following node and interface related identifiers:


- **Global_ID:** this is similar to the global ID that can be configured for Dynamic MS-PWs in Release 9.0 of SR OS. However, in MPLS-TP this should be set to the AS# of the node. If not explicitly configured, then it assumes the default value of 0. In SR OS, the source Global ID for an MPLS-TP Tunnel is taken to be the Global ID configured at the LER. The destination Global ID is optional in the tunnel configuration. If it is not configured, then it is taken as the same as the source Global ID.
- **Node_ID:** This is a 32-bit value assigned by the operator within the scope of the Global_ID. The system supports the configuration of an IPv4 formatted address <a.b.c.d> or an unsigned 32-bit integer for the MPLS-TP Node ID at each node. The node ID must be unique within the scope of the global ID, but there is no requirement for it to be a valid routable IP address. Indeed, a node-id can represent a separate IP-compatible addressing space that may be separate from the IP addressing plan of the underlying network. If no node ID is configured, then the node ID is taken to be the system interface IPv4 address of the node. When configuring a tunnel at an LER, either an IPv4 or an unsigned integer Node ID can be configured as the source and destination identifiers, but both ends must be of the same type.
- **IF_ID:** This is an MPLS-TP section layer identifier at the MPLS interface level. On the 7x50, this is used to provide an identifier for the LSP-Trace DSMAP when an IP identifier is not available.. The IF_ID is a 64-bit identifier of an MPLS-TP interface on a node that is unique within the scope of a Global_ID. It is composed of the Node_ID and the IF_Num. The IF_Num is a node-wide unique identifier for an MPLS-TP interface. On the 7x50, this is primarily used for supporting the DSMAP TLV in LSP Trace using MPLS-TP identifiers with unnumbered MPLS-TP interfaces.

Statically configured LSPs are identified using GMPLS-compatible identifiers with the addition of a Tunnel_Num and LSP_Num. As in RSVP-TE, tunnels represent, for example, a set of working and protect LSPs. These are GMPLS-compatible because GMPLS chosen by the IETF as the

control plane for MPLS-TP LSPs, although this is not supported in Release 11.0 of the 7750. PWs are identified using a PW Path ID which has the same structure as FEC129 AII Type 2.

SR OS derives the identifiers for MEPs and MIPs on LSPs and PWs based on the configured identifiers for the MPLS-TP Tunnel, LSP or PW Path ID, for use in MPLS-TP OAM and protection switching, as per RFC6370.

The information models for LSPs and PWs supported in Release 11.0 are illustrated in Figure 11 and Figure 12. The figures use the terminology defined in RFC6370.



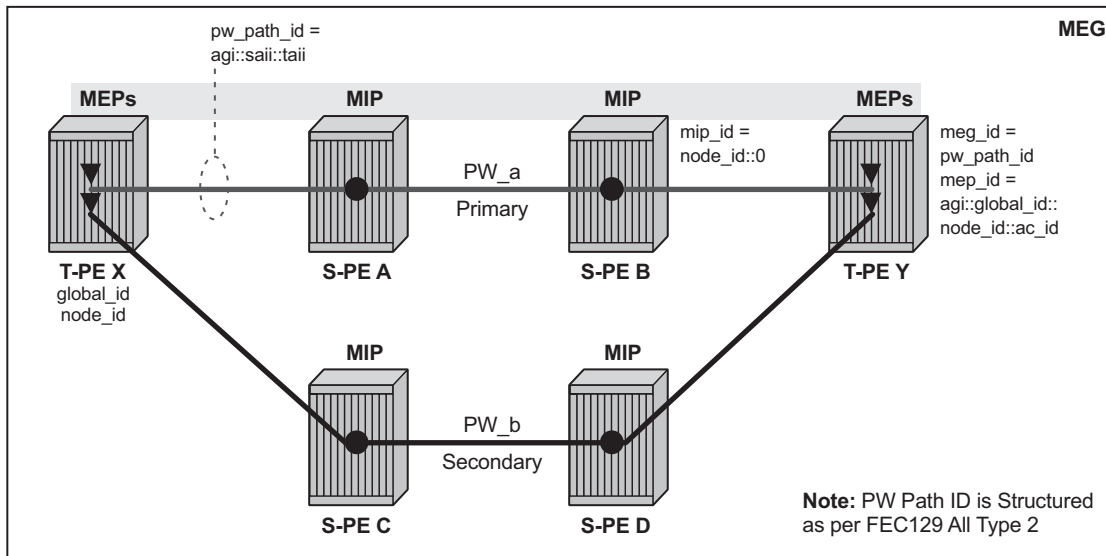
al_0227

Figure 11: MPLS-TP LSP and Tunnel Information Model

The MPLS-TP Tunnel ID and LSP ID are not to be confused with the RSVP-TE tunnel id implemented on the 7x50 system. Table 3 shows how these map to the X and Y ends of the tunnel shown in the above figure for the case of co-routed bidirectional LSPs.

Table 3: Mapping from RSVP-TE to MPLS-TP Maintenance Identifiers

RSVP-TE Identifier	MPLS-TP Maintenance Identifier
Tunnel Endpoint Address	Node ID (Y)
Tunnel ID (X)	Tunnel Num (X)
Extended Tunnel ID	Node ID (X)
Tunnel Sender Address	Node ID (X)
LSP ID	LSP Num



al_0228

Figure 12: MPLS-TP PW Information Model

In the PW information model shown in [Figure 12](#), the MS-PW is identified by the PW Path ID that is composed of the full AGI:SAII:TAII. The PW Path ID is also the MEP ID at the T-PEs, so a user does not have to explicitly configure a MEP ID; it is automatically derived by the system. For MPLS-TP PWs with static labels, although the PW is not signaled end-to-end, the directionality of the SAII and TAII is taken to be the same as for the equivalent label mapping message i.e. from downstream to upstream. This is to maintain consistency with signaled pseudowires using FEC 129.

On the system, an S-PE for an MS-PW with static labels is configured as a pair of spoke-sdps bound together in an VLL service using the VC-switching command. Therefore, the PW Path ID configured at the spoke-SDP level at an S-PE must contain the Global-ID, Node-ID and AC-ID at the far end T-PEs, not the local S-PE. Note that the ordering of the SAII:TAII in the PW Path ID where static PWs are used should be consistent with the direction of signaling of the egress label to a spoke-SDP forming that segment, if that label were signaled using T-LDP (in downstream unsolicited mode). VCCV Ping will check the PW ID in the VCCV Ping echo request message against the configured PW Path ID for the egress PW segment.

[Figure 13](#) shows an example of how the PW Path IDs can be configured for a simple two-segment MS-PW.




Figure 13: Example usage of PW Identifiers

Generic Associated Channel

MPLS-TP requires that all OAM traffic be carried in-band on both directions of an LSP or PW. This is to ensure that OAM traffic always shares fate with user data traffic. This is achieved by using an associated control channel on an LSP or PW, similar to that used today on PWs. This creates a channel, which is used for OAM, protection switching protocols (e.g. LSP linear protection switching coordination), and other maintenance traffic., and is known as the Generic Associated Channel (G-ACh).

RFC5586 specifies mechanisms for implementing the G-ACh, relying on the combination of a reserved MPLS label, the 'Generic-ACH Label (GAL)', as an alert mechanism (value=13) and Generic Associated Channel Header (G-ACH) for MPLS LSPs, and using the Generic Associated Channel Header, only, for MPLS PWs (although the GAL is allowed on PWs). The purpose of the GAL is to indicate that a G-ACh resides at the bottom of the label stack, and is only visible when the bottom non-reserved label is popped. The G-ACH channel type is used to indicate the packet type carried on the G-ACh. Packets on a G-ACh are targeted to a node containing a MEP by ensuring that the GAL is pushed immediately below the label that is popped at the MEP (e.g. LSP endpoint or PW endpoint), so that it can be inspected as soon as the label is popped. A G-ACh packet is targeted to a node containing a MIP by setting the TTL of the LSP or PW label, as applicable, so that it expires at that node, in a similar manner to the SROS implementation of VCCV for MS-PWs.

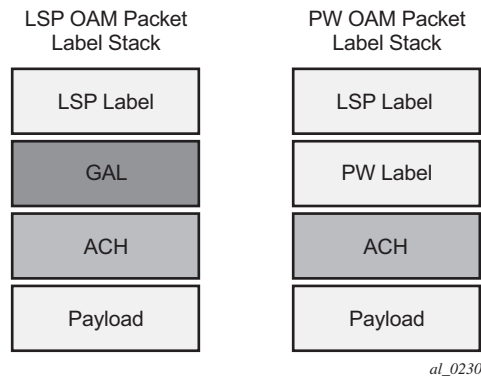


Figure 14: Label for LSP and PW G-ACh Packets

The system supports the G-ACh on static pseudowires and static LSPs.

MPLS-TP Operations, Administration and Maintenance (OAM)

This section details the MPLS-TP OAM mechanisms that are supported.

On-Demand Connectivity Verification (CV) using LSP-Ping

MPLS-TP supports mechanisms for on demand CC/CV as well as route tracing for LSPs and PWs. These are required to enable an operator to test the initial configuration of a transport path, or to assist with fault isolation and diagnosis. On demand CC/CV and route tracing for MPLS-TP is based on LSP-Ping and is described in RFC6426. Three possible encapsulations are specified in that RFC:

- IP encapsulation, using the same label stack as RFC4379, or encapsulated in the IPv4 G-ACh channel with a GAL/ACH
- and non-IP encapsulation with GAL/ACH for LSPs and ACH for PWs.

In IP-encapsulation, LSP-Ping packets are sent over the MPLS LSP for which OAM is being performed and contain an IP/UDP packet within them. The On-demand CV echo response message is sent on the reverse path of the LSP, and the reply contains IP/UDP headers followed by the On-demand CV payload.

In non-IP environments, LSP ping can be encapsulated with no IP/UDP headers in a G-ACh and use a source address TLV to identify the source node, using forward and reverse LSP or PW associated channels on the same LSP or PW for the echo request and reply packets. In this case, no IP/UDP headers are included in the LSP-Ping packets.

The 7750 support the following encapsulations:

- IP encapsulation with ACH for PWs (as per VCCV type 1).
- IP encapsulation without ACH for LSPs using labeled encapsulation
- Non-IP encapsulation with ACH for both PWs and LSPs.

LSP Ping and VCCV Ping for MPLS-TP use two new FEC sub-types in the target FEC stack in order to identify the static LSP or static PW being checked. These are the Static LSP FEC sub-type, which has the same format as the LSP identifier described above, and the Static PW FEC sub-type,. These are used in-place of the currently defined target FEC stack sub-TLVs.

In addition, MPLS-TP uses a source/destination TLV to carry the MPLS-TP global-id and node-id of the target node for the LSP ping packet, and the source node of the LSP ping packet.

LSP Ping and VCCV-Ping for MPLS-TP can only be launched by the LER or T-PE. The replying node therefore sets the TTL of the LSP label or PW label in the reply packet to 255 to ensure that it reaches the node that launched the LSP ping or VCCV Ping request.

Downstream Mapping Support

RFC 4379 specifies four address types for the downstream mapping TLV for use with IP numbered and unnumbered interfaces:

Type #	Address Type	K Octets	Reference
1	IPv4 Numbered	16	RFC 4379
2	IPv4 Unnumbered	16	RFC 4379
3	IPv6 Numbered	40	RFC 4379
4	IPv6 Unnumbered	28	RFC 4379

RFC 6426 adds address type 5 for use with Non IP interfaces, including MPLS-TP interfaces. In addition, this RFC specifies that type 5 must be used when non-IP ACH encapsulation is used for LSP Trace.

It is possible to send and respond to a DSMAP/DDMAP TLV in the LSP Trace packet for numbered IP interfaces as per RFC4379. In this case, the echo request message contains a downstream mapping TLV with address type 1 (IPv4 address) and the IPv4 address in the DDMAP/DSMAP TLV is taken to be the IP address of the IP interface that the LSP uses. The LSP

trace packet therefore contains a DSMAP TLV in addition to the MPLS-TP static LSP TLV in the target FEC stack.

DSMAP/DDMAP is not supported for pseudo wires.

Proactive CC, CV and RDI

Proactive Continuity Check (CC) is used to detect a loss of continuity defect (LOC) between two MEPs in a MEG. Proactive Connectivity Verification (CV) is used to detect an unexpected connectivity defect between two MEPs (e.g. mis-merging or mis-connection), as well as unexpected connectivity within the MEG with an unexpected MEP. This feature implements both functions using proactive generation of OAM packets by the source MEP that are processed by the peer sink MEP. CC and CV packets are always sent in-band such that they fate share with user traffic, either on an LSP, PW or section and are used to trigger protection switching mechanisms.

Proactive CC/CV based on bidirectional forwarding detection (BFD) for MPLS-TP is described in RFC6428. BFD packets are sent using operator configurable timers and encapsulated without UDP/IP headers on a standardized G-ACh channel on an LSP or PW. CC packets simply consist of a BFD control packet, while CV packets also include an identifier for the source MEP in order that the sink MEP can detect if it is receiving packets from an incorrect peer MEP, thus indicating a mis-connectivity defect. Other defect types (including period mis-configuration defect) should be supported. When a supported defect is detected, an appropriate alarm is generated (e.g. log, SNMP trap) at the receiving MEP and all traffic on the associated transport path (LSP or PW) is blocked. This is achieved using linear protection for CC defects, and by blocking the ingress data path for CV defects. The system supports both a CC-only mode and a combine CC / CV mode, as defined in RFC6428.

Note that when an LSP with CV is first configured, the LSP will be held in the CV defect state for 3.5 seconds after the first valid CV packet is received.

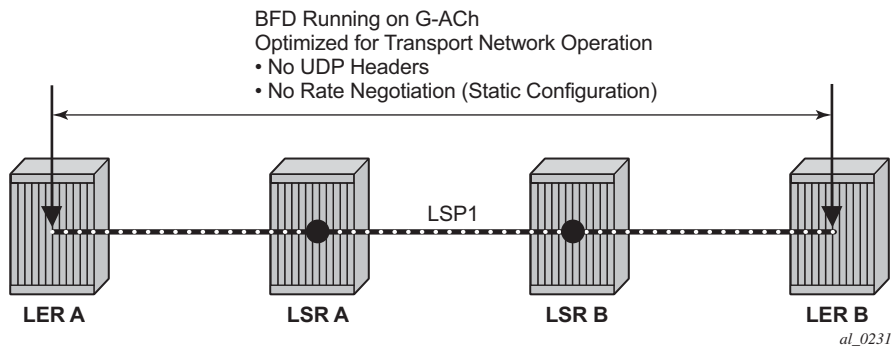



Figure 15: BFD used for proactive CC on MPLS-TP LSP



al_0232

Figure 16: BFD used for proactive CV on MPLS-TP LSP

Linear protection switching of LSPs (see below) is triggered based on a CC or CV defect detected by BFD CC/CV.

Note that RFC6428 defines two BFD session modes: Coordinated mode, in which the session state on both directions of the LSP is coordinated and constructed from a single, bidirectional BFD session, and independent mode, in which two independent sessions are bound together at a MEP. Coordinated mode is supported.

BFD is supported on MPLS-TP LSPs. When BFD_CV detects a mis-connectivity on an LSP, the system will drop all incoming non-OAM traffic with the LSP label (at the LSP termination point) instead of forwarding it to the associated SAP or PW segment.

The following GCh channel types are supported for the combined CC/CV mode:

- 0x22 for BFD CC with no IP encapsulation
- 0x23 for BFD CV

The following G-ACh channel types are used for the CC-only mode:

- 0x07

BFD-based RDI

RDI provides a mechanism whereby the source MEP can be informed of a downstream failure on an LSP, and can thus either raise an alarm, or initiate a protection switching operation. In the case of BFD based CC/CV, RDI is communicated using the BFD diagnostic field in BFC CC/CV messages. The following diagnostic codes are supported:

1 - Control Detection Time Expired

9 - mis-connectivity defect

PW Control Channel Status Notifications (Static Pseudowire Status Signaling)

MPLS-TP introduces the ability to support a full range of OAM and protection / redundancy on PWs for which no dynamic T-LDP control plane exists. Static PW status signaling is used to advertise the status of a PW with statically configured labels by encapsulating the PW status TLV in a G-ACh on the PW. This mechanism enables OAM message mapping and PW redundancy for such PWs, as defined in RFC6478. This mechanism is known as control channel status signaling in SR OS.

PW control channel status notifications use a similar model to T-LDP status signaling. That is, in general, status is always sent to the nearest neighbor T-PE or S-PE and relayed to the next segment by the S-PE. To achieve this, the PW label TTL is set to 1 for the G-ACh packet containing the status message.

Control channel status notifications are disabled by default on a spoke-SDP. If they are enabled, then the default refresh interval is set to zero (although this value should be configurable in CLI). That is, when a status bit changes, three control channel status packets will be sent consecutively at one-second intervals, and then the transmitter will fall silent. If the refresh timer interval is non-zero, then status messages will continue to be sent at that interval. The system supports the configuration of a refresh timer of 0, or from 10-65535 seconds. The recommended value is 600 seconds.

The system supports the optional acknowledgement of a PW control channel status message.

In order to constrain the CPU resources consumed processing control channel status messages, the system implements a credit-based mechanism. If a user enables control channel status on a PW[n], then a certain number of credits c_n are consumed from a CPM-wide pool of max_credit credits. The number of credits consumed is inversely proportional to the configured refresh timer (the first three messages at 1 second interval do not count against the credit). If the $current_credit \leq 0$, then control channel status signaling cannot be configured on a PW (but the PW can still be configured and no shutdown).

If a PE with a non-zero refresh timer configured does not receive control channel status refresh messages for 3.5 time the specified timer value, then by default it will time out and assume a PW status of zero.

A trap is generated if the refresh timer times-out.

If PW redundancy is configured, the system will always consider the literal value of the PW status; a time-out of the refresh timer will not impact the choice of the active transit object for the VLL service. The result of this is that if the refresh timer times-out, and a given PW is currently the active PW, then the system will not fail-over to an alternative PW if the status is zero and some lower-layer OAM mechanism e.g. BFD has not brought down the LSP due to a connectivity defect. It is recommended that the PW refresh timer be configured with a much longer interval than any proactive OAM on the LSP tunnel, so that the tunnel can be brought down before the refresh timer expires if there is a CC defect.

Note that a unidirectional continuity fault on a RSVP TE LSP may not result in the LSP being brought down before the received PW status refresh timer expires. It is therefore recommended that either bidirectional static MPLS-TP LSPs with BFD CC, or additional protection mechanisms e.g. FRR be used on RSVP-TE LSPs carrying MPLS-TP PWs. This is particularly important in active/standby PW dual homing configurations, where the active / standby forwarding state or operational state of every PW in the redundancy set must be accurately reflected at the redundant PE side of the configuration.

Note that a PW with a refresh timer value of zero is always treated as having not expired.

The system implements a hold-down timer for control-channel-status PW-status bits in order to suppress bouncing of the status of a PW. For a specific spoke-SDP, if the system receives 10 PW-status *change* events in 10 seconds, the system will *hold-down* the spoke-SDP on the local node with the last received non-zero PW-status bits for 20 seconds. It will update the local spoke with the most recently received PW-status. This hold down timer is not persistent across shutdown/no-shutdown events.

PW Control Channel Status Request Mechanism

The system implements an optional PW control channel status request mechanism. This enhances the existing control channel status mechanism so that a peer that has *stale* PW status for the far-end of a PW can request that the peer PE send a static PW status update. Accurate and current information about the far end status of a PW is important for proper operation of PW redundancy. This mechanism ensures a consistent view of the control plane is maintained, as far as possible, between peer nodes. It is not intended to act as a continuity check between peer nodes.

Pseudowire Redundancy and Active / Standby Dual Homing

PW redundancy is supported for static MPLS-TP pseudowires. However, instead of using T-LDP status signaling to signal the forwarding state of a PW, control channel status signaling is used.

The following PW redundancy scenarios must be supported:

- MC-LAG and MC-APS with single and multi-segment PWs interconnecting the PEs.
- MS-PW (S-PE) Redundancy between VLL PEs with single-homed CEs.
- Dual-homing of a VLL service into redundant IES or VPRN PEs, with active/standby PWs.
- Dual-homing of a VLL service into a VPLS with active/standby PWs.

Note that active/standby dual-homing into routed VPLS is not supported in for MPLS-TP PWs. This is because it relies on PW label withdrawal of the standby PW in order to take down the VPLS instance, and hence the associated IP interface. Instead, it is possible to enable BGP multi-homing on a routed VPLS that has MPLS-TP PWs as spokes, and for the PW status of each spoke-SDP to be driven (using control channel status) from the active or standby forwarding state assigned to each PW by BGP.

It is possible to configure inter-chassis backup (ICB) PWs as static MPLS-TP PWs with MPLS-TP identifiers. Only MPLS-TP PWs are supported in the same endpoint. That is, PWs in an endpoint must either be all MPLS-TP, or none of them must be MPLS-TP. This implies that an ICB used in an endpoint for which other PWs are MPLS TP must also be configured as an MPLS-TP PW.

A failover to a standby pseudowire is initiated based on the existing supported methods (e.g. failure of the SDP).

MPLS-TP LSP Protection

Linear 1-for-1 protection of MPLS-TP LSPs is supported, as defined in RFC. This applies only to LSPs (not PWs).

This is supported edge-to-edge on an LSP, between two LERs, where normal traffic is transported either on the working LSP or on the protection LSP using a logical selector bridge at the source of the protected LSP.

At the sink LER of the protected LSP, the LSP that carries the normal traffic is selected, and that LSP becomes the working LSP. A protection switching coordination (PSC) protocol coordinates between the source and sink bridge, which LSP will be used, as working path and protection path. The PSC protocol is always carried on a G-ACh on the protection LSP.

The system supports single-phased coordination between the LSP endpoints, in which the initiating LER performs the protection switchover to the alternate path and informs the far-end LER of the switch.

Bidirectional protection switching is achieved by the PSC protocol coordinating between the two end points to determine which of the two possible paths (i.e. the working or protect path), transmits user traffic at any given time.

It is possible to configure non-revertive or revertive behavior. For non-revertive, the LSP will not switch back to the working path when the PSC switchover requests end, while for revertive configurations, the LSP always returns back to the working path when the switchover requests end.

The following figures illustrate the behavior of linear protection in more detail.

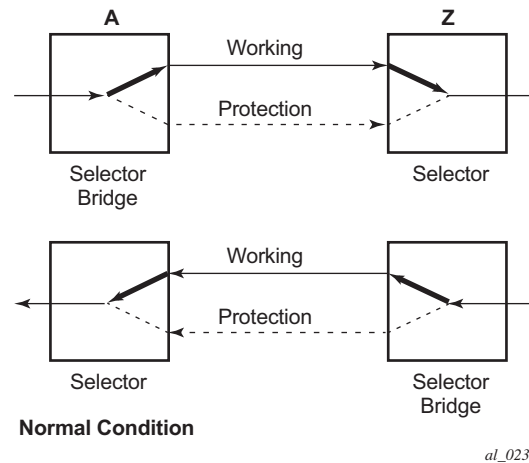


Figure 17: Normal Operation

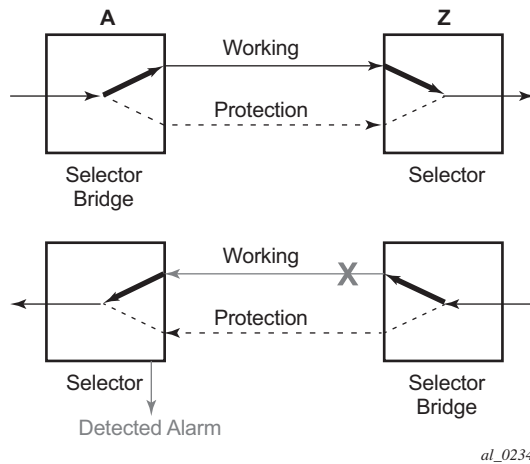


Figure 18: Failed Condition

In normal condition, user data packets are sent on the working path on both directions, from A to Z and Z to A.

A defect in the direction of transmission from node Z to node A impacts the working connection Z-to-A, and initiates the detection of a defect at the node A.

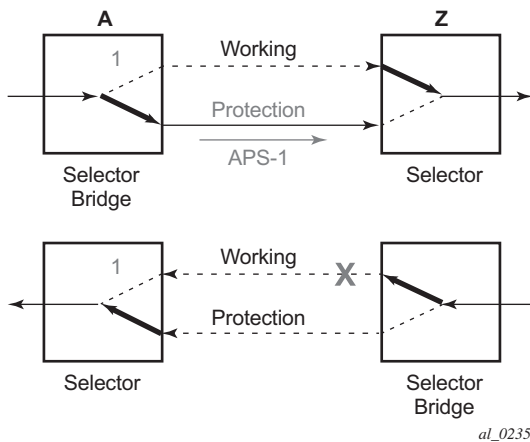


Figure 19: Failed Condition - Switching at A

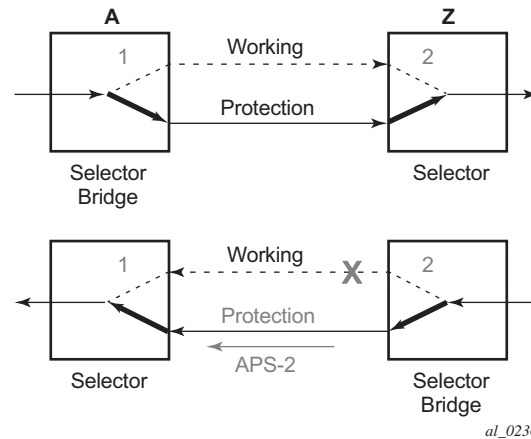


Figure 20: Failed Condition - Switching at Z

The unidirectional PSC protocol initiates protection switching: the selector bridge at node A is switched to protection connection A-to-Z and the selector at node A switches to protection connection Z to-A. The PSC packet, sent from node A to node Z, requests a protection switch to node Z.

After node Z validates the priority of the protection switch request, the selector at node Z is switched to protection connection A-to-Z and the selector bridge at the node Z is switched to protection connection Z-to-A. The PSC packet, sent from node Z to node A, is used as acknowledge, informing node A about the switching.

If BFD CC or CC/CV OAM packets are used to detect defects on the working and protection paths, they are inserted on both working and protection paths. It should be noted that they are sent regardless of whether the selected as the currently active path.

The 7750 supports the following operator commands:

- Forced Switch
- Manual Switch
- Clear,
- Lockout of protection

Switching Static MPLS-TP PWs to Dynamic, T-LDP Signaled PWs

Switching Static MPLS-TP to Dynamic T-LDP Signaled PWs

Some use cases for MPLS-TP require an MPLS-TP based aggregation network and an IP-based core network to interoperate, so providing the seamless transport of packet services across static MPLS-TP and dynamically signaled domains using an MS-PW. In this environment, end to end VCCV Ping and VCCV Trace may be used on the MS-PW. This is illustrated in the following figure:

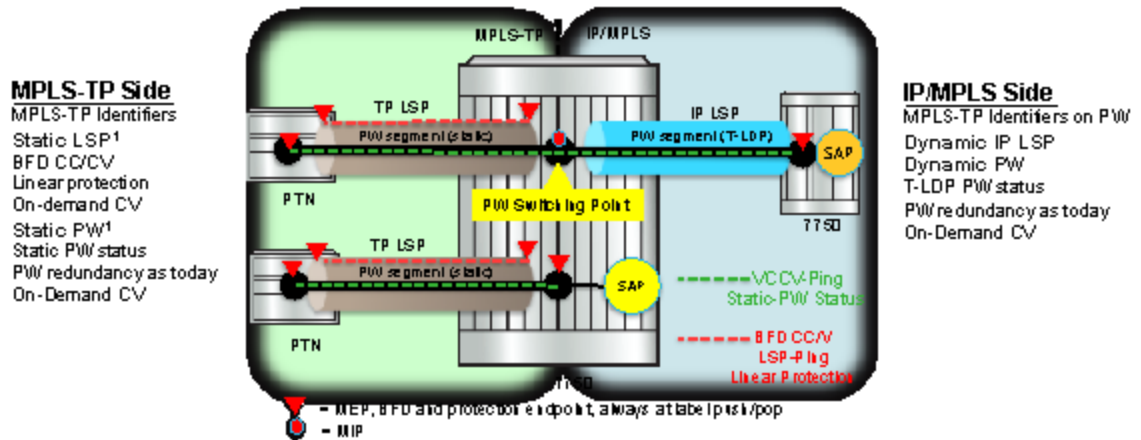


Figure 21: Static - Dynamic PW Switching with MPLS-TP

Services are backhauled from the static MPLS-TP network on the left to the dynamic IP/MPLS network on the right. The 7750 acts as an S-PE interconnecting the static and dynamic domains.

The 7x50 implementation supports such use cases through the ability to mate a static MPLS-TP spoke-sdp, with a defined pw-path-id, to a FEC128 spoke-sdp. The dynamically signaled spoke-sdp must be MPLS; GRE PWs are not supported, but the T-LDP signaled PW can use any supported MPLS tunnel type (e.g. LDP, RSVP-TE, static, BGP). The control-word must be enabled on both mate spoke-sdps.

Mapping of control channel status signaling to and from T-LDP status signaling at the 7x50 S-PE is also supported.

The use of VCCV Ping and VCCV Trace on an MS-PW composed of a mix of static MPLS-TP and dynamic FEC128 segments is described in more detail in the 7x50 SR OS OAM and Diagnostics Guide.

Configuring MPLS-TP

This section describes the steps required to configured MPLS-TP.

Configuration Overview

The following steps must be performed in order to configure MPLS-TP LSPs or PWs.

At the 7x50 LER and LSR:

1. Create an MPLS-TP context, containing nodal MPLS-TP identifiers. This is configured under **config>router>mpls>mpls-tp**.
2. Ensure that a sufficient range of labels is reserved for static LSPs and PWs. This is configured under **config>router>mpls-labels>static-labels**.
3. Ensure that a range of tunnel identifiers is reserved for MPLS-TP LSPs under **config>router>mpls-mpls-tp>tp-tunnel-id-range**.
4. A user may optionally configure MPLS-TP interfaces, which are interfaces that do not use IP addressing or ARP for next hop resolution. These can only be used by MPLS-TP LSPs.

At the 7x50 LER, configure:

1. OAM Templates. These contain generic parameters for MPLS-TP proactive OAM. An OAM template is configured under **config>router>mpls>mpls-tp>oam-template**.
2. BFD templates. These contain generic parameters for BFD used for MPLS-TP LSPs. A BFD template is configured under **config>router>bfd>bfd-template**.
3. Protection templates. These contain generic parameters for MPLS-TP 1-for-1 linear protection. A protection template is configured under **config>router>mpls>mpls-tp>protection-template**.
4. MPLS-TP LSPs are configured under **config>router>mpls>lsp mpls-tp**
5. Pseudowires using MPLS-TP are configured as spoke-sdps with static PW labels.

At an LSR, a user must configure an LSP transit-path under **config>router>mpls>mpls-tp>transit-path**.

The following sections describe these configuration steps in more detail.

Node-Wide MPLS-TP Parameter Configuration

Generic MPLS-TP parameters are configured under **config>router>mpls>mpls-tp**. If a user configures **no mpls**, normally the entire mpls configuration is deleted. However, in the case of mpls-tp a check that there is no other mpls-tp configuration e.g. services or tunnels using mpls-tp on the node, will be performed.

The mpls-tp context is configured as follows:

```
config
  router
    mpls
      [no] mpls-tp
      .
      .
      [no] shutdown
```

MPLS-TP LSPs may be configured if the mpls-tp context is administratively down (shutdown), but they will remain down until the mpls-tp context is configured as administratively up. No programming of the data path for an MPLS-TP Path occurs until the following are all true:

- Mpls-tp context is **no shutdown**
- Mpls-tp LSP context is **no shutdown**
- MPLS-TP Path context is **no shutdown**

A **shutdown** of mpls-tp will therefore bring down all MPLS-TP LSPs on the system.

The mpls-tp context cannot be deleted if MPLS-TP LSPs or SDPs exist on the system.

Node-Wide MPLS-TP Identifier Configuration

MPLS-TP identifiers are configured for a node under the following CLI tree:

```
config
  router
    mpls
      mpls-tp
        global-id <global-id>
        node-id {<ipv4address> | | <1.. .4,294,967,295>}
        [no] shutdown
      exit
```

The default value for the global-id is 0. This is used if the global-id is not explicitly configured. If a user expects that inter domain LSPs will be configured, then it is recommended that the global ID should be set to the local ASN of the node, as configured under **config>system**. If two-byte ASNs are used, then the most significant two bytes of the global-id are padded with zeros.

The default value of the node-id is the system interface IPv4 address. The MPLS-TP context cannot be administratively enabled unless at least a system interface IPv4 address is configured because MPLS requires that this value is configured.

These values are used unless overridden at the LSP or PW end-points, and apply only to static MPLS-TP LSPs and PWs.

In order to change the values, **config>router>mpls>mpls-tp** must be in the shutdown state. This will bring down all of the MPLS-TP LSPs on the node. New values are propagated to the system when a **no shutdown** is performed.

Static LSP and pseudowire (VC) Label and Tunnel Ranges

SR OS reserves a range of labels for use by static LSPs, and a range of labels for use by static pseudowires (SVCs) i.e. LSPs and pseudowires with no dynamic signaling of the label mapping. These are configured as follows:

```
config
  router
    mpls-labels
      [no] static-label max-lsp-labels <number>
      static-svc-label <number>
```

<number>: indicates the maximum number of labels for the label type.

The minimum label value for the static LSP label starts at 32 and expands all the way to the maximum number specified. The static VC label range is contiguous with this. The dynamic label range exists above the static VC label range (the label ranges for the respective label type are contiguous). This prevents fragmentation of the label range.

The MPLS-TP tunnel ID range is configured as follows:

```
config
  router
    mpls
      mpls-tp
        [no] tp-tunnel-id-range <start-id> <end-id>
```

The tunnel ID range referred to here is a contiguous range of RSVP-TE Tunnel IDs is reserved for use by MPLS TP, and these IDs map to the MPLS-TP Tunnel Numbers. There are some cases where the dynamic LSPs may have caused fragmentation to the number space such that contiguous range {max-min} is not available. In these cases, the command will fail.

There is no default value for the tunnel id range, and it must be configured to enable MPLS-TP.

If a configuration of the tunnel ID range fails, then the system will give a reason. This could be that the initially requested range, or the change to the allocated range, is not available i.e. tunnel

IDs in that range have already been allocated by RSVP-TE. Allocated Tunnel IDs are visible using a show command.

Note that changing the LSP or static VC label ranges does not require a reboot.

Note also that the static label ranges for LSPs, above, apply only to static LSPs configured using the CLI tree for MPLS-TP specified in this section. Different scalability constraints apply to static LSPs configured using the following CLI introduced in earlier SR OS releases:

```
config>router>mpls>static-lsp
```

```
config>router>mpls>interface>label-map
```

The scalability applying to labels configured using this CLI is enforced as follows:

- A maximum of 1000 static LSP names may be configured with a PUSH operation.
- A maximum of 1000 LSPs with a POP or SWAP operation may be configured.

These two limits are independent of one another, giving a combined limit of 1000 PUSH and 1000 POP/SAP operations configured on a node.

The static LSP and VC label spaces are contiguous. Therefore, the dimensioning of these label spaces requires careful planning by an operator as increasing the static LSP label space impacts the start of the static VC label space, which may already-deployed

Interface Configuration for MPLS-TP

It is possible for MPLS-TP paths to use both numbered IP numbered interfaces that use ARP/static ARP, or IP unnumbered interfaces. MPLS-TP requires no changes to these interfaces. It is also possible to use a new type of interface that does not require any IP addressing or next-hop resolution.

Draft-ietf-mpls-tp-next-hop-addressing provides guidelines for the usage of various Layer 2 next-hop resolution mechanisms with MPLS-TP. If protocols such as ARP are supported, then they should be used. However, in the case where no dynamic next hop resolution protocol is used, it should be possible to configure a unicast, multicast or broadcast next-hop MAC address. The rationale is to minimize the amount of configuration required for upstream nodes when downstream interfaces are changes. A default multicast MAC address for use by MPLS-TP point-to-point LSPs has been assigned by IANA (Value: 01-00-5e-90-00-00). This value is configurable on the 7x50 to support interoperability with 3rd party implementations that do not default to this value, and this no default value is implemented on the 7x50.

In order to support these requirements, a new interface type, known as an unnumbered MPLS-TP interface is introduced. This is an unnumbered interface that allows a broadcast or multicast

destination MAC address to be configured. An unnumbered MPLS-TP interface is configured using the **unnumbered-mpls-tp** keyword, as follows:

```
config
router
  interface <if-name> [unnumbered-mpls-tp]
    port <port-id>[:encap-val]
    mac <local-mac-address>
    static-arp <remote-mac-addr>
    //ieee-address needs to support mcast and bcast
    exit
```

The **remote-mac-address** may be any unicast, broadcast or multicast address. However, a broadcast or multicast remote-mac-address is only allowed in the **static-arp** command on Ethernet unnumbered interfaces when the **unnumbered-mpls-tp** keyword has been configured. This also allows the interface to accept packets on a broadcast or any multicast MAC address. Note that if a packet is received with a unicast destination MAC address, then it will be checked against the configured <local-mac-address> for the interface, and dropped if it does not match. When an interface is of type **unnumbered-mpls-tp**, only MPLS-TP LSPs are allowed on that interface; other protocols are blocked from using the interface.

An unnumbered MPLS-TP interface is assumed to be point-to-point, and therefore users must ensure that the associated link is not broadcast or multicast in nature if a multicast or broadcast remote MAC address is configured.

The following is a summary of the constraints of an unnumbered MPLS-TP interface:

- It is unnumbered and may borrow/use the system interface address
- It prevents explicit configuration of a borrowed address
- It prevents IP address configuration
- It prevents all protocols except mpls
- It prevents Deletion if an MPLS-TP LSP is bound to the Interface
- It is allowed only in network chassis mode D

MPLS-TP is only supported over Ethernet ports in Release 11.0. The system will block the association of an MPLS-TP LSP to an interface whose port is non-Ethernet.

If required, the IF_Num is configured under a MEP context under the MPLS interface. The **mpls-tp-mep** context is created under the interface as shown below. The *if-num* parameter, when concatenated with the Node ID, forms the IF_ID (as per RFC 6370), which is the identifier of this MEP. Note that it is possible to configure this context whether the interface is IP numbered, IP unnumbered or mpls-tp unnumbered:

```
config
router
  mpls
    interface <ip-int-name>
      mpls-tp-mep
```

```
[no] if-num <if-num>
[no] if-num-validation [enable|disable]
...
exit
```

The **if-num-validation** command is used to enable or disable validation of the if-num in LSP Trace packet against the locally configured if-num for the interface over which the LSP Trace packet was received at the egress LER. This is because some implementations, do not perform interface validation for unnumbered MPLS-TP interfaces and instead set the if-num in the dsmapi TLV to 0. The default is enabled.

LER Configuration for MPLS-TP

LSP and Path Configuration

MPLS-TP tunnels are configured using the **mpls-tp** LSP type at an LER under the LSP configuration, using the following CLI tree:

```

config
  router
    mpls
      lsp <xyz> [bypass-only|p2mp-lsp|mpls-tp <src-tunnel-num>]
        to node-id {<a.b.c.d> | <1.. .4,294,967,295>}
        dest-global-id <global-id>
        dest-tunnel-number <tunnel-num>
        [no] working-tp-path
          lsp-num <lsp-num>
          in-label <in-label>
          out-label <out-label> out-link <if-name>
            [next-hop <ipv4-address>]
        [no] mep
          [no] oam-template <name>
          [no] bfd-enable [cc | cc_cv] // defaults to cc
          [no] shutdown
          exit
        [no] shutdown
        exit
      [no] protect-tp-path
        lsp-num <lsp-num>
        in-label <in-label>
        out-label <out-label> out-link <if-name>
          [next-hop <ipv4-address> ]
        [no] mep
          [no] protection-template <name>
          [no] oam-template <name>
          [no] bfd-enable [cc | cc_cv] //defaults to cc
          [no] shutdown
          exit
        [no] shutdown
        exit

```

<if-name> could be numbered or unnumbered interface using an Ethernet port.

<src-tunnel-num> is a mandatory create time parameter for mpls-tp tunnels, and has to be assigned by the user based on the configured range of tunnel ids. The *src-global-id* used for the LSP ID is derived from the node-wide *global-id* value configured under `config>router>mpls>mpls-tp`. A tunnel can not be **un shutdown** unless the *global-id* is configured.

The from address of an LSP to be used in the tunnel identifier is taken to be the local node's node-id/global-id, as configured under `config>router>mpls>mpls-tp`. If that is not explicitly configured, either, then the default value of the system interface IPv4 address is used

The **to node-id** address may be entered in 4-octet IPv4 address format or unsigned 32-bit format. This is the far-end node-id for the LSP, and does not need to be routable IP addresses.

The **from** and **to** addresses are used as the from and to node-id in the MPLS-TP Tunnel Identifier used for the MEP ID.

Each LSP consists of a working-tp-path and, optionally, a protect-tp-path. The protect-tp-path provides protection for the working-tp-path if 1:1 linear protection is configured (see below). Proactive OAM, such as BFD, is configured under the MEP context of each path. Protection for the LSP is configured under the protect-tp-path MEP context.

The *to* global-id is an optional parameter. If it is not entered, then the destination global ID takes the default value of 0. Global ID values of 0 are allowed and indicate that the node's configured Global ID should be used. If the local global ID value is 0, then the remote **to** global ID must also be 0. The *to* global ID value cannot be changed if an LSP is in use by an SDP.

The *to* tunnel number is an optional parameter. If it is not entered, then it is taken to be the same value as the source tunnel number.

LSPs are assumed to be bidirectional and co-routed. Therefore, the system will assume that the incoming interface is the same as the out-link.

The next-hop *ip-address* can only be configured if the out-link if-name refers to a numbered IP interface. In this case, the system will determine the interface to use to reach the configured next-hop, but will check that the user-entered value for the out-link corresponds to the link returned by the system. If they do not correspond, then the path will not come up. Note that if a user changes the physical port referred to in the interface configuration, then BFD, if configured on the LSP, will go down. Users should therefore ensure that an LSP is moved to a different interface with a different port configuration in order to change the port that it uses. This is enforced by blocking the next-hop configuration for an unnumbered interface.

There is no check made that a valid ARP entry exists before allowing a path to be unshut. Therefore, a path will only be held down if BFD is down. If static ARP is not configured for the interface, then it is assumed that dynamic ARP is used. The result is that if BFD is not configured, a path can come up before ARP resolution has completed for an interface. If BFD is not used, then it is recommended that the connectivity of the path is explicitly checked using on-demand CC/CV prior to sending user traffic on it.

The following is a list of additional considerations for the configuration of MPLS-TP LSPs and paths:

- The working-tp-path must be configured before the protect-tp-path.
- Likewise, the protect-tp-path has to be deleted first before the working-tp-path.
- The *lsp-num* parameter is optional. The default values are 1 for the working-tp-path and 2 for protect-tp-path.
- The **mep** context must be deleted before a path can be deleted.

- An MPLS interface needs to be created under **config>router>mpls>interface** before using/specifying the out-label/out-link in the Forward path for an MPLS-TP LSP. Creation of the LSP will fail if the corresponding mpls interface doesn't exist even though the specified router interface may be valid.
- The system will program the MPLS-TP LSP information upon a **no shutdown** of the TP-Path only on the very first **no shutdown**. The Working TP-Path is programmed as the Primary and the Protection TP-Path is programmed as the backup.
- The system will not de-program the IOM on an admin shutdown of the MPLS-TP path. Traffic will gracefully move to the other TP-Path if valid, as determined by the proactive MPLS-TP OAM. This should not result in traffic loss. However it is recommended that the user does moves traffic to the other TP-Path through a tools command before doing 'admin shut' of an Active TP-Path.
- Deletion of the out-label/out-link sub-command under the MPLS-TP Path is not allowed once configured. These can only be modified.
- MPLS will allow the deletion of an 'admin shutdown' TP-Path. This will cause MPLS to de-program the corresponding TP-Path forwarding information from IOM. This can cause traffic loss for certain users that are bound to the MPLS-TP LSP.
- MPLS will not de-program the IOM on a specific interface admin shut/clear unless the interface is a System Interface. However, if mpls informs the TP-OAM module that the mpls interface has gone down, then it triggers a switch to the standby tp-path if the associated interface went down and if it is valid.
- If a MEP is defined and shutdown, then the corresponding path is also operationally down. Note, however, that the MEP admin state is applicable only when a MEP is created from an MPLS-TP path.
- It is not mandatory to configure BFD or protection on an MPLS-TP path in order to bring the LSP up.
- If **bfd-enable cc** is configured, then CC-only mode using ACh channel 0x07 is used. If **bfd-enable cc_v** is configured, then BFD CC packets use channel 0x22 and CV packets use channel 0x23.

The protection template is associated with a LSP as a part of the MEP on the protect path. If only a working path is configured, then the protection template is not configured.

BFD cannot be enabled under the MEP context unless a named BFD template is configured.

Support for Downstream Mapping Information

In order to validate the downstream mapping for an LSP, a node sending a DSMAP TLV must include the incoming and (optionally) outgoing IF_Num values for the interfaces that it expects the LSP to transit. Additionally, it will include the out-label for the LSP in the Label TLV for the DSMAP in the echo request message.

The incoming and outgoing if-num values correspond to the incoming and outgoing interfaces transited by an LSP at the next hop LER and LSR are configured using the **dsmap** command, as follows:

```
config
router
mpls
lsp
working-tp-path
mep
dsmap <in-if-num>[:<out-if-num>]

config
router
mpls
lsp
protect-tp-path
mep
dsmap <in-if-num>[:<out-if-num>]

config
router
mpls
mpls-tp
transit-path
forward-path
mip
dsmap <in-if-num>[:<out-if-num>]
exit
reverse-path
mip
dsmap <in-if-num>[:<out-if-num>]
exit
```

A node sending a DSMAP TLV will include these **in-if-num** and **out-if-num** (if configured) values. Additionally, it will include the out-label for the LSP in the Label TLV for the DSMAP in the echo request message.

Proactive CC/CV (using BFD) Configuration

Generally applicable proactive OAM parameters are configured using templates.

Proactive CC and CV uses BFD parameters such as Tx/Rx timer intervals, multiplier and other session/fault management parameters which are specific to BFD. These are configured using a BFD Template. The BFD Template may be used for non-MPLS-TP applications of BFD, and therefore contains the full set of possible configuration parameters for BFD. Only a sub-set of these may be used for any given application.

Generic MPLS-TP OAM and fault management parameters are configured in the OAM Template.

Named templates are referenced from the MPLS-TP Path MEP configuration, so different parameter values are possible for the working and protect paths of a tunnel.

The BFD Template is configured as follows:

```
config
  router
    bfd
      [no] bfd-template <name>
      [no] transmit-interval <transmit-interval>
      [no] receive-interval <receive-interval>
      [no] echo-receive <echo-interval>
      [no] multiplier <multiplier>
      [no] type <cpm-np>
    exit
```

The parameters are as follows:

- **transmit-interval** *transmit-interval* and the **rx** *receive-interval*: These are the transmit and receive timers for BFD packets. If the template is used for MPLS-TP, then these are the timers used by CC packets. Values are in milliseconds: 10ms to 100,000ms, with 1ms granularity. Default 10ms for CPM3 or better, 1 sec for other hardware. Note that for MPLS-TP CV packets, a transmit interval of 1 sec is always used.
- **multiplier** *multiplier*: Integer 3 – 20. Default: 3. This parameter is ignored for MPLS-TP combined cc-v BFD sessions, and the default of 3 used, as per RFC6428.
- **echo-receive** *echo-interval*: Sets the minimum echo receive interval, in milliseconds, for a session. Values: 100ms – 100,000ms. Default: 100. This parameter is not used by a BFD session for MPLS-TP.
- **type** **cpm-np**: This selects the CPM network processor as the local termination point for the BFD session. This is enabled by default.

Note that if the above BFD timer values are changed in a given template, any BFD sessions on MEPs to which that template is bound will try to renegotiate their timers to the new values. Note that the BFD implementations in some MPLS-TP peer nodes may not be able handle this renegotiation, as allowed by Section 3.7.1 of RFC6428 and may take the BFD session down. This

could result in undesired behavior, for example an unexpected protection switching event. It is therefore recommended that in these circumstances, user of the system exercise care in modifying the BFD timer values after a BFD session is UP.

Commands within the BFD-template use a begin-commit model. To edit any value within the BFD template, a *begin* needs to be executed once the template context has been entered. However, a value will still be stored temporarily until the commit is issued. Once the commit is issued, values will actually be used by other modules like the mpls-tp module and BFD module.

A BFD template is referenced from the OAM template. The OAM Template is configured as follows:

```
config
router
 mpls
  mpls-tp
    [no] oam-template <name>
    [no] bfd-template <name>
    [no] hold-time-down <interval>
    [no] hold-time-up <interval>
  exit
```

- **hold-time-down interval:** 0-5000 deciseconds, 10ms steps, default 0. This is equivalent to the standardized hold-off timer.
- **hold-time-up interval:** 0-500 centiseconds in 100ms steps, default 2 seconds This is an additional timer that can be used to reduce BFD bouncing.
- **bfd-template name:** This is the named BFD template to use for any BFD sessions enabled under a MEP for which the OAM template is configured.

An OAM template is then applied to a MEP as described above.

Protection templates and Linear Protection Configuration

Protection templates defines the generally applicable protection parameters for an MPLS-TP tunnel. Only linear protection is supported, and so the application of a named template to an MPLS-TP tunnel implies that linear protection is used.

A template is configured as follows:

```
config
router
 mpls
  mpls-tp
    protection-template <name>
    [no] revertive
    [no] wait-to-restore <interval>
    rapid-psc-timer <interval>
    slow-psc-timer <interval>
  exit
```


The allowed values are as follows:

- **wait-to-restore** *interval*: 0-720 seconds, 1 sec steps, default 300 seconds. This is applicable to revertive mode only.
- **rapid-psc-timer** *interval*: [10, 100, 1000ms]. Default 100ms
- **slow-psc-timer** *interval*: 5s-60s. Default: 5s
- **revertive**: Selects revertive behavior. Default: no revertive.

LSP Linear Protection operations are enacted using the following **tools>perform** commands.

```
tools>perform>router>mpls
  tp-tunnel
    clear {<lsp-name> | id <tunnel-id>}
    force {<lsp-name> | id <tunnel-id>}
    lockout {<lsp-name> | id <tunnel-id>}
    manual {<lsp-name> | id <tunnel-id>}
  exit
exit
```

To minimize outage times, users should use the “mpls-tp protection command” (e.g. force/manual) to switch all the relevant MPLS-TP paths before executing the following commands:

- clear router mpls interface <>
- config router mpls interface <> shut

Intermediate LSR Configuration for MPLS-TP LSPs

The forward and reverse directions of the MPLS-TP LSP Path at a transit LSR are configured using the following CLI tree:

```
config
  router
    mpls
      mpls-tp
        transit-path <path-name>
          [no] path-id {lsp-num <lsp-num>|working-path|protect-path
            [src-global-id <global-id>]
            src-node-id {<ipv4address> | <1.. .4,294,967,295>}
            src-tunnel-num <tunnel-num>
            [dest-global-id <global-id>]
            dest-node-id {<ipv4address> | <1.. .4,294,967,295>}
            [dest-tunnel-num <tunnel-num>]}
          forward-path
            in-label <in-label> out-label <out-label>
            out-link <if-name> [next-hop <ipv4-next-hop>]
          reverse-path
            in-label <in-label> out-label <out-label>
            [out-link <if-name> [next-hop <ipv4-next-hop>]
          [no] shutdown
```

Note that the *src-tunnel-num* and *dest-tunnel-num* are consistent with the source and destination of a label mapping message for a signaled LSP.

If *dest-tunnel-num* is not entered in CLI, the *dest-tunnel-num* value is taken to be the same as the SRC-tunnel-num value.

If any of the *global-id* values are not entered, the value is taken to be 0.

If the *src-global-id* value is entered, but the *dest-global-id* value is not entered, *dest-global-id* value is the same as the *src-global-id* value.

Note that the *lsp-num* must match the value configured in the LER for a given path. If no explicit *lsp-num* is configured, then *working-path* or *protect-path* must be specified (equating to 1 or 2 in the system).

The forward path must be configured before the reverse path. The configuration of the reverse path is optional.

The LSP-ID (*path-id*) parameters apply with respect to the downstream direction of the forward LSP path, and are used to populate the MIP ID for the path at this LSR.

The reverse path configuration must be deleted before the forward path.

The *forward-path* (and *reverse-path* if applicable) parameters can be configured with or without the *path-id*, but they must be configured if MPLS-TP OAM is to be able to identify the LSR MIP.

The *transit-path* can be no shutdown (as long as the *forward-path/reverse-path* parameters have been configured properly) with or without identifiers.

The *path-id* and *path-name* must be unique on the node. There is a one to one mapping between a given *path-name* and *path-id*.

Traffic can not pass through the *transit-path* if the *transit-path* is in the **shutdown** state.

MPLS-TP Show Commands

Static MPLS Labels

The following new commands show the details of the static MPLS labels.

```
show>router>mpls-labels>label <start-label> [<end-label> [in-use|<label-owner>]]
```

```
show>router>mpls-labels>label-range
```

An example output is as follows:

```
*A:mlstp-dutA# show router mpls
mpls          mpls-labels
*A:mlstp-dutA# show router mpls label
label        label-range
*A:mlstp-dutA# show router mpls label-range

=====
Label Ranges
=====
Label Type      Start Label      End Label      Aging      Total Available
-----
Static-lsp      32                16415          -           16364
Static-svc      16416             32799          -           16376
Dynamic         32800             131071         0           98268
=====
```

MPLS-TP Tunnel Configuration

These should show the configuration of a given tunnel.

```
show>router>mpls>tp-lsp
```

A sample output is as follows:

```
*A:mlstp-dutA# show router mpls tp-lsp
- tp-lsp [<lsp-name>] [status {up|down}] [from <ip-address>|to <ip-address>]
  [detail]
- tp-lsp [<lsp-name>] path [protect|working] [detail]
- tp-lsp [<lsp-name>] protection

<lsp-name>      : [32 chars max] - accepts * as wildcard char
<path>         : keyword - Display LSP path information.
<protection>   : keyword - Display LSP protection information.
<up|down>      : keywords - Specify state of the LSP
<ip-address>   : a.b.c.d
<detail>       : keyword - Display detailed information.
*A:mlstp-dutA# show router mpls tp-lsp
path
protection
```

```

to <a.b.c.d>
<lsp-name>
  "lsp-32" "lsp-33" "lsp-34" "lsp-35" "lsp-36" "lsp-37" "lsp-38" "lsp-39"
  "lsp-40" "lsp-41"
status {up|down}
from <ip-address>
detail

*A:mlstp-dutA# show router mpls tp-lsp "lsp-
"lsp-32" "lsp-33" "lsp-34" "lsp-35" "lsp-36" "lsp-37" "lsp-38" "lsp-39"
"lsp-40" "lsp-41"
*A:mlstp-dutA# show router mpls tp-lsp "lsp-32"

=====
MPLS MPLS-TP LSPs (Originating)
=====
LSP Name                To                Tun      Protect  Adm  Opr
                        Id                Id       Path
-----
lsp-32                   0.0.3.234        32       No       Up   Up
-----
LSPs : 1
=====
*A:mlstp-dutA# show router mpls tp-lsp "lsp-32" detail

=====
MPLS MPLS-TP LSPs (Originating) (Detail)
=====
Type : Originating
-----
LSP Name      : lsp-32
LSP Type      : MplsTp                LSP Tunnel ID : 32
From Node Id: 0.0.3.233+             To Node Id    : 0.0.3.234
Adm State     : Up                    Oper State     : Up
LSP Up Time  : 0d 04:50:47            LSP Down Time : 0d 00:00:00
Transitions  : 1                      Path Changes   : 2

DestGlobalId: 42                      DestTunnelNum  : 32

```

MPLS-TP Path configuration.

This can reuse and augment the output of the current show commands for static LSPs. They should also show if BFD is enabled on a given path. If this referring to a transit path, this should also display (among others) the path-id (7 parameters) for a given transit-path-name, or the transit-path-name for a given the path-id (7 parameters)

show>router>mpls>tp-lsp>path

A sample output is as follows:

```

=====
*A:mlstp-dutA# show router mpls tp-lsp path

```

```

=====
MPLS-TP LSP Path Information
=====
LSP Name      : lsp-32                               To      : 0.0.3.234
Admin State   : Up                                  Oper State : Up

-----
Path          NextHop          InLabel  OutLabel  Out I/F      Admin  Oper
-----
Working                               32       32       AtoB_1       Up     Down
Protect                               2080     2080     AtoC_1       Up     Up
-----
LSP Name      : lsp-33                               To      : 0.0.3.234
Admin State   : Up                                  Oper State : Up

-----
Path          NextHop          InLabel  OutLabel  Out I/F      Admin  Oper
-----
Working                               33       33       AtoB_1       Up     Down
Protect                               2082     2082     AtoC_1       Up     Up
-----
LSP Name      : lsp-34                               To      : 0.0.3.234
Admin State   : Up                                  Oper State : Up

-----
Path          NextHop          InLabel  OutLabel  Out I/F      Admin  Oper
-----
Working                               34       34       AtoB_1       Up     Down
Protect                               2084     2084     AtoC_1       Up     Up
-----
LSP Name      : lsp-35                               To      : 0.0.3.234
Admin State   : Up                                  Oper State : Up

-----
Path          NextHop          InLabel  OutLabel  Out I/F      Admin  Oper
-----
Working                               35       35       AtoB_1       Up     Down
Protect                               2086     2086     AtoC_1       Up     Up
-----
LSP Name      : lsp-36                               To      : 0.0.3.234
Admin State   : Up                                  Oper State : Up

-----
Path          NextHop          InLabel  OutLabel  Out I/F      Admin  Oper
-----
Working                               36       36       AtoB_1       Up     Down
Protect                               2088     2088     AtoC_1       Up     Up
-----
LSP Name      : lsp-37                               To      : 0.0.3.234
Admin State   : Up                                  Oper State : Up

-----
Path          NextHop          InLabel  OutLabel  Out I/F      Admin  Oper
-----
Working                               37       37       AtoB_1       Up     Down
Protect                               2090     2090     AtoC_1       Up     Up
-----
LSP Name      : lsp-38                               To      : 0.0.3.234
Admin State   : Up                                  Oper State : Up

```

```

-----
Path          NextHop          InLabel  OutLabel  Out I/F          Admin  Oper
-----
Working              38          38          AtoB_1          Up     Down
Protect             2092        2092        AtoC_1          Up     Up
=====
LSP Name       : lsp-39                      To           : 0.0.3.234
Admin State   : Up                      Oper State    : Up

```

```

-----
Path          NextHop          InLabel  OutLabel  Out I/F          Admin  Oper
-----
Working              39          39          AtoB_1          Up     Down
Protect             2094        2094        AtoC_1          Up     Up
=====
LSP Name       : lsp-40                      To           : 0.0.3.234
Admin State   : Up                      Oper State    : Up

```

```

-----
Path          NextHop          InLabel  OutLabel  Out I/F          Admin  Oper
-----
Working              40          40          AtoB_1          Up     Down
Protect             2096        2096        AtoC_1          Up     Up
=====
LSP Name       : lsp-41                      To           : 0.0.3.234
Admin State   : Up                      Oper State    : Up

```

```

-----
Path          NextHop          InLabel  OutLabel  Out I/F          Admin  Oper
-----
Working              41          41          AtoB_1          Up     Down
Protect             2098        2098        AtoC_1          Up     Up

```

```
*A:mlstp-dutA# show router mpls tp-lsp "lsp-32" path working
```

```

=====
MPLS-TP LSP Working Path Information
LSP: "lsp-32"

```

```

=====
LSP Name       : lsp-32                      To           : 0.0.3.234
Admin State   : Up                      Oper State    : Up

```

```

-----
Path          NextHop          InLabel  OutLabel  Out I/F          Admin  Oper
-----
Working              32          32          AtoB_1          Up     Down

```

```
*A:mlstp-dutA# show router mpls tp-lsp "lsp-32" path protect
```

```

=====
MPLS-TP LSP Protect Path Information
LSP: "lsp-32"

```

```

=====
LSP Name       : lsp-32                      To           : 0.0.3.234
Admin State   : Up                      Oper State    : Up

```

```

-----
Path          NextHop          InLabel  OutLabel  Out I/F          Admin  Oper

```

```

-----
Protect                2080      2080      AtoC_1      Up      Up
-----
*A:mlstp-dutA# show router mpls tp-lsp "lsp-32" path protect detail

=====
MPLS-TP LSP Protect Path Information
  LSP: "lsp-32" (Detail)
=====
LSP Name       : lsp-32                               To           : 0.0.3.234
Admin State    : Up                                   Oper State    : Up

Protect path information
-----
Path Type      : Protect                               LSP Num      : 2
Path Admin     : Up                                   Path Oper     : Up
Out Interface  : AtoC_1                               Next Hop Addr : n/a
In Label       : 2080                                 Out Label     : 2080
Path Up Time   : 0d 04:52:17                          Path Dn Time  : 0d 00:00:00
Active Path    : Yes                                   Active Time   : 0d 00:52:56

MEP information
MEP State      : Up                                   BFD           : cc
OAM Templ     : privatebed-oam-template              CC Status     : inService
                                                       CV Status     : unknown
Protect Templ  : privatebed-protection-template      WTR Count Down: 0 seconds
RX PDU        : SF (1,1)                             TX PDU        : SF (1,1)
Defects       :
=====
*A:mlstp-dutA# show router mpls tp-lsp "lsp-32" path working detail

=====
MPLS-TP LSP Working Path Information
  LSP: "lsp-32" (Detail)
=====
LSP Name       : lsp-32                               To           : 0.0.3.234
Admin State    : Up                                   Oper State    : Up

Working path information
-----
Path Type      : Working                               LSP Num      : 1
Path Admin     : Up                                   Path Oper     : Down
Down Reason    : ccFault ifDn
Out Interface  : AtoB_1                               Next Hop Addr : n/a
In Label       : 32                                   Out Label     : 32
Path Up Time   : 0d 00:00:00                          Path Dn Time  : 0d 00:53:01
Active Path    : No                                   Active Time   : n/a

MEP information
MEP State      : Up                                   BFD           : cc
OAM Templ     : privatebed-oam-template              CC Status     : outOfService
                                                       CV Status     : unknown
=====
*A:mlstp-dutA#

```

MPLS-TP Protection

These should show the protection configuration for a given tunnel, which path in a tunnel is currently working and which is protect, and whether the working or protect is currently active.

show>router>mpls>tp-lsp>protection

A sample output is as follows:

```
*A:mlstp-dutA# show router mpls tp-lsp protection

=====
MPLS-TP LSP Protection Information
Legend: W-Working, P-Protect,
=====
LSP Name                Admin Oper  Path      Ingr/Egr    Act. Rx PDU
                        State State   State Label    Path Tx PDU
-----
lsp-32                   Up    Up    W Down      32/32      No  SF (1,1)
                        P Up      2080/2080  Yes  SF (1,1)
lsp-33                   Up    Up    W Down      33/33      No  SF (1,1)
                        P Up      2082/2082  Yes  SF (1,1)
lsp-34                   Up    Up    W Down      34/34      No  SF (1,1)
                        P Up      2084/2084  Yes  SF (1,1)
lsp-35                   Up    Up    W Down      35/35      No  SF (1,1)
                        P Up      2086/2086  Yes  SF (1,1)
lsp-36                   Up    Up    W Down      36/36      No  SF (1,1)
                        P Up      2088/2088  Yes  SF (1,1)
lsp-37                   Up    Up    W Down      37/37      No  SF (1,1)
                        P Up      2090/2090  Yes  SF (1,1)
lsp-38                   Up    Up    W Down      38/38      No  SF (1,1)
                        P Up      2092/2092  Yes  SF (1,1)
lsp-39                   Up    Up    W Down      39/39      No  SF (1,1)
                        P Up      2094/2094  Yes  SF (1,1)
lsp-40                   Up    Up    W Down      40/40      No  SF (1,1)
                        P Up      2096/2096  Yes  SF (1,1)
lsp-41                   Up    Up    W Down      41/41      No  SF (1,1)
                        P Up      2098/2098  Yes  SF (1,1)
-----
No. of MPLS-TP LSPs: 10
=====
```

MPLS TP Node Configuration

Displays the Global ID, Node ID and other general MPLS-TP configurations for the node.

show>router>mpls>mpls-tp

A sample output is as follows:

```
*A:mlstp-dutA# show router mpls mpls-tp
- mpls-tp
```



```

oam-template      - Display MPLS-TP OAM Template information
protection-tem*   - Display MPLS-TP Protection Template information
status           - Display MPLS-TP system configuration
transit-path     - Display MPLS-TP Tunnel information

*A:mlstp-dutA# show router mpls mpls-tp oam-template

=====
MPLS-TP OAM Templates
=====
Template Name : privatebed-oam-template Router ID      : 1
BFD Template  : privatebed-bfd-template Hold-Down Time: 0 centiseconds
                                           Hold-Up Time  : 20 deciseconds
=====
*A:mlstp-dutA# show router mpls mpls-tp protection-template

=====
MPLS-TP Protection Templates
=====
Template Name : privatebed-protection-template Router ID      : 1
Protection Mode: one2one Direction                : bidirectional
Revertive     : revertive Wait-to-Restore: 300sec
Rapid-PSC-Timer: 10ms Slow-PSC-Timer : 5sec
=====
*A:mlstp-dutA# show router mpls mpls-tp status

=====
MPLS-TP Status
=====
Admin Status : Up
Global ID    : 42 Node ID              : 0.0.3.233
Tunnel Id Min : 1 Tunnel Id Max       : 4096
=====
*A:mlstp-dutA# show router mpls mpls-tp transit-path
  - transit-path [<path-name>] [detail]

<path-name>      : [32 chars max]
<detail>         : keyword - Display detailed information.

A:mlstp-dutC# show router mpls mpls-tp transit-path
  - transit-path [<path-name>] [detail]

<path-name>      : [32 chars max]
<detail>         : keyword - Display detailed information.

A:mlstp-dutC# show router mpls mpls-tp transit-path
<path-name>
"tp-32" "tp-33" "tp-34" "tp-35" "tp-36" "tp-37" "tp-38" "tp-39"
"tp-40" "tp-41"
detail

A:mlstp-dutC# show router mpls mpls-tp transit-path "tp-32"

=====

```

```

MPLS-TP Transit tp-32 Path Information
=====
Path Name      : tp-32
Admin State    : Up                               Oper State    : Up

-----
Path           NextHop           InLabel   OutLabel   Out I/F
-----
FP              2080           2081      CtoB_1
RP              2081           2080      CtoA_1
=====
A:mplstp-dutC# show router mpls mpls-tp transit-path "tp-32" detail

=====
MPLS-TP Transit tp-32 Path Information (Detail)
=====
Path Name      : tp-32
Admin State    : Up                               Oper State    : Up

-----
Path ID configuration
Src Global ID  : 42                               Dst Global ID : 42
Src Node ID   : 0.0.3.234                       Dst Node ID   : 0.0.3.233
LSP Number    : 2                               Dst Tunnel Num: 32

Forward Path configuration
In Label      : 2080                               Out Label     : 2081
Out Interface : CtoB_1                             Next Hop Addr : n/a

Reverse Path configuration
In Label      : 2081                               Out Label     : 2080
Out Interface : CtoA_1                             Next Hop Addr : n/a
=====
A:mplstp-dutC#

```

MPLS-TP Interfaces

The existing show>router>interface command should be enhanced to display mpls-tp specific information.

The following is a sample output:

```

*A:mlstp-dutA# show router interface "AtoB_1"

=====
Interface Table (Router: Base)
=====
Interface-Name      Adm      Opr (v4/v6)  Mode      Port/SapId
IP-Address          PfxState
-----
AtoB_1              Down     Down/--      Network  1/2/3:1
Unnumbered If[system]
n/a
-----
Interfaces : 1

```

MPLS-TP Debug Commands

The following command provides the debug command for an MPLS-TP tunnel:

```
tools>dump>router>mpls>tp-tunnel <lsp-name> [clear]
```

The following is a sample output:

```
A:mlstp-dutA# tools dump router mpls tp-tunnel
- tp-tunnel <lsp-name> [clear]
- tp-tunnel id <tunnel-id> [clear]

<lsp-name>          : [32 chars max]
<tunnel-id>         : [1..61440]
<clear>             : keyword - clear stats after reading

*A:mlstp-dutA# tools dump router mpls tp-tunnel "lsp-
"lsp-32" "lsp-33" "lsp-34" "lsp-35" "lsp-36" "lsp-37" "lsp-38" "lsp-39"
"lsp-40" "lsp-41"
*A:mlstp-dutA# tools dump router mpls tp-tunnel "lsp-32"

Idx: 1-32 (Up/Up): pgId 4, paths 2, operChg 1, Active: Protect
TunnelId: 42::0.0.3.233::32-42::0.0.3.234::32
PgState: Dn, Cnt/Tm: Dn 1/000 04:00:48.160 Up:3/000 00:01:25.840
MplsMsg: tpDn 0/000 00:00:00.000, tunDn 0/000 00:00:00.000
         wpDn 0/000 00:00:00.000, ppDn 0/000 00:00:00.000
         wpDel 0/000 00:00:00.000, ppDel 0/000 00:00:00.000
         tunUp 1/000 00:00:02.070

Paths:
Work (Up/Dn): Lsp 1, Lbl 32/32, If 2/128 (1/2/3 : 0.0.0.0)
  Tmpl: ptc: , oam: privatebed-oam-template (bfd: privatebed-bfd-template(np)-10 ms)
  Bfd: Mode CC state Dn/Up handle 160005/0
  Bfd-CC (Cnt/Tm): Dn 1/000 04:00:48.160 Up:1/000 00:01:23.970
  State: Admin Up (1::1::1) port Up , if Dn , operChg 2
  DnReasons: ccFault ifDn

Protect (Up/Up): Lsp 2, Lbl 2080/2080, If 3/127 (5/1/1 : 0.0.0.0)
  Tmpl: ptc: privatebed-protection-template, oam: privatebed-oam-template (bfd: pri-
vatebed-bfd-template(np)-10 ms)
  Bfd: Mode CC state Up/Up handle 160006/0
  Bfd-CC (Cnt/Tm): Dn 0/000 00:00:00.000 Up:1/000 00:01:25.410
  State: Admin Up (1::1::1) port Up , if Up , operChg 1

Aps: Rx - 5, raw 3616, nok 0(), txRaw - 3636, revert Y
Pdu: Rx - 0x1a-21::0101 (SF), Tx - 0x1a-21::0101 (SF)
State: PF:W:L LastEvt pdu (L-SFw/R-SFw)
Tmrs: slow
Defects: None Now: 000 05:02:19.130
Seq  Event  state      TxPdu      RxPdu      Dir      Act      Time
===  =====  =====  =====  =====  =====  =====  =====
000  start    UA:P:L     SF (0,0)   NR (0,0)   Tx-->   Work    000 00:00:02.080
001  pdu      UA:P:L     SF (0,0)   SF (0,0)   Rx<--   Work    000 00:01:24.860
002  pdu      UA:P:L     SF (0,0)   NR (0,0)   Rx<--   Work    000 00:01:26.860
003  pUp      NR         NR (0,0)   NR (0,0)   Tx-->   Work    000 00:01:27.440
004  pdu      NR         NR (0,0)   NR (0,0)   Rx<--   Work    000 00:01:28.760
005  wDn     PF:W:L     SF (1,1)   NR (0,0)   Tx-->   Prot    000 04:00:48.160
```

```

006 pdu PF:W:L SF (1,1) NR (0,1) Rx<-- Prot 000 04:00:48.160
007 pdu PF:W:L SF (1,1) SF (1,1) Rx<-- Prot 000 04:00:51.080

```

The following command shows the free mpls tunnel IDs available between two values, start-range and end-range.

```
tools>dump>router>mpls>free-tunnel-id <start-range> <end-range>
```

The following command provides a debug tool to view control-channel-status signaling packets.

```
*A:bksim1611# /debug service id 700 sdp 200:700 event-type ?{config-change|oper-status-change|neighbor-discovery|control-channel-status}
```

```
*A:bksim1611# /debug service id 700 sdp 200:700 event-type control-channel-status
```

```
*A:bksim1611#
```

```
1 2012/08/31 09:56:12.09 EST MINOR: DEBUG #2001 Base PW STATUS SIG PKT (RX):
```

```
"PW STATUS SIG PKT (RX)::
```

```
Sdp Bind 200:700 Instance 3
  Version          : 0x0
  PW OAM Msg Type  : 0x27
  Refresh Time     : 0xa
  Total TLV Length : 0x8
  Flags            : 0x0
  TLV Type         : 0x96a
  TLV Len          : 0x4
  PW Status Bits   : 0x0
"
```

```
2 2012/08/31 09:56:22.09 EST MINOR: DEBUG #2001 Base PW STATUS SIG PKT (RX):
```

```
"PW STATUS SIG PKT (RX)::
```

```
Sdp Bind 200:700 Instance 3
  Version          : 0x0
  PW OAM Msg Type  : 0x27
  Refresh Time     : 0xa
  Total TLV Length : 0x8
  Flags            : 0x0
  TLV Type         : 0x96a
  TLV Len          : 0x4
  PW Status Bits   : 0x0
"
```

```
3 2012/08/31 09:56:29.44 EST MINOR: DEBUG #2001 Base PW STATUS SIG PKT (TX):
```

```
"PW STATUS SIG PKT (TX)::
```

```
Sdp Bind 200:700 Instance 3
  Version          : 0x0
  PW OAM Msg Type  : 0x27
  Refresh Time     : 0x1e
  Total TLV Length : 0x8
  Flags            : 0x0
  TLV Type         : 0x96a
  TLV Len          : 0x4
  PW Status Bits   : 0x0
"
```

RSVP

The Resource Reservation Protocol (RSVP) is a network control protocol used by a host to request specific qualities of service from the network for particular application data streams or flows. RSVP is also used by routers to deliver quality of service (QoS) requests to all nodes along the path(s) of the flows and to establish and maintain state to provide the requested service. RSVP requests generally result in resources reserved in each node along the data path. MPLS leverages this RSVP mechanism to set up traffic engineered LSPs. RSVP is not enabled by default and must be explicitly enabled.

RSVP requests resources for simplex flows. It requests resources only in one direction (unidirectional). Therefore, RSVP treats a sender as logically distinct from a receiver, although the same application process may act as both a sender and a receiver at the same time. Duplex flows require two LSPs, to carry traffic in each direction.

RSVP is not a routing protocol. RSVP operates with unicast and multicast routing protocols. Routing protocols determine where packets are forwarded. RSVP consults local routing tables to relay RSVP messages.

RSVP uses two message types to set up LSPs, PATH and RESV. [Figure 22](#) depicts the process to establish an LSP.

- The sender (the ingress LER (ILER)), sends PATH messages toward the receiver, (the egress LER (ELER)) to indicate the FEC for which label bindings are desired. PATH messages are used to signal and request label bindings required to establish the LSP from ingress to egress. Each router along the path observes the traffic type.

PATH messages facilitate the routers along the path to make the necessary bandwidth reservations and distribute the label binding to the router upstream.

- The ELER sends label binding information in the RESV messages in response to PATH messages received.
- The LSP is considered operational when the ILER receives the label binding information.




Figure 22: Establishing LSPs

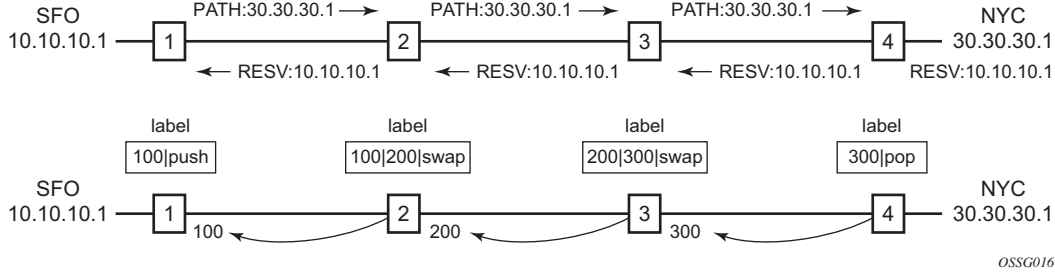


Figure 23: LSP Using RSVP Path Set Up

Figure 23 displays an example of an LSP path set up using RSVP. The ingress label edge router (ILER 1) transmits an RSVP path message (path: 30.30.30.1) downstream to the egress label edge router (ELER 4). The path message contains a label request object that requests intermediate LSRs and the ELER to provide a label binding for this path.

In addition to the label request object, an RSVP PATH message can also contain a number of optional objects:

- Explicit route object (ERO) — When the ERO is present, the RSVP path message is forced to follow the path specified by the ERO (independent of the IGP shortest path).
- Record route object (RRO) — Allows the ILER to receive a listing of the LSRs that the LSP tunnel actually traverses.
- A session attribute object controls the path set up priority, holding priority, and local-rerouting features.

Upon receiving a path message containing a label request object, the ELER transmits a RESV message that contains a label object. The label object contains the label binding that the downstream LSR communicates to its upstream neighbor. The RESV message is sent upstream towards the ILER, in a direction opposite to that followed by the path message. Each LSR that processes the RESV message carrying a label object uses the received label for outgoing traffic associated with the specific LSP. When the RESV message arrives at the ingress LSR, the LSP is established.

Using RSVP for MPLS

Hosts and routers that support both MPLS and RSVP can associate labels with RSVP flows. When MPLS and RSVP are combined, the definition of a flow can be made more flexible. Once an LSP is established, the traffic through the path is defined by the label applied at the ingress node of the LSP. The mapping of label to traffic can be accomplished using a variety of criteria. The set of packets that are assigned the same label value by a specific node are considered to belong to the same FEC which defines the RSVP flow.

For use with MPLS, RSVP already has the resource reservation component built-in which makes it ideal to reserve resources for LSPs.

RSVP Traffic Engineering Extensions for MPLS

RSVP has been extended for MPLS to support automatic signaling of LSPs. To enhance the scalability, latency, and reliability of RSVP signaling, several extensions have been defined. Refresh messages are still transmitted but the volume of traffic, the amount of CPU utilization, and response latency are reduced while reliability is supported. None of these extensions result in backward compatibility problems with traditional RSVP implementations.

- [Hello Protocol on page 87](#)
 - [MD5 Authentication of RSVP Interface on page 88](#)
 - [RSVP Overhead Refresh Reduction on page 90](#)
-

Hello Protocol

The Hello protocol detects the loss of a neighbor node or the reset of a neighbor's RSVP state information. In standard RSVP, neighbor monitoring occurs as part of RSVP's soft-state model. The reservation state is maintained as cached information that is first installed and then periodically refreshed by the ingress and egress LSRs. If the state is not refreshed within a specified time interval, the LSR discards the state because it assumes that either the neighbor node has been lost or its RSVP state information has been reset.

The Hello protocol extension is composed of a hello message, a hello request object and a hello ACK object. Hello processing between two neighbors supports independent selection of failure detection intervals. Each neighbor can automatically issue hello request objects. Each hello request object is answered by a hello ACK object.

MD5 Authentication of RSVP Interface

When enabled on an RSVP interface, authentication of RSVP messages operates in both directions of the interface.

A node maintains a security association with its neighbors for each authentication key. The following items are stored in the context of this security association:

- The HMAC-MD5 authentication algorithm.
- Key used with the authentication algorithm.
- Lifetime of the key. A key is user-generated key using a third party software/hardware and enters the value as static string into CLI configuration of the RSVP interface. The key will continue to be valid until it is removed from that RSVP interface.
- Source Address of the sending system.
- Latest sending sequence number used with this key identifier.

The RSVP sender transmits an authenticating digest of the RSVP message, computed using the shared authentication key and a keyed-hash algorithm. The message digest is included in an Integrity object which also contains a Flags field, a Key Identifier field, and a Sequence Number field. The RSVP sender complies to the procedures for RSVP message generation in RFC 2747, *RSVP Cryptographic Authentication*.

An RSVP receiver uses the key together with the authentication algorithm to process received RSVP messages.

When a PLR node switches the path of the LSP to a bypass LSP, it does not send the Integrity object in the RSVP messages over the bypass tunnel. If an integrity object is received from the MP node, then the message is discarded since there is no security association with the next-next-hop MP node.

The MD5 implementation does not support the authentication challenge procedures in RFC 2747.

Reservation Styles

LSPs can be signaled with explicit reservation styles. A reservation style is a set of control options that specify a number of supported parameters. The style information is part of the LSP configuration. SR OS supports two reservation styles:

- **Fixed Filter (FF)** — The Fixed Filter (FF) reservation style specifies an explicit list of senders and a distinct reservation for each of them. Each sender has a dedicated reservation that is not shared with other senders. Each sender is identified by an IP address and a local identification number, the LSP ID. Because each sender has its own reservation, a unique label and a separate LSP can be constructed for each sender-receiver pair. For traditional RSVP applications, the FF reservation style is ideal for a video distribution application in which each channel (or source) requires a separate pipe for each of the individual video streams.
- **Shared Explicit (SE)** — The Shared Explicit (SE) reservation style creates a single reservation over a link that is shared by an explicit list of senders. Because each sender is explicitly listed in the RESV message, different labels can be assigned to different sender-receiver pairs, thereby creating separate LSPs.

Note that if FRR option is enabled for the LSP and selects the facility FRR method at the head-end node, only the SE reservation style is allowed. Furthermore, if a PLR node receives a path message with fast-reroute requested with facility method and the FF reservation style, it will reject the reservation. The one-to-one detour method supports both FF and SE styles.

RSVP Message Pacing

When a flood of signaling messages arrive because of topology changes in the network, signaling messages can be dropped which results in longer set up times for LSPs. RSVP message pacing controls the transmission rate for RSVP messages, allowing the messages to be sent in timed intervals. Pacing reduces the number of dropped messages that can occur from bursts of signaling messages in large networks.

RSVP Overhead Refresh Reduction

The RSVP refresh reduction feature consists of the following capabilities implemented in accordance to RFC 2961, *RSVP Refresh Overhead Reduction Extensions*:

- RSVP message bundling — This capability is intended to reduce overall message handling load. The system supports receipt and processing of bundled message only, but no transmission of bundled messages.
- Reliable message delivery: — This capability consists of sending a message-id and returning a message-ack for each RSVP message. It can be used to detect message loss and support reliable RSVP message delivery on a per hop basis. It also helps reduce the refresh rate since the delivery becomes more reliable.
- Summary refresh — This capability consists of refreshing multiples states with a single message-id list and sending negative ACKs (NACKs) for a message_id which could not be matched. The summary refresh capability reduce the amount of messaging exchanged and the corresponding message processing between peers. It does not however reduce the amount of soft state to be stored in the node.

These capabilities can be enabled on a per-RSVP-interface basis are referred to collectively as “refresh overhead reduction extensions”. When the refresh-reduction is enabled on a system RSVP interface, the node indicates this to its peer by setting a refresh-reduction- capable bit in the flags field of the common RSVP header. If both peers of an RSVP interface set this bit, all the above three capabilities can be used. Furthermore, the node monitors the settings of this bit in received RSVP messages from the peer on the interface. As soon as this bit is cleared, the node stops sending summary refresh messages. If a peer did not set the “refresh-reduction-capable” bit, a node does not attempt to send summary refresh messages.

The RSVP Overhead Refresh Reduction is supported with both RSVP P2P LSP path and the S2L path of an RSVP P2MP LSP instance over the same RSVP interface.

RSVP Graceful Restart Helper

This **gr-helper** command enables the RSVP Graceful Restart Helper feature.

The RSVP-TE Graceful Restart helper mode allows the SR OS based system (the helper node) to provide another router that has requested it (the restarting node) a grace period, during which the system will continue to use RSVP sessions to neighbors requesting the grace period. This is typically used when another router is rebooting its control plane but its forwarding plane is expected to continue to forward traffic based on the previously available Path and Resv states.

The user can enable Graceful Restart helper on each RSVP interface separately. When the GR helper feature is enabled on an RSVP interface, the node starts inserting a new Restart_Cap Object in the Hello packets to its neighbor. The restarting node does the same and indicates to the helper node the desired Restart Time and Recovery Time.

The GR Restart helper consists of a couple of phases. Once it loses Hello communication with its neighbor, the helper node enters the Restart phase. During this phase, it preserves the state of all RSVP sessions to its neighbor and waits for a new Hello message.

Once the Hello message is received indicating the restarting node preserved state, the helper node enters the recovery phase in which it starts refreshing all the sessions that were preserved. The restarting node will activate all the stale sessions that are refreshed by the helper node. Any Path state that did not get a Resv message from the restarting node once the Recovery Phase time is over is considered to have expired and is deleted by the helper node causing the proper Path Tear generation downstream.

The duration of the restart phase (recovery phase) is equal to the minimum of the neighbor's advertised Restart Time (Recovery Time) in its last Hello message and the locally configured value of the max-restart (max-recovery) parameter.

When GR helper is enabled on an RSVP interface, its procedures apply to the state of both P2P and P2MP RSVP LSP to a neighbor over this interface.

Enhancements to RSVP control plane congestion control

The RSVP control plane makes use of a global flow control mechanism to adjust the rate of Path messages for unmapped LSP paths sent to the network under congestion conditions. When a Path message for establishing a new LSP path or retrying an LSP path that failed is sent out, the control plane keeps track of the rate of successful establishment of these paths and adjusts the number of Path messages it sends per second to reflect the success ratio.

In addition, an option to enable an exponential back-off retry-timer is available. When an LSP path establishment attempt fails, the path is put into retry procedures and a new attempt will be performed at the expiry of the user-configurable retry-timer. By default, the retry time is constant. The exponential back-off timer procedures will double the value of the user configurable retry-timer value at every failure of the attempt to adjust to the potential network congestion that caused the failure. An LSP establishment fails if no Resv message was received and the Path message retry-timer expired, or a PathErr message was received before the timer expired.

Three enhancements to this flow-control mechanism to improve congestion handling in the rest of the network are supported.

The first enhancement is the change to the LSP path retry procedure. If the establishment attempt failed due to a Path message timeout and no Resv was received, the next attempt will be performed at the expiry of a new LSP path initial retry-timer instead of the existing retry-timer. While the LSP path initial retry-timer is still running, a refresh of the Path message using the same path and the same LSP-id is performed according to the configuration of the refresh-timer. Once the LSP path initial retry-timer expires, the ingress LER then puts this path on the regular retry-timer to schedule the next path signaling using a new computed path by CSPF and a new LSP-id.

The benefits of this enhancement is that the user can now control how many refreshes of the pending PATH state can be performed before starting a new retry-cycle with a new LSP-id. This is all done without affecting the ability to react faster to failures of the LSP path, which will continue to be governed by the existing retry-timer. By configuring the LSP path initial retry-timer to values that are larger than the retry-timer, the ingress LER will decrease the probability of overwhelming a congested LSR with new state while the previous states installed by the same LSP are lingering and will only be removed after the refresh timeout period expires.

The second enhancement consists of applying a jitter +/- 25% to the value of the retry-timer similar to how it is currently done for the refresh timer. This will further decrease the probability that ingress LER nodes synchronize their sending of Path messages during the retry-procedure in response to a congestion event in the network.

The third enhances the RSVP flow control mechanism by taking into account new parameters: outstanding CSPF requests, Resv timeouts and Path timeouts.

RSVP LSP Statistics

This feature provides the following counters:

- Per forwarding class forwarded in-profile packet count
- Per forwarding class forwarded in-profile byte count
- Per forwarding class forwarded out of profile packet count
- Per forwarding class forwarded out of profile byte count

The counters are available for an RSVP LSP at the egress datapath of an ingress LER and at the ingress datapath of an egress LER. No LSR statistics are provided.

P2MP RSVP-TE LSP Statistics

This feature provides the following counters for a RSVP P2MP LSP instance:

- Per forwarding class forwarded in-profile packet count.
- Per forwarding class forwarded in-profile byte count.
- Per forwarding class forwarded out of profile packet count.
- Per forwarding class forwarded out of profile byte count.

The above counters are provided for the following LSR roles:

1. At ingress LER, a set of per P2MP LSP instance counters for packets forwarded to the P2MP LSP instance without counting the replications is provided. In other words, a packet replicated over multiple branches of the same P2MP LSP instance will count once as long as at least one LSP branch forwarded it.
2. At BUD LSR and egress LER, per ILM statistics are provided. These counters will include all packets received on the ILM, whether they match a L2/L3 MFIB record or not. ILM stats will work the same way as for a P2P LSP. In other words, they will count all packets received on the primary ILM, including packets received over the bypass LSP.

When MBB is occurring for an S2L path of an RSVP P2MP LSP, paths of the new and old S2L will both receive packets on the egress LER. Both packets are forwarded to the fabric and outgoing PIM/IGMP interfaces until the older path is torn down by the ingress LER. In this case, packet duplication should be counted.

3. No branch LSR statistics are provided.
4. The P2MP LSP statistics share the same pool of counters and stat indices the P2P LSP share on the node. Each P2P/P2MP RSVP LSP or LDP FEC consumes one stat index for egress stats and one stat index for ingress stats.
5. The user can retrieve the above counters in four different ways:
 - In CLI display of the output of the show command applied to a specific instance, or a specific template instance, of a RSVP P2MP.
 - In CLI display of the output of the monitor command applied to a specific instance, or a specific template instance, of a RSVP P2MP.
 - Via an SNMP interface by querying the MIB.
 - Via an accounting file if statistics collection with the default or user specified accounting policy is enabled for the MPLS LSP stats configuration contexts.
6. OAM packets that are forwarded using the LSP encapsulation, for example, P2MP LSP Ping and P2MP LSP Trace, are also included in the above counters.

The user can determine if packets are dropped for a given branch of a P2MP RSVP LSP by comparing the egress counters at the ingress LER with the ILM counters at the egress LER or BUD LSR.

Octet counters are for the entire frame and thus include the label stack and the L2 header and padding similar to the existing P2P RSVP LSP and LDP FEC counters. Thus ingress and egress octet counters for an LSP may slightly differ if the type of interface or encapsulation is different (POS, Ethernet NULL, Ethernet Dot1.Q).

Configuring RSVP P2MP LSP Egress Statistics

At ingress LER, the configuration of the egress statistics is under the MPLS P2MP LSP context when carrying multicast packets over a RSVP P2MP LSP in the base routing instance. This is the same configuration as the one already supported with P2P RSVP LSP.

```
config
  router
    [no] mpls
      [no] lsp lsp-name p2mp-lsp
        [no] egress-statistics
          accounting-policy policy-id
          no accounting-policy
          [no] collect-stats
        [no] shutdown
```

If there are no stat indices available when the user performs the 'no shutdown' command for the egress statistics node, the command will be failed.

The configuration is in the P2MP LSP template when the RSVP P2MP LSP is used as an I-PMSI or S-PMSI in multicast VPN or in VPLS/B-VPLS.

```
config
  router
    [no] mpls
      lsp-template template-name p2mp
      no lsp-template template-name
        [no] egress-statistics
          accounting-policy policy-id
          no accounting-policy
          [no] collect-stats
```

If there are no stat indices available at the time an instance of the P2MP LSP template is signaled, no stats are allocated to the instance, but the LSP is brought up. In this case, an operational state of out-of-resources is shown for the egress stats in the show output of the P2MP LSP S2L path.

Configuring RSVP P2MP LSP Ingress Statistics

When the ingress LER signals the path of the S2L sub-LSP, it includes the name of the LSP and that of the path in the Session Name field of the Session Attribute object in the Path message. The encoding is as follows:

Session Name: *lsp-name::path-name*, where *lsp-name* component is encoded as follows:

1. P2MP LSP via user configuration for L3 multicast in global routing instance: “LspNameFromConfig”
2. P2MP LSP as I-PMSI or S-PMSI in L3 mVPN: *templateName-SvcId-mTTmIndex*
3. P2MP LSP as I-PMSI in VPLS/B-VPLS: *templateName-SvcId-mTTmIndex*

The ingress statistics CLI configuration allows the user to match either on the exact name of the P2MP LSP as configured at the ingress LER or on a context which matches on the template name and the service-id as configured at the ingress LER.

```
config
  router
    [no] mpls
      ingress-statistics
        [no] lsp lsp-name sender sender-address
          accounting-policy policy-id
          no accounting-policy
          [no] collect-stats
          [no] shutdown

        [no] p2mp-template-lsp rsvp-session-name
          SessionNameString sender sender-address
          accounting-policy policy-id
          no accounting-policy
          [no] collect-stats
          max-stats integer<1-8192|max, default max>
          no max-stats
          [no] shutdown
```

When the matching is performed on a context, the user must enter the RSVP session name string in the format “*templateName-svcId*” to include the LSP template name as well as the mVPN VPLS/B-VPLS service ID as configured at the ingress LER. In this case, one or more P2MP LSP instances signaled by the same ingress LER could be associated with the ingress statistics configuration. In this case, the user is provided with CLI parameter **max-stats** to limit the maximum number of stat indices which can be assigned to this context. If the context matches more than this value, the additional request for stat indices from this context will be rejected.

Note: The rules when configuring an ingress statistics context based on template matching are the following:

1. **max-stats** once allocated can be increased but not decreased unless the entire ingress statistics context matching a template name is deleted.
2. In order to delete ingress statistics context matching a template name, a shutdown is required.
3. An accounting policy cannot be configured or de-configured until the ingress statistics context matching a template name is shutdown.
4. After deleting an accounting policy from an ingress statistics context matching a template name, the policy is not removed from the log until a 'no shut' is performed on the ingress statistics context.

If there are no stat indices available at the time the session of the P2MP LSP matching a template context is signaled and the session state installed by the egress LER, no stats are allocated to the session.

Furthermore, the assignment of stat indices to the LSP names that match the context will also be not deterministic. The latter is due to the fact that a stat index is assigned and released following the dynamics of the LSP creation or deletion by the ingress LER. For example, a multicast stream crosses the rate threshold and is moved to a newly signaled S-PMSI dedicated to this stream. Later on, the same stream crosses the threshold downwards and is moved back to the shared I-PMSI and the P2MP LSP corresponding to the S-PMSI is deleted by the ingress LER.

Configuring Implicit Null

The implicit null label option allows a 7x50 egress LER to receive MPLS packets from the previous hop without the outer LSP label. The operation of the previous hop is referred to as penultimate hop popping (PHP).

This option is signaled by the egress LER to the previous hop during the LSP signaling with RSVP control protocol. In addition, the egress LER can be configured to receive MPLS packet with the implicit null label on a static LSP.

The user can configure your router to signal the implicit null label value over all RSVP interfaces and for all RSVP LSPs for which this node is the egress LER using the **implicit-null-label** command in the **config>router>rsvp** context.

The user must shutdown RSVP before being able to change the implicit null configuration option.

The user can also override the RSVP level configuration for a specific RSVP interface:

```
config>router>rsvp>interface>implicit-null-label {enable | disable}
```

All LSPs for which this node is the egress LER and for which the path message is received from the previous hop node over this RSVP interface will signal the implicit null label. This means that if the egress LER is also the merge-point (MP) node, then the incoming interface for the path refresh message over the bypass dictates if the packet will use the implicit null label or not. The same for a 1-to-1 detour LSP.

By default, an RSVP interface inherits the RSVP level configuration. The user must shutdown the RSVP interface before being able to change the implicit null configuration option. Note that the RSVP interface must be shutdown regardless if the new value for the interface is the same or different than the one it is currently using.

The egress LER does not signal the implicit null label value on P2MP RSVP LSPs. However, the PHP node can honor a Resv message with the label value set to the implicit null value when the egress LER is a third party implementation.

The implicit null label option is also supported on a static label LSP. The following commands can be used to cause the node to push or to swap to an implicit null label on the MPLS packet:

```
config>router>mpls>static-lsp>push implicit-null-label nexthop ip-address
```

```
config>router>mpls>interface>label-map>swap implicit-null-label nexthop ip-address
```

Using Unnumbered Point-to-Point Interface in RSVP

This feature introduces the use of unnumbered IP interface as a Traffic Engineering (TE) link for the signaling of RSVP P2P LSP and P2MP LSP.

An unnumbered IP interface is identified uniquely on a router in the network by the tuple {router-id, ifIndex}. Each side of the link assigns a system-wide unique interface index to the unnumbered interface. ISIS, OSPF, RSVP, and OAM modules will use this tuple to advertise the link information, signal LSP paths over this unnumbered interface, or send and respond to an MPLS echo request message over an unnumbered interface.

The interface borrowed IP address is used exclusively as the source address for IP packets that are originated from the interface and needs to be configured to an address different from system interface for the FRR bypass LSP to come up at the ingress LER.

The borrowed IP address for an unnumbered interface is configured using the following CLI command with a default value set to the system interface address:

```
configure> router>interface>unnumbered [ip-int-name | ip-address].
```

The support of unnumbered TE link in IS-IS consists of adding a new sub-TLV of the extended IS reachability TLV, which encodes the Link Local and Link Remote Identifiers as defined in RFC 5307.

The support of unnumbered TE link in OSPF consists of adding a new sub-TLV, which encodes the same Link Local and Link Remote Identifiers in the Link TLV of the TE area opaque LSA and sends the local Identifier in the Link Local Identifier TLV in the TE link local opaque LSA as per RFC 4203.

The support of unnumbered TE link in RSVP implements the signaling of unnumbered interfaces in ERO/RRO as per RFC 3477 and the support of IF_ID RSVP_HOP object with a new CType as per Section 8.1.1 of RFC 3473. The IPv4 Next/Previous Hop Address field is set to the borrowed IP interface address.

The unnumbered IP is advertised by IS-IS TE and OSPF TE, and CSPF can include them in the computation of a path for a P2P LSP or for the S2L of a P2MP LSP. This feature does not, however, support defining an unnumbered interface a hop in the path definition of an LSP.

A router creates an RSVP neighbor over an unnumbered interface using the tuple {router-id, ifIndex}. The router-id of the router that advertised a given unnumbered interface index is obtained from the TE database. As a result, if traffic engineering is disabled in IS-IS or OSPF, a non-CSPF LSP with the next-hop for its path is over an unnumbered interface will not come up at the ingress LER since the router-id of the neighbor that has the next-hop of the path message cannot be looked up. In this case, the LSP path will remain in operationally down state with a reason noRouteToDestination. If a PATH message was received at the LSR in which traffic engineering was disabled and the next-hop for the LSP path is over an unnumbered interface, a

PathErr message will be sent back to the ingress LER with the *Routing Problem* error code of 24 and an error value of 5 “No route available toward destination”.

All MPLS features available for numbered IP interfaces are supported, with the exception of the following:

- Configuring a router-id with a value other than system.
- Signaling of an LSP path with an ERO based a loose/strict hop using an unnumbered TE link in the path hop definition.
- Signaling of one-to-one detour LSP over unnumbered interface.
- Unnumbered RSVP interface registration with BFD.
- RSVP Hello and all Hello related capabilities such as Graceful-restart helper.
- The user SRLG database feature. The user-srlg-db option under MPLS allows the user to manually enter the SRLG membership of any link in the network in a local database at the ingress LER. The user cannot enter an unnumbered interface into this database and as such, all unnumbered interfaces will be considered as having no SRLG membership if the user enabled the user-srlg-db option.

This feature also extends the support of lsp-ping, p2mp-lsp-ping, lsp-trace, and p2mp-lspttrace to P2P and P2MP LSPs that have unnumbered TE links in their path.

Operation of RSVP FRR Facility Backup over Unnumbered Interface

When the Point-of-Local Repair (PLR) node activates the bypass LSP by sending a PATH message to refresh the path state of protected LSP at the Merge-Point (MP) node, it must use an *IPv4 tunnel sender address* in the sender template object that is different than the one used by the ingress LER in the PATH message. These are the procedures specified in RFC 4090 and that are followed in the 7x50 implementation.

The 7x50 uses the address of the outgoing interface of the bypass LSP as the *IPv4 tunnel sender address* in the sender template object. This address will be different from the system interface address used in the sender template of the protected LSP by the ingress LER and thus there are no conflicts when the ingress LER acts as a PLR.

When the PLR is the ingress LER node and the outgoing interface of the bypass LSP is unnumbered, it is required that the user assigns to the interface a borrowed IP address that is different from the system interface. If not, the bypass LSP will not come up.

In addition, the PLR node will include the IPv4 RSVP_HOP object (C-Type=1) or the IF_ID RSVP_HOP object (C-Type=3) in the PATH message if the outgoing interface of the bypass LSP is numbered or unnumbered respectively.

When the MP node receives the PATH message over the bypass LSP, it will create the merge-point context for the protected LSP and associate it with the existing state if any of the following is satisfied:

- Change in C-Type of the RSVP_HOP object, or
- C-Type is IF_ID RSVP_HOP and did not change but IF_ID TLV is different, or
- Change in IPv4 Next/Previous Hop Address in RSVP_HOP object regardless of the C-Type value.

These procedures at PLR and MP nodes are followed in both link-protect and node-protect FRR. Note that if the MP node is running a pre-R11 implementation, it will reject the new IF_ID C-Type and will drop the PATH over bypass. This will result in the protected LSP state expiring at the MP node, which will tear down the path. This will be the case in general when node-protect FRR is enabled and the MP node does not support unnumbered RSVP interface.

Traffic Engineering

Without traffic engineering, routers route traffic according to the SPF algorithm, disregarding congestion or packet types.

With traffic engineering, network traffic is routed efficiently to maximize throughput and minimize delay. Traffic engineering facilitates traffic flows to be mapped to the destination through a different (less congested) path other than the one selected by the SPF algorithm.

MPLS directs a flow of IP packets along a label switched path (LSP). LSPs are simplex, meaning that the traffic flows in one direction (unidirectional) from an ingress router to an egress router. Two LSPs are required for duplex traffic. Each LSP carries traffic in a specific direction, forwarding packets from one router to the next across the MPLS domain.

When an ingress router receives a packet, it adds an MPLS header to the packet and forwards it to the next hop in the LSP. The labeled packet is forwarded along the LSP path until it reaches the destination point. The MPLS header is removed and the packet is forwarded based on Layer 3 information such as the IP destination address. The physical path of the LSP is not constrained to the shortest path that the IGP would choose to reach the destination IP address.

TE Metric (IS-IS and OSPF)

When the use of the TE metric is selected for an LSP, the shortest path computation after the TE constraints are applied will select an LSP path based on the TE metric instead of the IGP metric. The user configures the TE metric under the MPLS interface. Both the TE and IGP metrics are advertised by OSPF and IS-IS for each link in the network. The TE metric is part of the traffic engineering extensions of both IGP protocols.

A typical application of the TE metric is to allow CSPF to represent a dual TE topology for the purpose of computing LSP paths.

An LSP dedicated for real-time and delay sensitive user and control traffic has its path computed by CSPF using the TE metric. The user configures the TE metric to represent the delay figure, or a combined delay/jitter figure, of the link. In this case, the shortest path satisfying the constraints of the LSP path will effectively represent the shortest delay path.

An LSP dedicated for non delay sensitive user and control traffic has its path computed by CSPF using the IGP metric. The IGP metric could represent the link bandwidth or some other figure as required.

When the use of the TE metric is enabled for an LSP, CSPF will first prune all links in the network topology that do not meet the constraints specified for the LSP path. These constraints include bandwidth, admin-groups, and hop limit. CSPF will then run an SPF on the remaining links. The shortest path among the all SPF paths will be selected based on the TE metric instead of the IGP metric which is used by default. Note that the TE metric is only used in CSPF computations for MPLS paths and not in the regular SPF computation for IP reachability.

Admin Group Support on Facility Bypass Backup LSP

This feature provides for the inclusion of the LSP primary path admin-group constraints in the computation of a Fast ReRoute (FRR) facility bypass backup LSP to protect the primary LSP path by all nodes in the LSP path.

This feature is supported with the following LSP types and in both intra-area and inter-area TE where applicable:

- Primary path of a RSVP P2P LSP.
 - S2L path of an RSVP P2MP LSP instance
 - LSP template for an S2L path of an RSVP P2MP LSP instance.
 - LSP template for auto-created RSVP P2P LSP in intra-area TE.
-

Procedures at Head-End Node

The user enables the signaling of the primary LSP path admin-group constraints in the FRR object at the ingress LER with the following CLI command:

```
configure>router>mpls>lsp>fast-reroute>propagate-admin-group
```

When this command is enabled at the ingress LER, the admin-group constraints configured in the context of the P2P LSP primary path, or the ones configured in the context of the LSP and inherited by the primary path, are copied into the FAST_REROUTE object. The admin-group constraints are copied into the *include-any* or *exclude-any* fields.

The ingress LER thus propagates these constraints to the downstream nodes during the signaling of the LSP to allow them to include the admin-group constraints in the selection of the FRR backup LSP for protecting the LSP primary path.

The ingress LER will insert the FAST_REROUTE object by default in a primary LSP path message. If the user disables the object using the following command, the admin-group constraints will not be propagated: **configure>router>mpls>no frr-object**.

Note that the same admin-group constraints can be copied into the Session Attribute object. They are intended for the use of an LSR, typically an ABR, to expand the ERO of an inter-area LSP path. They are also used by any LSR node in the path of a CSPF or non-CSPF LSP to check the admin-group constraints against the ERO regardless if the hop is strict or loose. These are governed strictly by the command:

```
configure>router>mpls>lsp>propagate-admin-group
```


In other words, the user may decide to copy the primary path admin-group constraints into the FAST_REROUTE object only, or into the Session Attribute object only, or into both.

Note however, that the PLR rules for processing the admin-group constraints can make use of either of the two object admin-group constraints.

Procedures at PLR Node

The user enables the use of the admin-group constraints in the association of a manual or dynamic bypass LSP with the primary LSP path at a Point-of-Local Repair (PLR) node using the following global command:

```
configure>router>mpls>admin-group-frr
```

When this command is enabled, each PLR node reads the admin-group constraints in the FAST_REROUTE object in the Path message of the LSP primary path. If the FAST_REROUTE object is not included in the Path message, then the PLR will read the admin-group constraints from the Session Attribute object in the Path message.

If the PLR is also the ingress LER for the LSP primary path, then it just uses the admin-group constraint from the LSP and/or path level configurations.

Whether the PLR node is also the ingress LER or just an LSR for the protected LSP primary path, the outcome of the ingress LER configuration dictates the behavior of the PLR node and is summarized in [Table 4](#).

Table 4: Bypass LSP Admin-Group Constraint Behavior

	Ingress LER Configuration	Session Attribute	FRR Object	Bypass LSP at PLR (LER/LSF) follows admin-group constraints
1	frr-object lsp>no propagate-admin-group lsp>frr>propagate-admin-group	Admin color constraints not sent	Admin color constraints sent	yes
2	frr-object lsp>propagate-admin-group lsp>frr>propagate-admin-group	Admin color constraints sent	Admin color constraints sent	yes

Table 4: Bypass LSP Admin-Group Constraint Behavior

3	frr-object lsp>propagate-admin group lsp>frr>no propagate- admin-group	Admin color constraints sent	Admin color constraints not sent	no
4	frr-object lsp>propagate-admin group lsp>frr>no propagate- admin-group	Admin color constraints sent	Not present	yes
5	No frr-object lsp>no propagate-admin group lsp>frr>propagate-admin- group	Admin color constraints not sent	Not present	no
6	No frr-object lsp>propagate-admin group lsp>frr>no propagate- admin-group	Admin color constraints sent	Not present	yes

The PLR node then uses the admin-group constraints along with other constraints, such as hop-limit and SRLG, to select a manual or dynamic bypass among those that are already in use.

If none of the manual or dynamic bypass LSP satisfies the admin-group constraints, and/or the other constraints, the PLR node will request CSPF for a path that merges the closest to the protected link or node and that includes or excludes the specified admin-group IDs.

If the user changes the configuration of the above command, it will not have any effect on existing bypass associations. The change will only apply to new attempts to find a valid bypass.

Diff-Serv Traffic Engineering

Diff-Serv traffic engineering provides the ability to manage bandwidth on a per Traffic Engineering (TE) class basis as per RFC 4124. In the base traffic engineering, LER computes LSP paths based on available BW of links on the path. Diff-Serv TE adds ability to perform this on a per TE class basis.

A TE class is a combination of Class Type and LSP priority. A Class Type is mapped to one or more system Forwarding Classes using a configuration profile. The operator sets different limits for admission control of LSPs in each TE class over each TE link. Eight TE classes are supported. Admission control of LSP paths bandwidth reservation is performed using the Maximum Allocation Bandwidth Constraint Model as per RFC 4125.

Mapping of Traffic to a Diff-Serv LSP

An LER will allow the operator to map traffic to a Diff-Serv LSP through one of the following methods:

1. Explicit RSVP SDP configuration of a VLL, VPLS, or VPRN service.
 2. Class-based forwarding in an RSVP SDP. The operator can enable the checking by RSVP that a Forwarding Class (FC) mapping to an LSP under the SDP configuration is compatible with the Diff-Serv Class Type (CT) configuration for this LSP.
 3. Auto-bind RSVP-TE option in a VPRN service.
 4. Static routes with indirect next-hop being an RSVP LSP name.
-

Admission Control of Classes

There are a couple of admission control decisions made when an LSP with a specified bandwidth is to be signaled. The first is in the head-end node. CSPF will only consider network links that have sufficient bandwidth. Link bandwidth information is provided by IGP TE advertisement by all nodes in that network.

Another decision made is local CAC and is performed when the RESV message for the LSP path is received in the reverse direction by a SR OS node in that path. The bandwidth value selected by the egress LER will be checked against link bandwidth, otherwise the reservation is rejected. If accepted, the new value for the remaining link bandwidth will be advertised by IGP at the next advertisement event.

Both of these admission decisions are enhanced to be performed at the TE class level when Diff-Serv TE is enabled. In other words, CSPF in the head-end node will need to check the LSP bandwidth against the 'unreserved bandwidth' advertised for all links in the path of the LSP for that TE class which consists of a combination of a CT and a priority. Same for the admission control at SR OS node receiving the Resv message.

Maximum Allocation Model

The admission control rules for this model are described in RFC 4125. Each CT shares a percentage of the Maximum Reservable Link Bandwidth through the user-configured BC for this CT. The Maximum Reservable Link Bandwidth is the link bandwidth multiplied by the RSVP interface subscription factor.

The sum of all BC values across all CTs will not exceed the Maximum Reservable Link Bandwidth. In other words, the following rule is enforced:

$$\text{SUM}(BC_c) \leq \text{Max-Reservable-Bandwidth}, 0 \leq c \leq 7$$

An LSP of class-type CT_c, setup priority p, holding priority h (h ≤ p), and bandwidth B is admitted into a link if the following condition is satisfied:

$$B \leq \text{Unreserved Bandwidth for TE-Class}[i]$$

where TE-Class [i] maps to < CT_c, p > in the definition of the TE classes on the node. The bandwidth reservation is effected at the holding priority, i.e., in TE-class [j] = < CT_c, h >. Thus, the reserved bandwidth for CT_c and the unreserved bandwidth for the TE classes using CT_c are updated as follows:

$$\text{Reserved}(CT_c) = \text{Reserved}(CT_c) + B$$

$$\text{Unreserved TE-Class [j]} = BC_c - \text{SUM}(\text{Reserved}(CT_c, q)) \text{ for } 0 \leq q \leq h$$

$$\text{Unreserved TE-Class [i]} = BC_c - \text{SUM}(\text{Reserved}(CT_c, q)) \text{ for } 0 \leq q \leq p$$

The same is done to update the unreserved bandwidth for any other TE class making use of the same CT_c. These new values are advertised to the rest of the network at the next IGP-TE flooding.

When Diff-Serv is disabled on the node, this model degenerates into a single default CT internally with eight pre-emption priorities and a non-configurable BC equal to the Maximum Reservable Link Bandwidth. This would behave exactly like CT₀ with eight pre-emption priorities and BC = Maximum Reservable Link Bandwidth if Diff-Serv was enabled.

Russian Doll Model

The RDM model is defined using the following equations:

$$\text{SUM (Reserved (CTc))} \leq \text{BCb},$$

where the SUM is across all values of c in the range $b \leq c \leq (\text{MaxCT} - 1)$, and BCb is the bandwidth constraint of CTb .

$\text{BC0} = \text{Max-Reservable-Bandwidth}$, so that:

$$\text{SUM (Reserved(CTc))} \leq \text{Max-Reservable-Bandwidth},$$

where the SUM is across all values of c in the range $0 \leq c \leq (\text{MaxCT} - 1)$

An LSP of class-type CTc , setup priority p , holding priority h ($h \leq p$), and bandwidth B is admitted into a link if the following condition is satisfied:

$$B \leq \text{Unreserved Bandwidth for TE-Class}[i],$$

where $\text{TE-Class}[i]$ maps to $\langle \text{CTc}, p \rangle$ in the definition of the TE classes on the node. The bandwidth reservation is effected at the holding priority, i.e., in $\text{TE-class}[j] = \langle \text{CTc}, h \rangle$. Thus, the reserved bandwidth for CTc and the unreserved bandwidth for the TE classes using CTc are updated as follows:

$$\text{Reserved(CTc)} = \text{Reserved(CTc)} + B$$

$$\begin{aligned} \text{Unreserved TE-Class}[j] = \text{Unreserved (CTc, h)} = \text{Min} [\\ & \text{BCc} - \text{SUM (Reserved (CTb, q) for } 0 \leq q \leq h, c \leq b \leq 7, \\ & \text{BC(c-1)} - \text{SUM (Reserved (CTb, q) for } 0 \leq q \leq h, (c-1) \leq b \leq 7, \\ & \dots\dots \\ & \text{BC0} - \text{SUM (Reserved (CTb, q) for } 0 \leq q \leq h, 0 \leq b \leq 7] \end{aligned}$$

$$\begin{aligned} \text{Unreserved TE-Class}[i] = \text{Unreserved (CTc, p)} = \text{Min} [\\ & \text{BCc} - \text{SUM (Reserved (CTb, q) for } 0 \leq q \leq p, c \leq b \leq 7, \\ & \text{BC(c-1)} - \text{SUM (Reserved (CTb, q) for } 0 \leq q \leq p, (c-1) \leq b \leq 7, \\ & \dots\dots \\ & \text{BC0} - \text{SUM (Reserved (CTb, q) for } 0 \leq q \leq p, 0 \leq b \leq 7] \end{aligned}$$

The same is done to update the unreserved bandwidth for any other TE class making use of the same CTc . These new values are advertised to the rest of the network at the next IGP-TE flooding.

Example CT Bandwidth Sharing with RDM

Below is a simple example with two CT values (CT0, CT1) and one priority 0 as shown in [Figure 24](#).

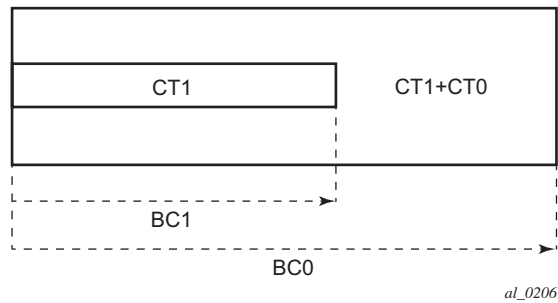


Figure 24: RDM with Two Class Types

Suppose CT1 bandwidth, or the CT1 percentage of Maximum Reservable Bandwidth to be more accurate is 100 Mbps and CT2 bandwidth is 100 Mbps and link bandwidth is 200 Mbps. BC constraints can be calculated as follows.

$BC1 = CT1 \text{ Bandwidth} = 100 \text{ Mbps}$.

$BC0 = \{CT1 \text{ Bandwidth}\} + \{CT0 \text{ Bandwidth}\} = 200 \text{ Mbps}$.

Suppose an LSP comes with CT1, setup and holding priorities of 0 and a bandwidth of 50 Mbps.




Figure 25: First LSP Reservation

According to the RDM admission control policy:

$$\text{Reserved (CT1, 0)} = 50 \leq 100 \text{ Mbps}$$

$$\text{Reserved (CT0, 0)} + \text{Reserved (CT1, 0)} = 50 \leq 200 \text{ Mbps}$$

This results in the following unreserved bandwidth calculation.

$$\text{Unreserved (CT1, 0)} = \text{BC1} - \text{Reserved (CT1, 0)} = 100 - 50 = 50 \text{ Mbps}$$

$$\text{Unreserved (CT0, 0)} = \text{BC0} - \text{Reserved (CT0, 0)} - \text{Reserved (CT1, 0)} = 200 - 0 - 50 = 150 \text{ Mbps.}$$

Note that bandwidth reserved by a doll is not available to itself as well any of the outer dolls.

Suppose now another LSP comes with CT0, setup and holding priorities of 0 and a bandwidth 120 Mbps.




Figure 26: Second LSP Reservation

$$\text{Reserved (CT0, 0)} = 120 \leq 150 \text{ Mbps}$$

$$\text{Reserved (CT0, 0)} + \text{Reserved (CT1, 0)} = 120 + 50 = 170 \leq 200 \text{ Mbps}$$

$$\text{Unreserved (CT0, 0)} = 150 - 120 = 30 \text{ Mbps}$$

If we simply checked BC1, the formula would yield the wrong results:

$$\text{Unreserved (CT1, 0)} = \text{BC1} - \text{Reserved (CT1, 0)} = 100 - 50 = 50 \text{ Mbps}$$

Because of the encroaching of CT0 into CT1, we would need to deduct the overlapping reservation. This would then yield:

$$\text{Unreserved (CT1, 0)} = \text{BC0} - \text{Reserved (CT0, 0)} - \text{Reserved (CT1, 0)} = 200 - 120 - 50 = 30 \text{ Mbps,}$$

which is the correct figure.

Extending the formula with both equations:

$$\text{Unreserved (CT1, 0)} = \text{Min} [\text{BC1} - \text{Reserved (CT1, 0)}, \text{BC0} - \text{Reserved (CT0, 0)} - \text{Reserved (CT1, 0)}] = \text{Min} [100 - 50, 200 - 120 - 50] = 30 \text{ Mbps}$$

Note that an outer doll can encroach into inner doll reducing the bandwidth available for inner dolls.

RSVP Control Plane Extensions

RSVP will use the Class Type object to carry LSP class-type information during path setup. Eight values will be supported for class-types 0 through 7 as per RFC 4124. Class type 0 is the default class which is supported today on the router.

One or more forwarding classes will map to a Diff-Serv class type through a system level configuration.

IGP Extensions

IGP extensions are defined in RFC 4124. Diff-Serv TE advertises link available bandwidth, referred to as unreserved bandwidth, by OSPF TE or IS-IS TE on a per TE class basis. A TE class is a combination of a class type and an LSP priority. In order to reduce the amount of per TE class flooding required in the network, the number of TE classes is set to eight. This means that eight class types can be supported with a single priority or four class types with two priorities, etc. In that case, the operator configures the desired class type on the LSP such that RSVP-TE can signal it in the class-type object in the path message.

IGP will continue to advertise the existing Maximum Reservable Link Bandwidth TE parameter to mean the maximum bandwidth that can be booked on a given interface by all classes. The value advertised is adjusted with the link subscription factor.

Diff-Serv TE Configuration and Operation

RSVP Protocol Level

The following are the configuration steps at the RSVP protocol level:

1. The operator enables Diff-Serv TE by executing the **diffserv-te** command in the **config>router>rsvp** context. When this command is enabled, IS-IS and OSPF will start advertising available bandwidth for each TE class configured under the **diffserv-te** node. The operator can disable Diff-Serv TE globally by using the no form of the command.
2. The enabling or disabling of Diff-Serv on the system requires that the RSVP and MPLS protocol be shutdown. The operator must execute the **no shutdown** command in each context once all parameters under both protocols are defined. When saved in the configuration file, the **no shutdown** command is automatically inserted under both protocols to make sure they come up after a node reboot.
3. IGP will advertise the available bandwidth in each TE class in the unreserved bandwidth TE parameter for that class for each RSVP interface in the system.
4. In addition, IGP will continue to advertise the existing Maximum Reservable Link Bandwidth TE parameter so the maximum bandwidth that can be booked on a given interface by all classes. The value advertised is adjusted with the link subscription factor configured in the **config>router>rsvp>interface>subscription percentage** context.
5. The operator can overbook (underbook) the maximum reservable bandwidth of a given CT by overbooking (underbooking) the interface maximum reservable bandwidth by configuring the appropriate value for the **subscription percentage** parameter.
6. The **diffserv-te** command will only have effect if the operator has already enabled traffic engineering at the IS-IS and/or OSPF routing protocol levels:
config>router>isis>traffic-engineering
and/or:
config>router>ospf>traffic-engineering
7. The following Diff-Serv TE parameters are configured globally under the **diffserv-te** node. They apply to all RSVP interfaces on the system. Once configured, these parameters can only be changed after shutting down the MPLS and RSVP protocols:
 - a. Definition of TE classes, TE Class = {Class Type (CT), LSP priority}. Eight TE classes can be supported. There is no default TE class once Diff-Serv is enabled. The operator must explicitly define each TE class. However, when Diff-Serv is disabled there will be an internal use of the default CT (CT0) and eight pre-emption priorities as shown in [Table 5](#).

Table 5: Internal TE Class Definition when Diff-Serv TE is Disabled

Class Type (CT internal)	LSP Priority
0	7
0	6
0	5
0	4
0	3
0	2
0	1
0	0

b. A mapping of the system forwarding class to CT. The default settings are shown in [Table 6](#).

Table 6: Default Mapping of Forwarding Class to TE Class

FC ID	FC Name	FC Designation	Class Type (CT)
7	Network Control	NC	7
6	High-1	H1	6
5	Expedited	EF	5
4	High-2	H2	4
3	Low-1	L1	3
2	Assured	AF	2
1	Low-2	L2	1
0	Best Effort	BE	0

c. Configuration of the percentage of RSVP interface bandwidth each CT shares, for example, the Bandwidth Constraint (BC), using the **class-type-bw** command. The absolute value of the CT share of the interface bandwidth is derived as the percentage of the bandwidth advertised by IGP in the maximum reservable link bandwidth TE parameter, for example, the link bandwidth multiplied by the RSVP interface **subscription percentage** parameter. Note that this configuration also exists at the RSVP

interface level and the interface specific configured value overrides the global configured value. The BC value can be changed at any time. The operator can specify the BC for a CT which is not used in any of the TE class definition but that does not get used by any LSP originating or transiting this node.

d. Configuration of the Admission Control Policy to be used: only the Maximum Allocation Model (MAM) is supported. The MAM value represents the bandwidth constraint models for the admission control of an LSP reservation to a link.

RSVP Interface Level

The following are the configuration steps at the RSVP interface level:

1. The operator configures the percentage of RSVP interface bandwidth each CT shares, for example, the BC, using the **class-type-bw** command. The value entered at the interface level overrides the global value configured under the **diffserv-te** node.
 2. The operator can overbook (underbook) the maximum reservable bandwidth of a given CT by overbooking (underbooking) the interface maximum reservable bandwidth via configuring the appropriate value for the **subscription percentage** parameter in the **config>router>rsvp>interface** context.
 3. .Both the BC value and the subscription parameter can be changed at any time.
-

LSP and LSP Path Levels

The following are the configuration steps at the LSP and LSP path levels:

1. The operator configures the CT in which the LSP belongs by configuring the **class-type ct-number** command at the LSP level and/or the path level. The path level value overrides the LSP level value. By default, an LSP belongs to CT0.
2. Only one CT per LSP path is allowed per RFC 4124, *Protocol Extensions for Support of Diffserv-aware MPLS Traffic Engineering*. A multi-class LSP path is achieved through mapping multiple system Forwarding Classes to a CT.
3. The signaled CT of a dynamic bypass must always be CT0 regardless of the CT of the primary LSP path. The setup and hold priorities must be set to default values, for example, 7 and 0 respectively. This assumes that the operator configured a couple of TE classes, one which combines CT0 and a priority of 7 and the other which combines CT0 and a priority of 0. If not, the bypass LSP will not be signaled and will go into the down state.
4. The operator cannot configure the CT, setup priority, and holding priority of a manual bypass. They are always signaled with CT0 and the default setup and holding priorities.

5. The signaled CT, setup priority and holding priority of a detour LSP matches those of the primary LSP path it is associated with.
6. The operator can also configure the setup and holding priorities for each LSP path.
7. An LSP which does not have the CT explicitly configured will behave like a CT0 LSP when Diff-Serv is enabled.

If the operator configured a combination of a CT and a setup priority and/or a combination of a CT and a holding priority for an LSP path that are not supported by the user-defined TE classes, the LSP path will be kept in a down state and error code will be shown within the show command output for the LSP path.

Diff-Serv TE LSP Class Type Change under Failure

An option to configure a main Class Type (CT) and a backup CT for the primary path of a Diff-Serv TE LSP is provided. The main CT is used under normal operating conditions, for example, when the LSP is established the first time and when it gets re-optimized due to timer based or manual re-signal. The backup CT is used when the LSP retries under failure.

The use of backup Class Type (CT) by an LSP is enabled by executing the **config>router>mpls>lsp>primary>backup-class-type** *ct-number* command at the LSP primary path level.

When this option is enabled, the LSP will use the CT configured using the following commands (whichever is inherited at the primary path level) as the main CT:

- **config>router>mpls>lsp>class-type** *ct-number*
- **config>router>mpls>lsp>primary>class-type** *ct-number*

The main CT is used at initial establishment and during a manual or a timer based re-signal Make-Before-Break (MBB) of the LSP primary path. The backup CT is used temporarily to signal the LSP primary path when it fails and goes into retry.

Note that any valid values may be entered for the backup CT and main CT, but they cannot be the same. No check is performed to make sure that the backup CT is a lower CT in Diff-Serv Russian-Doll Model (RDM) admission control context.

The secondary paths of the same LSP are always signaled using the main CT as in existing implementation.

LSP Primary Path Retry Procedures

This feature behaves according to the following procedures.

- When a LSP primary path retries due a failure, for example, it fails after being in the up state, or undergoes any type of MBB, MPLS will retry a new path for the LSP using the main CT. If the first attempt failed, the head-end node performs subsequent retries using the backup CT. This procedure must be followed regardless if the currently used CT by this path is the main or backup CT. This applies to both CSPF and non-CSPF LSPs.
- The triggers for using the backup CT after the first retry attempt are:
 - A local interface failure or a control plane failure (hello timeout, etc.).
 - Receipt of a PathErr message with a notification of a FRR protection becoming active downstream and/or receipt of a Resv message with a 'Local-Protection-In-Use' flag set. This invokes the FRR Global Revertive MBB.

- Receipt of a PathErr message with error code=25 (Notify) and sub-code=7 (Local link maintenance required) or a sub-code=8 (Local node maintenance required). This invokes the TE Graceful Shutdown MBB. Note that in this case, only a single attempt is performed by MBB as in current implementation; only the main CT will be retried.
- Receipt of a Resv refresh message with the 'Preemption pending' flag set or a PathErr message with error code=34 (Reroute) and a value=1 (Reroute request soft preemption). This invokes the soft pre-emption MBB.
- Receipt of a ResvTear message.
- A configuration change MBB.
- When an unmapped LSP primary path goes into retry, it uses the main CT until the number of retries reaches the value of the new main-ct-retry-limit parameter. If the path did not come up, it must start using the backup CT at that point in time. By default, this parameter is set to infinite value. The new main-ct-retry-limit parameter has no effect on an LSP primary path, which retries due to a failure event. This parameter is configured using the **main-ct-retry-limit** command in the **config>router>mpls>lsp** context. If the user entered a value of the **main-ct-retry-limit** parameter that is greater than the LSP retry-limit, the number of retries will still stop when the LSP primary path reaches the value of the LSP retry-limit. In other words, the meaning of the LSP retry-limit parameter is not changed and always represents the upper bound on the number of retries. The unmapped LSP primary path behavior applies to both CSPF and non-CSPF LSPs.
- An unmapped LSP primary path is a path that never received a Resv in response to the first path message sent. This can occur when performing a "shut/no-shut" on the LSP or LSP primary path or when the node reboots. An unmapped LSP primary path goes into retry if the retry timer expired or the head-end node received a PathErr message before the retry timer expired.
- When the **clear>router>mpls>lsp** command is executed, the retry behavior for this LSP is the same as in the case of an unmapped LSP.
- If the value of the parameter main-ct-retry-limit is changed, the new value will only be used at the next time the LSP path is put into a "no-shut" state.
- The following is the behavior when the user changes the main or backup CT:
 - If the user changes the LSP level CT, all paths of the LSP are torn down and re-signaled in a break-before-make fashion. Specifically, the LSP primary path will be torn down and re-signaled even if it is currently using the backup CT.
 - If the user changes the main CT of the LSP primary path, the path will be torn down and re-signaled even if it is currently using the backup CT.
 - If the user changes the backup CT of an LSP primary path when the backup CT is in use, the path is torn down and is re-signaled.
 - If the user changes the backup CT of an LSP primary path when the backup CT is not in use, no action is taken. If however, the path was in global Revertive, gshut, or soft pre-emption MBB, the MBB is restarted. This actually means the first attempt will be with the main CT and subsequent ones, if any, with the new value of the backup CT.

- Consider the following priority of the various MBB types from highest to lowest: Delayed Retry, Preemption, Global Revertive, Configuration Change, and TE Graceful Shutdown. If an MBB request occurs while a higher priority MBB is in progress, the latter MBB will be restarted. This actually means the first attempt will be with the main CT and subsequent ones, if any, with the new value of the backup CT.
- If the least-fill option is enabled at the LSP level, then CSPF must use least-fill equal cost path selection when the main or backup CT is used on the primary path.
 - When the re-signal timer expires, CSPF will try to find a path with the main CT. The head-end node must re-signal the LSP even if the new path found by CSPF is identical to the existing one since the idea is to restore the main CT for the primary path. If a path with main CT is not found, the LSP remains on its current primary path using the backup CT. This means that the LSP primary path with the backup CT may no longer be the most optimal one. Furthermore, if the least-fill option was enabled at the LSP level, CSPF will not check if there is a more optimal path, with the backup CT, according to the least-fill criterion and will thus raise no trap to indicate the LSP path is eligible for least-fill re-optimization.
 - When the user performs a manual re-signal of the primary path, CSPF will try to find a path with the main CT. The head-end node must re-signal the LSP as in current implementation.
 - If a CPM switchover occurs while an the LSP primary path was in retry using the main or backup CT, for example, was still in operationally down state, the path retry will be restarted with the main CT until it comes up. This is because the LSP path retry count is not synchronized between the active and standby CPMs until the path becomes up.
 - When the user configured secondary standby and non-standby paths on the same LSP, the switchover behavior between primary and secondary is the same as in existing implementation.

This feature is not supported on a P2MP LSP.

Bandwidth Sharing Across Class Types

In order to allow different levels of booking of network links under normal operating conditions and under failure conditions, it is necessary to allow sharing of bandwidth across class types.

This feature introduces the Russian-Doll Model (RDM) Diff-Serv TE admission control policy described in RFC 4127, *Russian Dolls Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering*. This mode is enabled using the following command:

```
config>router>rsvp>diffserv-te rdm.
```

The Russian Doll Model (RDM) LSP admission control policy allows bandwidth sharing across Class Types (CTs). It provides a hierarchical model by which the reserved bandwidth of a CT is the sum of the reserved bandwidths of the numerically equal and higher CTs.

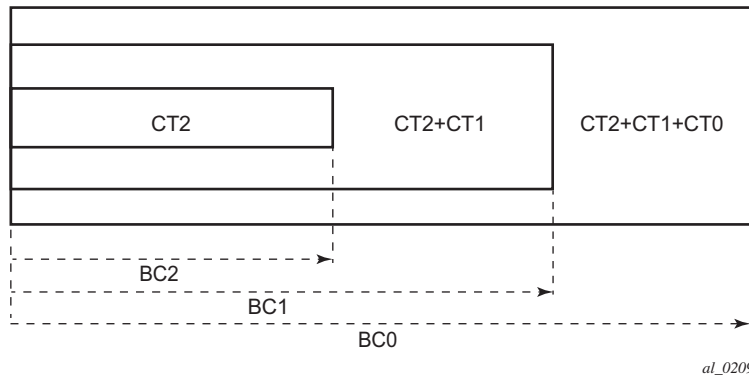


Figure 27: RDM Admission Control Policy Example


CT2 has a bandwidth constraint BC2 which represents a percentage of the maximum reservable link bandwidth. Both CT2 and CT1 can share BC1 which is the sum of the percentage of the maximum reservable bandwidth values configured for CT2 and CT1 respectively. Finally, CT2, CT1, and CT0 together can share BC0 which is the sum of the percentage of the maximum reservable bandwidth values configured for CT2, CT1, and CT0 respectively. The maximum value for BC0 is of course the maximum reservable link bandwidth.

What this means in practice is that CT0 LSPs can use up to BC0 in the absence of LSPs in CT1 and CT2. When this occurs and a CT2 LSP with a reservation less than or equal to BC2 requests admission, it is only admitted by preempting one or more CT0 LSPs of lower holding priority than this LSP setup priority. Otherwise, the reservation request for the CT2 LSP will be rejected.

It is required that multiple paths of the same LSP share common link bandwidth since they are signaled using the Shared Explicit (SE) style. Specifically, two instances of a primary path, one with the main CT and the other with the backup CT, must temporarily share bandwidth while MBB is in progress. Also, a primary path and one or many secondary paths of the same LSP must share bandwidth whether they are configured with the same or different CTs.

Downgrading the CT of Bandwidth Sharing LSP Paths

Consider a link configured with two class types CT0 and CT1 and making use of the RDM admission control model as shown in [Figure 28](#).



al_0210

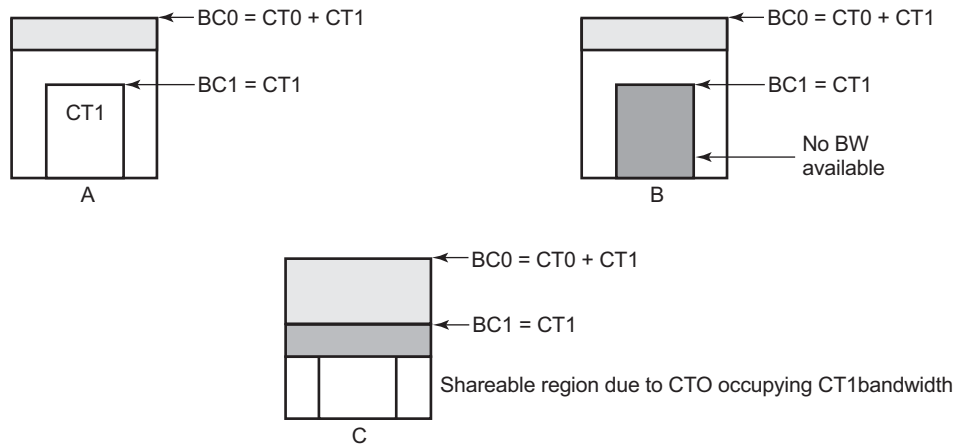
Figure 28: Sharing bandwidth when an LSP primary path is downgraded to backup CT

Consider an LSP path Z occupying bandwidth B at CT1. BC0 being the sum of all CTs below it, the bandwidth occupied in CT1 is guaranteed to be available in CT0. Thus when new path X of the same LSP for CT0 is setup, it will use the same bandwidth B as used by path Z as shown in [Figure 28](#) (a). When path Z is torn down the same bandwidth now occupies CT0 as shown in [Figure 28](#) (b). Even if there were no new BW available in CT0 as can be seen in [Figure 28](#) (c), path X can always share the bandwidth with path Z.

CSPF at the head-end node and CAC at the transit LSR node will share bandwidth of an existing path when its CT is downgraded in the new path of the same LSP.

Upgrading the CT of Bandwidth Sharing LSP Paths

When upgrading the CT the following issue can be apparent. Assume an LSP path X exists with CT0. An attempt is made to upgrade this path to a new path Z with CT1 using an MBB.



al_0211

Figure 29: Sharing Bandwidth When an LSP Primary Path is Upgraded to Main CT

In [Figure 29](#) (a), if the path X occupies the bandwidth as shown it can not share the bandwidth with the new path Z being setup. If a condition exists, as shown in [Figure 29](#), (b) the path Z can never be setup on this particular link.

Consider [Figure 29](#) (c). The CT0 has a region that overlaps with CT1 as CT0 has incursion into CT1. This overlap can be shared. However, in order to find whether such an incursion has occurred and how large the region is, it is required to know the reserved bandwidths in each class. Currently, IGP-TE advertises only the unreserved bandwidths. Hence, it is not possible to compute these overlap regions at the head end during CSPF. Moreover, the head end needs to then try and mimic each of the traversed links exactly which increases the complexity.

CSPF at the head-end node will only attempt to signal the LSP path with an upgraded CT if the advertised bandwidth for that CT can accommodate the bandwidth. In other words, it will assume that in the worst case this path will not share bandwidth with another path of the same LSP using a lower CT.

Advanced MPLS/RSVP Features

- [Extending RSVP LSP to use Loopback Interfaces Other Than router-id on page 125](#)
 - [LSP Path Change on page 125](#)
 - [RSVP-TE LSP Shortcut for IGP Resolution on page 132](#)
 - [Shared Risk Link Groups on page 140](#)
 - [TE Graceful Shutdown on page 144](#)
 - [Soft Pre-emption of Diff-Serv RSVP LSP on page 144](#)
 - [Least-Fill Bandwidth Rule in CSPF ECMP Selection on page 145](#)
-

Extending RSVP LSP to use Loopback Interfaces Other Than router-id

It is possible to configure the address of a loopback interface, other than the router-id, as the destination of an RSVP LSP, or a P2MP S2L sub-LSP. In the case of a CSPF LSP, CSPF searches for the best path that matches the constraints across all areas and levels of the IGP where this address is reachable. If the address is the router-id of the destination node, then CSPF selects the best path across all areas and levels of the IGP for that router-id; regardless of which area and level the router-id is reachable as an interface.

In addition, the user can now configure the address of a loopback interface, other than the router-id, as a hop in the LSP path hop definition. If the hop is strict and corresponds to the router-id of the node, the CSPF path can use any TE enabled link to the downstream node, based on best cost. If the hop is strict and does not correspond to the router-id of the node, then CSPF will fail.

LSP Path Change

The **tools perform router mpls update-path** `{lsp lsp-name path current-path-name new-path new-path-name}` command instructs MPLS to replace the path of the primary or secondary LSP.

The primary or secondary LSP path is indirectly identified via the current-path-name value. In existing implementation, the same path name cannot be used more than once in a given LSP name.

This command is also supported on an SNMP interface.

This command applies to both CSPF LSP and to a non-CSPF LSP. However, it will only be honored when the specified current-path-name has the adaptive option enabled. The adaptive option can be enabled the LSP level or at the path level.

The new path must be first configured in CLI or provided via SNMP. The **configure router mpls path** *path-name* CLI command is used to enter the path.

The command fails if any of the following conditions are satisfied:

- The specified current-path-name of this LSP does not have the adaptive option enabled.
- The specified new-path-name value does not correspond to a previously defined path.
- The specified new-path-name value exists but is being used by any path of the same LSP, including this one.

When the command is executed, MPLS performs the following procedures:

- MPLS performs a single MBB attempt to move the LSP path to the new path.
- If the MBB is successful, MPLS updates the new path.
 - MPLS writes the corresponding NHLFE in the data path if this path is the current backup path for the primary.
 - If the current path is the active LSP path, it will update the path, write the new NHLFE in the data path, which will cause traffic to switch to the new path.
- If the MBB is not successful, the path retains its current value.
- The update-path MBB has the same priority as the manual re-signal MBB.

Manual LSP Path Switch

This feature provides a new command to move the path of an LSP from a standby secondary to another standby secondary.

The base version of the command allows the path of the LSP to move from a standby (or an active secondary) to another standby of the same priority. If a new standby path with a higher priority or a primary path comes up after the **tools perform** command is executed, the path re-evaluation command runs and the path is moved to the path specified by the outcome of the re-evaluation.

The CLI command for the base version is:

```
tools perform router mpls switch-path lsp lsp-name path path-name
```

The sticky version of the command can be used to move from a standby path to any other standby path regardless of priority. The LSP remains in the specified path until this path goes down or the user performs the no form of the **tools perform** command.

The CLI commands for the sticky version are:

```
tools perform router mpls force-switch-path lsp lsp-name path path-name  
tools perform router mpls no force-switch-path lsp lsp-name
```

Make-Before-Break (MBB) Procedures for LSP/Path Parameter Configuration Change

When an LSP is switched from an existing working path to a new path, it is desirable to perform this in a hitless fashion. The Make-Before-Break (MBB) procedure consist of first signaling the new path when it is up, and having the ingress LER move the traffic to the new path. Only then the ingress LER tears down the original path.

MBB procedure is invoked during the following operations:

1. Timer based and manual re-signal of an LSP path.
2. Fast-ReRoute (FRR) global revertive procedures.
3. Soft Pre-emption of an LSP path.
4. Traffic-Engineering (TE) graceful shutdown procedures.
5. Update of secondary path due to an update to primary path SRLG.
6. LSP primary or secondary path name change.
7. LSP or path configuration parameter change.

In a prior implementation, item (7) covers the following parameters:

1. Changing the primary or secondary path **bandwidth** parameter on the fly.
2. Enabling the **frr** option for an LSP.

This feature extends the coverage of the MBB procedure to most of the other LSP level and Path level parameters as follows:

1. Changes to include/exclude of admin groups at LSP and path levels.
2. Enabling/disabling LSP level cspf option.
3. Enabling/disabling LSP level use-te-metric parameter when cspf option is enabled.
4. Enabling/disabling LSP level propagate-admin-group option.
5. Enabling/disabling LSP level hop-limit option in the fast-reroute context.
6. Enabling the LSP level least-fill option.
7. Enabling/disabling LSP level adspec option.
8. Changing between node-protect and “no node-protect” (link-protect) values in the LSP level fast-reroute option.
9. Changing LSP primary or secondary path priority values (setup-priority and hold-priority).

10. Changing LSP primary or secondary path class-type value and primary path backup-class-type value.
11. Changing LSP level and path level hop-limit parameter value.
12. Enabling/disabling primary or secondary path record and record-label options.

This feature is not supported on a manual bypass LSP.

P2MP Tree Level Make-before-break operation is supported if changes are made to the following parameters on LSP-Template:

1. Changing Bandwidth on P2MP LSP-Template.
2. Enabling Fast -Re-Route on P2MP LSP-Template.

Automatic Creation of RSVP-TE LSP Mesh

This feature enables the automatic creation of an RSVP point-to-point LSP to a destination node whose router-id matches a prefix in the specified peer prefix policy. This LSP type is referred to as auto-LSP of type mesh.

The user can associate multiple templates with the same or different peer prefix policies. Each application of an LSP template with a given prefix in the prefix list will result in the instantiation of a single CSPF computed LSP primary path using the LSP template parameters as long as the prefix corresponds to a router-id for a node in the TE database. Each instantiated LSP will have a unique LSP-id and a unique tunnel-ID.

Up to five (5) peer prefix policies can be associated with a given LSP template at all times. Each time the user executes the above command with the same or different prefix policy associations, or the user changes a prefix policy associated with an LSP template, the system re-evaluates the prefix policy. The outcome of the re-evaluation will tell MPLS if an existing LSP needs to be torn down or if a new LSP needs to be signaled to a destination address that is already in the TE database.

If a /32 prefix is added to (removed from) or if a prefix range is expanded (shrunk) in a prefix list associated with a LSP template, the same prefix policy re-evaluation described above is performed.

The trigger to signal the LSP is when the router with a router-id the matching a prefix in the prefix list appears in the Traffic Engineering database. The signaled LSP is installed in the Tunnel Table Manager (TTM) and is available to applications such as LDP-over-RSVP, resolution of BGP label routes, resolution of BGP, IGP, and static routes. It can also be used for provisioning an SDP; it is however, not available to be used as a provisioned SDP for explicit binding or auto-binding by services.

If the **one-hop** option is specified instead of a prefix policy, this command enables the automatic signaling of one-hop point-to-point LSPs using the specified template to all directly connected neighbors. This LSP type is referred to as auto-LSP of type one-hop. Although the provisioning model and CLI syntax differ from that of a mesh LSP only by the absence of a prefix list, the actual behavior is quite different. When the above command is executed, the TE database will keep track of each TE link that comes up to a directly connected IGP neighbor whose router-id is discovered. It then instructs MPLS to signal an LSP with a destination address matching the router-id of the neighbor and with a strict hop consisting of the address of the interface used by the TE link. Thus, the **auto-lsp** command with the **one-hop** option will result in one or more LSPs signaled to the neighboring router.

An auto-created mesh or one-hop LSP can have egress statistics collected at the ingress LER by adding the **egress-statistics** node configuration into the LSP template. The user can also have ingress statistics collected at the egress LER using the same **ingress-statistics** node in CLI used

with a provisioned LSP. The user must specify the full LSP name as signaled by the ingress LER in the RSVP session name field of the Session Attribute object in the received Path message.

RSVP-TE LSP Shortcut for IGP Resolution

RSVP-TE LSP shortcut for IGP route resolution allows forwarding of packets to IGP learned routes using an RSVP-TE LSP. This is also referred to as IGP shortcut. This feature is enabled by entering the following command at the IS-IS routing protocol level or at the OSPF routing protocol instance level:

- **config>router>isis>rsvp-shortcut**
- **config>router>ospf>rsvp-shortcut**

These commands instruct IS-IS or OSPF to include RSVP LSPs originating on this node and terminating on the router-id of a remote node as direct links with a metric equal to the operational metric provided by MPLS. If the user enabled the relative-metric option for this LSP, IGP will apply the shortest IGP cost between the endpoints of the LSP plus the value of the offset, instead of the LSP operational metric, when computing the cost of a prefix which is resolved to the LSP.

When a prefix is resolved to a tunnel next-hop, the packet is sent labeled with the label stack corresponding to the NHLFE of the RSVP LSP. Any network event causing an RSVP LSP to go down will trigger a full SPF computation which may result in installing a new route over another RSVP LSP shortcut as tunnel next-hop or over a regular IP next-hop.

When rsvp-shortcut is enabled at the IGP instance level, all RSVP LSPs originating on this node are eligible by default as long as the destination address of the LSP, as configured in **config-ure>router>mpls>lsp>to**, corresponds to a router-id of a remote node. RSVP LSPs with a destination corresponding to an interface address or any other loopback interface address of a remote node are automatically not considered by IS-IS or OSPF. The user can, however, exclude a specific RSVP LSP from being used as a shortcut for resolving IGP routes by entering the command:

- **config>router>mpls>lsp>no igp-shortcut**

The SPF in OSPF or IS-IS will only use RSVP LSPs as forwarding adjacencies, IGP shortcuts, or as endpoints for LDP-over-RSVP. These applications of RSVP LSPs are mutually exclusive at the IGP instance level. If the user enabled two or more options in the same IGP instance, then forwarding adjacency takes precedence over the shortcut application that takes precedence over the LDP-over-RSVP application.

[Table 7](#) summarizes the outcome in terms of RSVP LSP role of mixing these configuration options.

Table 7: RSVP LSP Role As Outcome of LSP level and IGP level configuration options

LSP level configuration	IGP Instance level configurations					
	advertise-tunnel-link enabled / rsvp-shortcut enabled / ldp-over-rsvp enabled	advertise-tunnel-link enabled / rsvp-shortcut enabled / ldp-over-rsvp disabled	advertise-tunnel-link enabled / rsvp-shortcut disabled / ldp-over-rsvp disabled	advertise-tunnel-link disabled / rsvp-shortcut disabled / ldp-over-rsvp disabled	advertise-tunnel-link disabled / rsvp-shortcut enabled / ldp-over-rsvp enabled	advertise-tunnel-link disabled / rsvp-shortcut disabled / ldp-over-rsvp enabled
igp-shortcut enabled / ldp-over-rsvp enabled	Forwarding Adjacency	Forwarding Adjacency	Forwarding Adjacency	None	IGP Shortcut	LDP-over-RSVP
igp-shortcut enabled / ldp-over-rsvp disabled	Forwarding Adjacency	Forwarding Adjacency	Forwarding Adjacency	None	IGP Shortcut	None
igp-shortcut disabled / ldp-over-rsvp enabled	None	None	None	None	None	LDP-over-RSVP
igp-shortcut disabled / ldp-over-rsvp disabled	None	None	None	None	None	None

The resolution and forwarding of IPv6 prefixes to IPv4 IGP shortcuts is not supported.

The **no** form of this command disables the resolution of IGP routes using RSVP shortcuts.

Using LSP Relative Metric with IGP Shortcut

By default, the absolute metric of the LSP is used to update the SPF tree when the user enables IGP shortcut by configuring the `rsvp-shortcut` option in IGP. The absolute metric is the operational metric of the LSP populated by MPLS in the Tunnel Table Manager (TTM). This corresponds to the cumulative IGP-metric of the LSP path returned by CSPF or the static admin metric value of the LSP if the user configured one using the `config>router>mpls>lsp>metric` command. Note that MPLS populates the TTM with the maximum metric value of 16777215 in the case of a CSPF LSP using the TE-metric and a non-CSPF LSP with a loose or strict hop in the path. A non-CSPF LSP with an empty hop in the path definition returns the IGP cost for the destination of the LSP.

The user enables the use of the relative metric for an IGP shortcut with the following new CLI command:

```
config>router>mpls>lsp>igp-shortcut relative-metric [offset]
```

IGP will apply the shortest IGP cost between the endpoints of the LSP plus the value of the offset, instead of the LSP operational metric, when computing the cost of a prefix which is resolved to the LSP.

The offset value is optional and it defaults to zero. An offset value of zero is used when the **relative-metric** option is enabled without specifying the offset parameter value.

The minimum net cost for a prefix is capped to the value of one (1) after applying the offset:

Prefix cost = max(1, IGP cost + relative metric offset)

Note that the TTM continues to show the LSP operational metric as provided by MPLS. In other words, applications such as LDP-over-RSVP (when IGP shortcut is disabled) and BGP and static route shortcuts will continue to use the LSP operational metric.

The **relative-metric** option is mutually exclusive with the **lfa-protect** or the **lfa-only** options. In other words, an LSP with the **relative-metric** option enabled cannot be included in the LFA SPF and vice-versa when the **rsvp-shortcut** option is enabled in the IGP.

Finally, it should be noted that the **relative-metric** option is ignored when forwarding adjacency is enabled in IS-IS or OSPF by configuring the **advertise-tunnel-link** option. In this case, IGP advertises the LSP as a point-to-point unnumbered link along with the LSP operational metric capped to the maximum link metric allowed in that IGP.

The resolution and forwarding of IPv6 prefixes to IPv4 forwarding adjacency LSP is not supported.

ECMP Considerations

When ECMP is enabled on the system and multiple equal-cost paths exist for a prefix, the following selection criteria are used to pick up the set of next-hops to program in the data path:

- for a destination = tunnel-endpoint (including external prefixes with tunnel-endpoint as the next-hop):
 - select tunnel with lowest tunnel-index (ip next-hop is never used in this case)
- for a destination != tunnel-endpoint:
 - exclude LSPs with metric higher than underlying IGP cost between the endpoint of the LSP
 - prefer tunnel next-hop over ip next-hop

- within tunnel next-hops:
 - i. select lowest endpoint to destination cost
 - ii. if same endpoint to destination cost, select lowest endpoint node router-id
 - iii. if same router-id, select lowest tunnel-index
- within ip next-hops:
 - i. select lowest downstream router-id
 - ii. if same downstream router-id, select lowest interface-index
- Note though no ECMP is performed across both the IP and tunnel next-hops the tunnel end-point lies in one of the shortest IGP paths for that prefix. In that case, the tunnel next-hop is always selected as long as the prefix cost using the tunnel is equal or lower than the IGP cost.

The ingress IOM will spray the packets for a prefix over the set of tunnel next-hops and IP next-hops based on the hashing routine currently supported for IPv4 packets.

Handling of Control Packets

All control plane packets that require an RTM lookup and whose destination is reachable over the RSVP shortcut will be forwarded over the shortcut. This is because RTM keeps a single route entry for each prefix unless there is ECMP over different outgoing interfaces.

Interface bound control packets are not impacted by the RSVP shortcut since RSVP LSPs with a destination address different than the router-id are not included by IGP in its SPF calculation.

Forwarding Adjacency

The forwarding adjacency feature can be enabled independently from the IGP shortcut feature in CLI. To enable forwarding adjacency, the user enters the following command in IS-IS or OSPF:

- **configure>router>isis>advertise-tunnel-link**
- **configure>router>ospf>advertise-tunnel-link**

If both **rsvp-shortcut** and **advertise-tunnel-link** options are enabled for a given IGP instance, then the **advertise-tunnel-link** will win. With this feature, ISIS or OSPF advertises an RSVP LSP as a link so that other routers in the network can include it in their SPF computations. The RSVP LSP is advertised as an unnumbered point-to-point link and the link LSP/LSA has no Traffic Engineering opaque sub-TLVs as per RFC 3906 *Calculating Interior Gateway Protocol (IGP) Routes Over Traffic Engineering Tunnels*.

The forwarding adjacency feature can be enabled independently from the IGP shortcut feature in CLI. If both **rsvp-shortcut** and **advertise-tunnel-link** options are enabled for a given IGP instance, then the **advertise-tunnel-link** will win.

When the forwarding adjacency feature is enabled, each node advertises a p2p unnumbered link for each best metric tunnel to the router-id of any endpoint node. The node does not include the tunnels as IGP shortcuts in SPF computation directly. Instead, when the LSA/LSP advertising the corresponding P2P unnumbered link is installed in the local routing database, then the node performs an SPF using it like any other link LSA/LSP. The link bi-directional check requires that a link, regular link or tunnel link, exists in the reverse direction for the tunnel to be used in SPF.

Note that the **igp-shortcut** option under the LSP name governs the use of the LSP with both the **rsvp-shortcut** and the **advertise-tunnel-link** options in IGP. The interactions of these options are summarized in [Table 8](#):

Table 8: Impact of LSP level configuration on IGP shortcut and forwarding adjacency features

LSP level configuration	Actions with IGP Shortcut Feature	Actions with Forwarding Adjacency Feature
igp-shortcut	Tunnel is used in main SPF, but is not used in LFA SPF	Tunnel is advertised as p2p link if it has best LSP metric, is used in main SPF if advertised, but is not used in LFA SPF
igp-shortcut lfa-protect	Tunnel is used in main SPF, and is used in LFA SPF	Tunnel is advertised as p2p link if it has best LSP metric, is used in main SPF if advertised, and is used in LFA SPF regardless if it is advertised or not
igp-shortcut lfa-only	Tunnel is not used in main SPF, but is used in LFA SPF	Tunnel is not advertised as p2p link, if not used in main SPF, but is used in LFA SPF

LDP Forwarding over IGP Shortcut

The user can enable LDP FECs over IGP shortcuts by configuring T-LDP sessions to the destination of the RSVP LSP. In this case, LDP FEC is tunneled over the RSVP LSP, effectively implementing LDP-over-RSVP without having to enable the **ldp-over-rsvp** option in OSPF or IS-IS. The **ldp-over-rsvp** and **igp-shortcut** options are mutually exclusive under OSPF or IS-IS.

Handling of Multicast Packets

This feature supports multicast Reverse-Path Check (RPF) in the presence of IGP shortcuts. When the multicast source for a packet is reachable via an IGP shortcut, the RPF check fails since PIM requires a bi-directional path to the source but IGP shortcuts are unidirectional.

The implementation of the IGP shortcut feature provides IGP with the capability to populate the multicast RTM with the prefix IP next-hop when both the **rsvp-shortcut** option and the **multicast-import** option are enabled in IGP.

This change is made possible with the enhancement introduced by which SPF keeps track of both the direct first hop and the tunneled first hop of a node that is added to the Dijkstra tree.

Note that IGP will not pass LFA next-hop information to the mcast RTM in this case. Only ECMP next-hops are passed. As a consequence, features such as PIM Multicast-Only FRR (MoFRR) will only work with ECMP next-hops when IGP shortcuts are enabled.

Finally, note that the concurrent enabling of the **advertise-tunnel-link** option and the **multicast-import** option will result a multicast RTM that is a copy of the unicast RTM and is thus populated with mix of IP and tunnel NHs. RPF will succeed for a prefix resolved to a IP NH, but will fail for a prefix resolved to a tunnel NH. [Table 9](#) summarizes the interaction of the **rsvp-shortcut** and **advertise-tunnel-link** options with unicast and multicast RTMs.

Table 9: Impact of IGP Shortcut and Forwarding Adjacency on Unicast and Multicast RTM

		Unicast RTM (Primary SPF)	Multicast RTM (Primary SPF)	Unicast RTM (LFA SPF)	Multicast RTM (LFA SPF)
OSPF	rsvp-shortcut	√	√ (1)	√	X (3)
	advertise-tunnel-link	√	√ (2)	√	√ (4)
IS-IS	rsvp-shortcut	√	√ (1)	√	X (3)
	advertise-tunnel-link	√	√ (2)	√	√ (4)

Notes:

1. Multicast RTM is different from unicast RTM as it is populated with IP NHs only, including ECMP IP NHs. RPF check can be performed for all prefixes.
2. Multicast RTM is a copy of the unicast RTM and is thus populated with mix of IP and tunnel NHs. RPF will succeed for a prefix resolved to a IP NH but will fail for a prefix resolved to a tunnel NH.

3. LFA NH is not computed for the IP primary next-hop of a prefix passed to multicast RTM even if the same IP primary next-hop ends up being installed in the unicast RTM. The LFA next-hop will, however, be computed and installed in the unicast RTM for a primary IP next-hop of a prefix.
 4. Multicast RTM is a copy of the unicast RTM and is thus populated with mix of IP and tunnel LFA NHs. RPF will succeed for a prefix resolved to a primary or LFA IP NH but will fail for a prefix resolved to a primary or LFA tunnel NH.
-

Disabling TTL Propagation in an LSP Shortcut

This feature provides the option for disabling TTL propagation from a transit or a locally generated IP packet header into the LSP label stack when an RSVP LSP is used as a shortcut for BGP next-hop resolution, a static-route next-hop resolution, or for an IGP route resolution.

A transit packet is a packet received from an IP interface and forwarded over the LSP shortcut at ingress LER.

A locally-generated IP packet is any control plane packet generated from the CPM and forwarded over the LSP shortcut at ingress LER.

TTL handling can be configured for all RSVP LSP shortcuts originating on an ingress LER using the following global commands:

```
config>router>mpls>[no] shortcut-transit-ttl-propagate  
config>router>mpls>[no] shortcut-local-ttl-propagate
```

These commands apply to all RSVP LSPs which are used to resolve static routes, BGP routes, and IGP routes.

When the **no** form of the above command is enabled for local packets, TTL propagation is disabled on all locally generated IP packets, including ICMP Ping, trace route, and OAM packets that are destined to a route that is resolved to the LSP shortcut. In this case, a TTL of 255 is programmed onto the pushed label stack. This is referred to as pipe mode.

Similarly, when the **no** form is enabled for transit packets, TTL propagation is disabled on all IP packets received on any IES interface and destined to a route that is resolved to the LSP shortcut. In this case, a TTL of 255 is programmed onto the pushed label stack.

RSVP-TE LSP Signaling using LSP Template

LSP template can be used for signaling RSVP-TE LSP to far-end PE node that is detected based on auto-discovery method by a client application. RSVP-TE P2MP LSP signaling based on LSP template is supported for Multicast VPN application on SROS platform. LSP template avoids an explicit LSP or LSP S2L configuration for a node that is dynamically added as a receiver.

LSP template has option to configure traffic engineering parameters that apply to LSP that is setup using the template. Traffic engineering options that are currently supported are:

- adaptive
- admin-group
- bandwidth
- CSPF calculation
- fast-reroute
- hop-limit
- record-label
- retry-timer

Shared Risk Link Groups

Shared Risk Link Groups (SRLGs) is a feature that allows the user to establish a backup secondary LSP path or a FRR LSP path which is disjoint from the path of the primary LSP. Links that are members of the same SRLG represent resources sharing the same risk, for example, fiber links sharing the same conduit or multiple wavelengths sharing the same fiber.

When the SRLG option is enabled on a secondary path, CSPF includes the SRLG constraint in the computation of the secondary LSP path. This requires that the primary LSP already be established and up since the head-end LER needs the most current ERO computed by CSPF for the primary path. CSPF would return the list of SRLG groups along with the ERO during primary path CSPF computation. At a subsequent establishment of a secondary path with the SRLG constraint, the MPLS/RSVP task will query again CSPF providing the list of SLRG group numbers to be avoided. CSPF prunes all links with interfaces which belong to the same SRLGs as the interfaces included in the ERO of the primary path. If CSPF finds a path, the secondary is setup. If not, MPLS/RSVP will keep retrying the requests to CSPF.

When the SRLG option is enabled on FRR, CSPF includes the SRLG constraint in the computation of a FRR detour or bypass for protecting the primary LSP path. CSPF prunes all links with interfaces which belong to the same SRLG as the interface which is being protected, for example, the outgoing interface at the PLR the primary path is using. If one or more paths are found, the MPLS/RSVP task will select one based on best cost and will signal the bypass/detour. If not and the user included the strict option, the bypass/detour is not setup and the MPLS/RSVP task will keep retrying the request to CSPF. Otherwise, if a path exists which meets the other TE constraints, other than the SRLG one, the bypass/detour is setup.

A bypass or a detour LSP path is not guaranteed to be SRLG disjoint from the primary path. This is because only the SRLG constraint of the outgoing interface at the PLR that the primary path is using is avoided.

Enabling Disjoint Backup Paths


A typical application of the SRLG feature is to provide for an automatic placement of secondary backup LSPs or FRR bypass/detour LSPs that minimizes the probability of fate sharing with the path of the primary LSP ([Figure 30](#)).

The following details the steps necessary to create shared risk link groups:

- For primary/standby SRLG disjoint configuration:
 - Create an SRLG-group, similar to admin groups.
 - Link the SRLG-group to MPLS interfaces.

- Configure primary and secondary LSP paths and enable SRLG on the secondary LSP path. Note that the SRLG secondary LSP path(s) will *always* perform a strict CSPF query. The **srlg-frr** command is irrelevant in this case (see [srlg-frr on page 227](#)).
- For FRR detours/bypass SRLG disjoint configuration:
 - Create an SRLG group, similar to admin groups.
 - Link the SRLG group to MPLS interfaces.
 - Enable the **srlg-frr** (strict/non-strict) option, which is a system-wide parameter, and it force every LSP path CSPF calculation, to take the configured SRLG membership(s) (and propagated through the IGP opaque-te-database) into account.
 - Configure primary FRR (one-to-one/facility) LSP path(s). Consider that each PLR will create a detour/bypass that will only avoid the SRLG membership(s) configured on the primary LSP path egress interface. In a one-to-one case, detour-detour merging is out of the control of the PLR, thus the latter will not ensure that its detour will be prohibited to merge with a colliding one. For facility bypass, with the presence of several bypass type to bind to, the following priority rules will be followed:
 1. Manual bypass disjoint
 2. Manual bypass non-disjoint (eligible only if srlg-frr is non-strict)
 3. Dynamic disjoint
 4. Dynamic non-disjoint (eligible only if srlg-frr is non-strict)

Non-CSPF manual bypass is not considered.



Fig_33

Figure 30: Shared Risk Link Groups

This feature is supported on OSPF and IS-IS interfaces on which RSVP is enabled.

Static Configurations of SRLG Memberships

This feature provides operations with the ability to manually enter the link members of SRLG groups for the entire network at any SR OS node which will need to signal LSP paths (for example, a head-end node).

The operator may explicitly enable the use by CSPF of the SRLG database. In that case, CSPF will not query the TE database for IGP advertised interface SRLG information.

Note, however, that the SRLG secondary path computation and FRR bypass/detour path computation remains unchanged.

There are deployments where the SR OS will interoperate with routers that do not implement the SRLG membership advertisement via IGP SRLG TLV or sub-TLV.

In these situations, the user is provided with the ability to enter manually the link members of SRLG groups for the entire network at any SR OS node which will need to signal LSP paths, for example, a head-end node.

The user enters the SRLG membership information for any link in the network by using the **interface** *ip-int-name* **srlg-group** *group-name* command in the **config>router>mpls>srlg-database>router-id** context. An interface can be associated with up to 5 SRLG groups for each execution of this command. The user can associate an interface with up to 64 SRLG groups by executing the command multiple times. The user must also use this command to enter the local interface SRLG membership into the user SRLG database. The user deletes a specific interface entry in this database by executing the **no** form of this command.

The *group-name* must have been previously defined in the **srlg-group** *group-name* **value** *group-value* command in the **config>router>mpls if-attribute**. The maximum number of distinct SRLG groups the user can configure on the system is 1024.

The parameter value for *router-id* must correspond to the router ID configured under the base router instance, the base OSPF instance or the base IS-IS instance of a given node. Note however that a single user SRLG database is maintained per node regardless if the listed interfaces participate in static routing, OSPF, IS-IS, or both routing protocols. The user can temporarily disable the use by CSPF of all interface membership information of a specific router ID by executing the **shutdown** command in the **config>router>mpls>srlg-database>router-id** context. In this case, CSPF will assume these interfaces have no SRLG membership association. The operator can delete all interface entries of a specific router ID entry in this database by executing the **no router-id** *router-address* command in the **config>router>mpls>srlg-database** context.

CSPF will not use entered SRLG membership if an interface is not listed as part of a router ID in the TE database. If an interface was not entered into the user SRLG database, it will be assumed that it does not have any SRLG membership. CSPF will not query the TE database for IGP advertised interface SRLG information.

The operator enables the use by CSPF of the user SRLG database by entering the `user-srlg-db enable` command in the **config>router>mpls** context. When the MPLS module makes a request to CSPF for the computation of an SRLG secondary path, CSPF will query the local SRLG and computes a path after pruning links which are members of the SRLG IDs of the associated primary path. Similarly, when MPLS makes a request to CSPF for a FRR bypass or detour path to associate with the primary path, CSPF queries the user SRLG database and computes a path after pruning links which are members of the SRLG IDs of the PLR outgoing interface.

The operator can disable the use of the user SRLG database by entering the `user-srlg-db disable` in command in the **config>router>mpls** context. CSPF will then resumes queries into the TE database for SRLG membership information. However, the user SRLG database is maintained

The operator can delete the entire SRLG database by entering the **no srlg-database** command in the **config>router>mpls** context. In this case, CSPF will assume all interfaces have no SRLG membership association if the user has not disabled the use of this database.

TE Graceful Shutdown

Graceful shutdown provides a method to bulk re-route transit LSPs away from the node during software upgrade of a node. A solution is described in RFC 5817, *Graceful Shutdown in MPLS and Generalized MPLS Traffic Engineering Networks*. This is achieved in this RFC by using a PathErr message with a specific error code Local Maintenance on TE link required flag. When a LER gets this message, it performs a make-before-break on the LSP path to move the LSP away from the links/nodes which IP addresses were indicated in the PathErr message.

Graceful shutdown can flag the affected link/node resources in the TE database so other routers will signal LSPs using the affected resources only as a last resort. This is achieved by flooding an IGP TE LSA/LSP containing link TLV for the links under graceful shutdown with the traffic engineering metric set to 0xffffffff and 0 as unreserved bandwidth.

Soft Pre-emption of Diff-Serv RSVP LSP

A Diff-Serv LSP can pre-empt another LSP of the same or of a different CT if its setup priority is strictly higher (numerically lower) than the holding priority of that other LSP.

Least-Fill Bandwidth Rule in CSPF ECMP Selection

When multiples equal-cost paths satisfy the constraints of a given RSVP LSP path, CSPF in the router head-end node will select a path so that LSP bandwidth is balanced across the network links. In releases prior to R7.0, CSPF used a random number generator to select the path and returned it to MPLS. In the course of time, this method actually balances the number of LSP paths over the links in the network; it does not necessarily balance the bandwidth across those links.

The least-fill path selection algorithm identifies the single link in each of the equal cost paths which has the least available bandwidth in proportion to its maximum reserved bandwidth. It then selects the path which has the largest value of this figure. The net affect of this algorithm is that LSP paths will be spread over the network links over time such that percentage link utilization is balanced. When the least-fill option is enabled on an LSP, during a manual reset CSPF will apply this method to all path calculations of the LSP, also at the time of the initial configuration.

Inter Area TE LSP (ERO Expansion Method)

Inter area contiguous LSP scheme provides end-to-end TE path. Each transit node in an area can set up a TE path LSP based on TE information available within its local area.

A PE node initiating an inter area contiguous TE LSP does partial CSPF calculation to include its local area border router as a loose node.

Area border router on receiving a PATH message with loose hop ERO does a partial CSPF calculation to the next domain border router as loose hop or CSPF to reach the final destination.

Area Border Node FRR Protection for Inter Area LSP

This feature enhances the prior implementation of an inter-area RSVP P2P LSP by making the ABR selection automatic at the ingress LER. The user will not need to include the ABR as a loose-hop in the LSP path definition.

CSPF adds a new capability to compute all segments of a multi-segment intra-area or inter-area LSP path in one operation. In previous releases, MPLS makes a request to CSPF for each segment separately.

Figure 7 1 illustrates the role of each node in the signaling of an inter-area LSP with automatic ABR node selection.

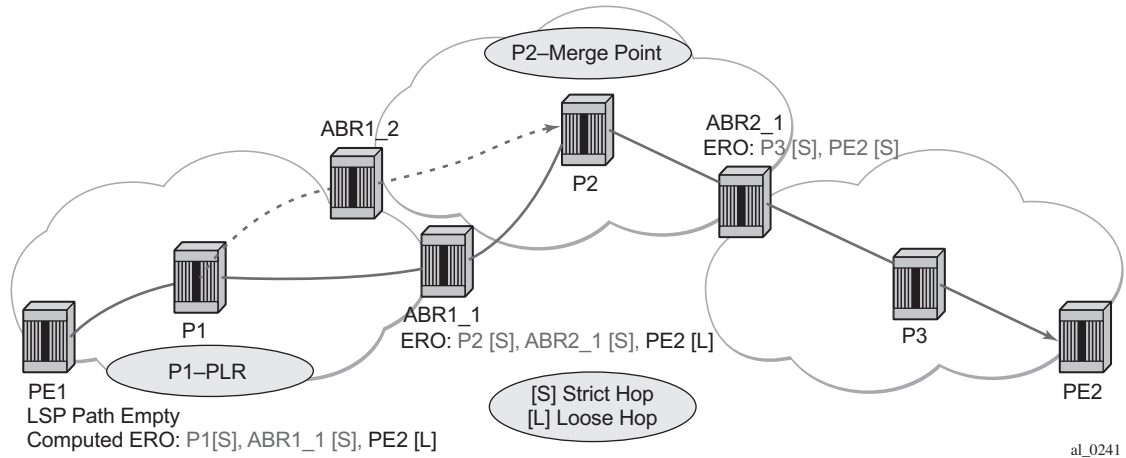



Figure 31: Automatic ABR Node Selection for Inter-Area LSP

CSPF for an inter-area LSP operates as follows:

1. CSPF in the Ingress LER node determines that an LSP is inter-area by doing a route lookup with the destination address of a P2P LSP (i.e., the address in the to field of the LSP configuration). If there is no intra-area route to the destination address, the LSP is considered as inter-area.
2. When the path of the LSP is empty, CPSF will compute a single-segment intra-area path to an ABR node that advertised a prefix matching with the destination address of the LSP.
3. When the path of the LSP contains one or more hops, CSPF will compute a multi-segment intra-area path including the hops that are in the area of the Ingress LER node.
4. When all hops are in the area of the ingress LER node, the calculated path ends on an ABR node that advertised a prefix matching with the destination address of the LSP.
5. When there are one or more hops that are not in the area of the ingress LER node, the calculated path ends on an ABR node that advertised a prefix matching with the first hop-address that is not in the area of the ingress LER node.
6. Note the following special case of a multi-segment inter-area LSP. If CSPF hits a hop that can be reached via an intra-area path but that resides on an ABR, CSPF only calculates a path up to that ABR. This is because there is a better chance to reach the destination of the LSP by first signaling the LSP up to that ABR and continuing the path calculation from there on by having the ABR expand the remaining hops in the ERO.

This behavior can be illustrated in the following example. The TE link between ABR nodes D and E is in area 0. When node C computes the path for LSP from C to B which path specified nodes C and D as loose hops, it would fail the path computation if CSPF attempted a path all the way to the last hop in the local area, node E. Instead, CSPF stops

the path at node A which will further expand the ERO by including link D-E as part of the path in area 0.



7. If there is more than 1 ABR that advertised a prefix, CSPF will calculate a path for all ABRs. Only the shortest path will be withheld. If more than one path has the shortest path, CSPF will pick a path randomly or based on the least-fill criterion if enabled. If more than one ABR satisfies the least-fill criterion, CSPF will also pick one path randomly.
8. The path for an intra-area LSP path will not be able to exit and re-enter the local area of the ingress LER. This behavior was possible in prior implementation when the user specified a loose hop outside of the local area or when the only available path was via TE links outside of the local area.

Rerouting of Inter-Area LSP

In prior implementation, an inter-area LSP path would have been re-routed if a failure or a topology change occurred in the local or a remote area while the ABR loose-hop in the path definition was still up. If the exit ABR node went down, went into IS-IS overload, or was put into node TE graceful shutdown, the LSP path will remain down at the ingress LER.

One new behavior introduced by the automatic selection of ABR is the ability of the ingress LER to reroute an inter-area LSP primary path via a different ABR in the following situations:

- When the local exit ABR node fails, There are two cases to consider:
 - The primary path is not protected at the ABR and is thus torn down by the previous hop in the path. In this case the ingress LER will retry the LSP primary path via the ABR which currently has the best path for the destination prefix of the LSP.
 - The primary path is protected at the ABR with a manual or dynamic bypass LSP. In this case the ingress LER will receive a Path Error message with a notification of a protection becoming active downstream and a RESV with a *Local-Protection-In-Use* flag set. At the receipt of first of these two messages, the ingress LER will then

perform a Global Revertive Make-Before-Break (MBB) to re-optimize the LSP primary path via the ABR which currently has the best path for the destination prefix of the LSP.

- When the local exit ABR node goes into IS-IS overload or is put into node TE Graceful Shutdown. In this case, the ingress LER will perform a MBB to re-optimize the LSP primary path via the ABR which currently has the best path for the destination prefix of the LSP. The MBB is performed at the receipt of the PathErr message for the node TE shutdown or at the next timer or manual re-optimization of the LSP path in the case of the receipt of the IS-IS overload bit.

Behavior of MPLS Options in Inter-Area LSP

The automatic ABR selection for an inter-area LSP does not change prior implementation inter-area LSP behavior of many of the LSP and path level options. There is, however, a number of enhancements introduced by the automatic ABR selection feature as explained in the following.

- Features such as path bandwidth reservation and admin-groups continue to operate within the scope of all areas since they rely on propagating the parameter information in the Path message across the area boundary.
- The TE graceful shutdown and soft pre-emption features will continue to support MBB of the LSP path to avoid the link or node that originated the PathErr message as long as the link or node is in the local area of the ingress LER. If the PathErr originated in a remote area, the ingress LER will not be able to avoid the link or node when it performs the MBB since it computes the path to the local ABR exit router only. There is, however, an exception to this for the TE graceful shutdown case only. An enhancement has been added to cause the upstream ABR nodes in the current path of the LSP to record the link or node to avoid and will use it in subsequent ERO expansions. This means that if the ingress LER computes a new MBB path which goes via the same exit ABR router as the current path and all ABR upstream nodes of the node or link which originated the PathErr message are also selected in the new MBB path when the ERO is expanded, the new path will indeed avoid this link or node. The latter is a new behavior introduced with the automatic ABR selection feature.
- The support of MBB to avoid the ABR node when the node is put into TE Graceful Shutdown is a new behavior introduced with the automatic ABR selection feature.
- The **use-te-metric** option in CSPF cannot be propagated across the area boundary and thus will operate within the scope of the local area of the ingress LER node. This is a new behavior introduced with the automatic ABR selection feature.
- The **srlg** option on bypass LSP will continue to operate locally at each PLR within each area. The PLR node protecting the ABR will check the SRLG constraint for the path of the bypass within the local area.
- The **srlg** option on secondary path is allowed to operate within the scope of the local area of the ingress LER node with the automatic ABR selection feature.

- The **least-fill** option support with an inter-area LSP is introduced with the automatic ABR selection feature. When this option is enabled, CSPF applies the least-fill criterion to select the path segment to the exit ABR node in the local area.
 - The PLR node must indicate to CSPF that a request to one-to-one detour LSP path must remain within the local area. If the destination for the detour, which is the same as that of the LSP, is outside of the area, CSPF must return no path.
 - The **propagate-admin-group** option under the LSP will still need to be enabled on the inter-area LSP if the user wants to have admin-groups propagated across the areas.
 - With the automatic ABR selection feature, timer based re-signal of the inter-area LSP path will be supported and will re-signal the path if the cost of the path segment to the local exit ABR changed. The cost shown for the inter-area LSP at ingress LER will be the cost of the path segments to the ABR node.
-

Inter-Area LSP support of OSPF Virtual Links

The OSPF virtual link extends area 0 for a router that is not connected to area 0. As a result, it makes all prefixes in area 0 reachable via an intra-area path but in reality, they are not since the path crosses the transit area through which the virtual link is set up to reach the area 0 remote nodes.

The TE database in a router learns all of the remote TE links in area 0 from the ABR connected to the transit area, but an intra-area LSP path using these TE links cannot be signaled within area 0 since none of these links is directly connected to this node.

This inter-area LSP feature can identify when the destination of an LSP is reachable via a virtual link. In that case, CSPF will automatically compute and signal an inter-area LSP via the ABR nodes that is connected to the transit area.

However, when the ingress LER for the LSP is the ABR connected to the transit area and the destination of the LSP is the address corresponding to another ABR router-id in that same transit area, CSPF will compute and signal an intra-area LSP using the transit area TE links, even when the destination router-id is only part of area 0.

Area Border Node FRR Protection for Inter Area LSP

For protection of the area border router, the upstream node of the area border router acts as a point-of-local-repair (PLR), and the next-hop node to the protected domain border router is the merge-point (MP). Both manual and dynamic bypass are available to protect area border node.

Manual bypass protection works only when a proper completely strict path is provisioned that avoids the area border node.

Dynamic bypass protection provides for the automatic computation, signaling, and association with the primary path of an inter-area P2P LSP to provide ABR node protection. Figure 32 illustrates the role of each node in the ABR node protection using a dynamic bypass LSP.

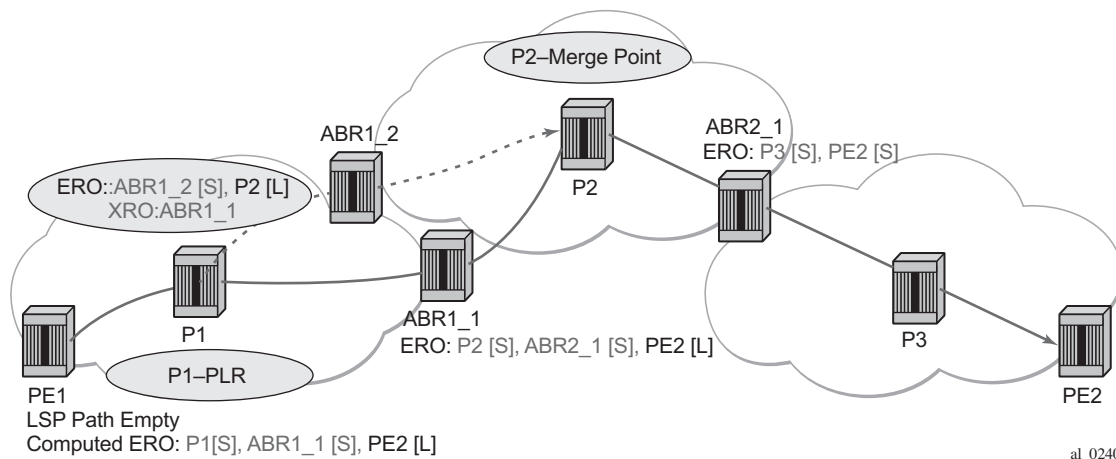


Figure 32: ABR Node Protection Using Dynamic Bypass LSP

In order for a PLR node within the local area of the ingress LER to provide ABR node protection, it must dynamically signal a bypass LSP and associate it with the primary path of the inter-area LSP using the following new procedures:

- The PLR node must inspect the node-id RRO of the LSP primary path to determine the address of the node immediately downstream of the ABR in the other area.
- The PLR signals an inter-area bypass LSP with a destination address set to the address downstream of the ABR node and with the XRO set to exclude the node-id of the protected ABR node.
- The request to CSPF is for a path to the merge-point (i.e., the next-next-hop in the RRO received in the RESV for the primary path) along with the constraint to exclude the protected ABR node and the include/exclude admin-groups of the primary path. If CSPF returns a path that can only go to an intermediate hop, then the PLR node signals the dynamic bypass and will automatically include the XRO with the address of the protected ABR node and propagate the admin-group constraints of the primary path into the Session Attribute object of the bypass LSP. Otherwise, the PLR signals the dynamic bypass directly to the merge-point node with no XRO object in the Path message.
- If a node-protect dynamic bypass cannot be found or signaled, the PLR node attempts a link-protect dynamic bypass LSP. As in existing implementation of dynamic bypass within the same area, the PLR attempts in the background to signal a node-protect bypass at the receipt of every third Resv refresh message for the primary path.

- Refresh reduction over dynamic bypass will only work if the node-id RRO also contains the interface address. Otherwise the neighbor will not be created once the bypass is activated by the PLR node. The Path state will then time out after three refreshes following the activation of the bypass backup LSP.

Note that a one-to-one detour backup LSP cannot be used at the PLR for the protection of the ABR node. As a result, a 7x50 PLR node will not signal a one-to-one detour LSP for ABR protection. In addition, an ABR node will reject a Path message, received from a third party implementation, with a detour object and with the ERO having the next-hop loose. This is performed regardless if the **cspf-on-loose** option is enabled or not on the 7x50 node. In other words, the 7x50 as a transit ABR for the detour path will reject the signaling of an inter-area detour backup LSP.

Automatic Creation of a RSVP Mesh LSP

Feature Configuration

The user first creates an LSP template of type mesh P2P:

```
config>router>mpls>lsp-template template-name mesh-p2p
```

Inside the template the user configures the common LSP and path level parameters or options shared by all LSPs using this template.

Then the user references the peer prefix list which is defined inside a policy statement defined in the global policy manager.

```
config>router>mpls>auto-lsp lsp-template template-name policy peer-prefix-policy
```

The user can associate multiple templates with same or different peer prefix policies. Each application of an LSP template with a given prefix in the prefix list will result in the instantiation of a single CSPF computed LSP primary path using the LSP template parameters as long as the prefix corresponds to a router-id for a node in the TE database. This feature does not support the automatic signaling of a secondary path for an LSP. If the user requires the signaling of multiple LSPs to the same destination node, he/she must apply a separate LSP template to the same or different prefix list which contains the same destination node. Each instantiated LSP will have a unique LSP-id and a unique tunnel-ID. This feature also does not support the signaling of a non-CSPF LSP. The selection of the **'no cspf'** option in the LSP template is thus blocked.

Up to 5 peer prefix policies can be associated with a given LSP template at all times. Each time the user executes the above command, with the same or different prefix policy associations, or the user changes a prefix policy associated with an LSP template, the system re-evaluates the prefix policy. The outcome of the re-evaluation will tell MPLS if an existing LSP needs to be torn down or a new LSP needs to be signaled to a destination address which is already in the TE database.

If a /32 prefix is added to (removed from) or if a prefix range is expanded (shrunk) in a prefix list associated with a LSP template, the same prefix policy re-evaluation described above is performed.

The user must perform a **no shutdown** of the template before it takes effect. Once a template is in use, the user must shutdown the template before effecting any changes to the parameters except for those LSP parameters for which the change can be handled with the Make-Before-Break (MBB) procedures. These parameters are **bandwidth** and enabling **fast-reroute** without the **hop-limit** or **node-protect** options. For all other parameters, the user shuts down the template and once it is added, removed or modified, the existing instances of the LSP using this template are torn down and re-signaled.

Finally the auto-created mesh LSP can be signaled over both numbered and unnumbered RSVP interfaces.

Feature Behavior

Whether the prefix list contains one or more specific /32 addresses or a range of addresses, an external trigger is required to indicate to MPLS to instantiate an LSP to a node which address matches an entry in the prefix list. The objective of the feature is to provide an automatic creation of a mesh of RSVP LSP to achieve automatic tunneling of LDP-over-RSVP. The external trigger is when the router with the router-id matching an address in the prefix list appears in the Traffic Engineering database. In the latter case, the TE database provides the trigger to MPLS which means this feature operates with CSPF LSP only.

Each instantiation of an LSP template results in RSVP signaling and installing state of a primary path for the LSP to the destination router. The auto- LSP is installed in the Tunnel Table Manager (TTM) and is available to applications such as LDP-over-RSVP, resolution of BGP label routes, resolution of BGP, IGP, and static routes. The auto-LSP can also be used for auto-binding by a VPRN service. The auto-LSP is however not available to be used in a provisioned SDP for explicit binding by services. A consequence of this is that an auto-LSP can also not be used directly for auto-binding of a PW template with the **use-provisioned-sdp** option in BGP-AD VPLS, or FEC129 VLL service. However, an auto-binding of a PW template to an LDP LSP, which is then tunneled over an RSVP auto-LSP is supported.

If the user changes the **bandwidth** parameter in the LSP template, an MBB is performed for all LSPs using the template. If however the **auto-bandwidth** option was enabled in the template, the bandwidth **parameter** change will be saved but will only take effect at the next time the LSP bounces or is re-signaled.

Except for the MBB limitations to the configuration parameter change in the LSP template, MBB procedures for manual and timer based re-signaling of the LSP, for TE Graceful Shutdown and for soft pre-emption are supported.

Note that the use of the '**tools perform router mpls update-path**' command with a mesh LSP is not supported.

The **one-to-one** option under **fast-reroute** is also not supported.

If while the LSP is UP, with the bypass backup path activated or not, the TE database loses the router-id, it will perform an update to MPLS module which will state router-id is no longer in TE database. This will cause MPLS to tear down all mesh LSPs to this router-id. Note however that if the destination router is not a neighbor of the ingress LER and the user shuts down the IGP instance in the destination router, the router-id corresponding to the IGP instance will only be deleted from the TE database in the ingress LER after the LSA/LSP ages out. If the user brought back up the IGP instance before the LSA/LSP aged out, the ingress LER will delete and re-install the same router-id at the receipt of the updated LSA/LSP. In other words, the RSVP LSPs destined

to this router-id will get deleted and re-established. All other failure conditions will cause the LSP to activate the bypass backup LSP or to go down without being deleted.

There is no overall chassis mode restrictions enforced with the mesh LSP feature. If the chassis-mode, network chassis-mode or IOM type requirements for a feature are not met, the configuration of the corresponding command will not be allowed into the LSP template on the system.

Multi-Area and Multi-Instance Support

A router which does not have TE links within a given IGP area/level will not have its router-id discovered in the TE database by other routers in this area/level. In other words, an auto-LSP of type P2P mesh cannot be signaled to a router which does not participate in the area/level of the ingress LER.

A mesh LSP can however be signaled using TE links all belonging to the same IGP area even if the router-id of the ingress and egress routers are interfaces reachable in a different area. In this case, the LSP is considered to be an intra-area LSP.

If multiple instances of ISIS or OSPF are configured on a router, each with its own router-id value, the TE database in other routers will be able to discover TE links advertised by each instance. In such a case, an instance of an LSP can be signaled to each router-id with a CSPF path computed using TE links within each instance.

Finally, if multiple instances of ISIS or OSPF are configured on a destination router each with the same router-id value, a single instance of LSP will be signaled from other routers. If the user shuts down one IGP instance, this will be **no op** as long as the other IGP instances remain up. The LSP will remain up and will forward traffic using the same TE links. The same behavior exists with a provisioned LSP.

Mesh LSP Name Encoding and Statistics

When the ingress LER signals the path of a mesh auto-LSP, it includes the name of the LSP and that of the path in the Session Name field of the Session Attribute object in the Path message. The encoding is as follows:

Session Name: <lsp-name::path-name>, where lsp-name component is encoded as follows:

TemplateName-DestIpv4Address-TunnelId

Where ***DestIpv4Address*** is the address of the destination of the auto-created LSP.

At ingress LER, the user can enable egress statistics for the auto-created mesh LSP by adding the following configuration to the LSP template:

```
config
  router
```

```
[no] mpls
    lsp-template template-name mesh-p2p]
no lsp-template template-name
    [no] egress-statistics
        accounting-policy policy-id
        no accounting-policy
    [no] collect-stats
```

If there are no stat indices available when an LSP is instantiated, the assignment is failed and the egress-statistics field in the show command for the LSP path will be in operational DOWN state but in admin UP state.

An auto-created mesh LSP can also have ingress statistics enabled on the egress LER as long as the user specifies the full LSP name following the above syntax.

```
configure>router>mpls>ingress-statistics>lsp lsp-name sender ip-address
```

Automatic Creation of an RSVP One-Hop LSP

Feature Configuration

The user first creates an LSP template of type one-hop:

```
config>router>mpls>lsp-template template-name one-hop-p2p
```

Then the user enables the automatic signaling of one-hop LSP to all direct neighbors using the following command:

```
config>router>mpls>auto-lsp lsp-template template-name one-hop
```

The LSP and path parameters and options supported in a LSP template of type **one-hop-p2p** are that same as in the LSP template of type **mesh-p2p** except for the parameter **from** which is not allowed in a template of type **one-hop-p2p**. The show command for the auto-LSP will display the actual outgoing interface address in the 'from' field. The full list of template parameters is shown in the CLI Section. Also, the rules for adding or modifying the template parameters are as described in 7.1.1.

Finally the auto-created one-hop LSP can be signaled over both numbered and unnumbered RSVP interfaces.

Feature Behavior

Although the provisioning model and CLI syntax differ from that of a mesh LSP only by the absence of a prefix list, the actual behavior is quite different. When the above command is executed, the TE database will keep track of each TE link which comes up to a directly connected IGP neighbor which router-id is discovered. It then instructs MPLS to signal an LSP with a destination address matching the router-id of the neighbor and with a strict hop consisting of the address of the interface used by the TE link. Thus the **auto-lsp** command with the **one-hop** option will result in one or more LSPs signaled to the IGP neighbor.

Only the router-id of the first IGP instance of the neighbor which advertises a TE link will cause the LSP to be signaled. If subsequently another IGP instance with a different router-id advertises the same TE link, no action is taken and the existing LSP is kept up. If the router-id originally used disappears from the TE database, the LSP is kept up and is associated now with the other router-id.

The state of a one-hop LSP once signaled follows the following behavior:

- If the interface used by the TE link goes down or BFD times out and the RSVP interface registered with BFD, the LSP path moves to the bypass backup LSP if the primary path is associated with one.
- If while the one-hop LSP is UP, with the bypass backup path activated or not, the association of the TE-link with a router-id is removed in the TE databases, the one-hop

LSP is torn down. This would be the case if the interface used by the TE link is deleted or if the interface is shutdown in the context of RSVP.

- If while the LSP is UP, with the bypass backup path activated or not, the TE database loses the router-id, it will perform two separate updates to MPLS module. The first one updates the loss of the TE link association which will cause action (B) above for the one-hop LSP. The other update will state router-id is no longer in TE database which will cause MPLS to tear down all mesh LSPs to this router-id as explained in Section 7.1.2. Note however that a shutdown at the neighbor of the IGP instance which advertised the router-id will cause the router-id to be removed from the ingress LER node immediately after the last IGP adjacency is lost and is not subject to age-out as for a non-directly connected destination router.

All other feature behavior, limitations, and statistics support are the same as for an auto-LSP of type **mesh-p2p**.

Point-to-Multipoint (P2MP) RSVP LSP

Point-to-multipoint (P2MP) RSVP LSP allows the source of multicast traffic to forward packets to one or many multicast receivers over a network without requiring a multicast protocol, such as PIM, to be configured in the network core routers. A P2MP LSP tree is established in the control plane which path consists of a head-end node, one or many branch nodes, and the leaf nodes. Packets injected by the head-end node are replicated in the data plane at the branching nodes before they are delivered to the leaf nodes.

Application in Video Broadcast

Figure 33 illustrates the use of the SR product family in triple play application (TPSDA). The Broadband Service Router (BSR) is a 7750 SR and the Broadband Service Aggregator (BSA) is the 7450 ESS.




Figure 33: Application of P2MP LSP in Video Broadcast

A PIM-free core network can be achieved by deploying P2MP LSPs using other core routers. The router can act as the ingress LER receiving the multicast packets from the multicast source and forwarding them over the P2MP LSP.

A router can act as a leaf for the P2MP LSP tree initiated from the head-end router co-located with the video source. The router can also act as a branch node serving other leaf nodes and supports the replication of multicast packets over P2MP LSPs.

P2MP LSP Data Plane

A P2MP LSP is a unidirectional label switched path (LSP) which inserts packets at the root (ingress LER) and forwards the exact same replication of the packet to one or more leaf nodes (egress LER). The packet can be replicated at the root of P2MP LSP tree and/or at a transit LSR which acts as a branch node for the P2MP LSP tree.

Note that the data link layer code-point, for example Ethertype when Ethernet is the network port, continues to use the unicast codepoint defined in RFC 3032, *MPLS Label Stack Encoding*, and which is used on P2P LSP. This change is specified in draft-ietf-mpls-multicast-encaps, *MPLS Multicast Encapsulations*.

When a router sends a packet over a P2MP LSP which egresses on an Ethernet-based network interface, the Ethernet frame uses a MAC unicast destination address when sending the packet over the primary P2MP LSP instance or over a P2P bypass LSP). Note that a MAC multicast destination address is also allowed in the draft-ietf-mpls-multicast-encaps. Thus, at the ingress network interface on an Ethernet port, the router can accept both types of Ethernet destination addresses.

Procedures at Ingress LER Node

The following procedures occur at the root of the P2MP LSP (head-end or ingress LER node):

1. First, the P2MP LSP state is established via the control plane. Each leaf of the P2MP LSP will have a next-hop label forwarding entry (NHLFE) configured in the forwarding plane for each outgoing interface.
1. The user maps a specific multicast destination group address to the P2MP LSP in the base router instance by configuring a static multicast group under a tunnel interface representing the P2MP LSP.
2. An FTN entry is programmed at the ingress of the head-end node that maps the FEC of a received user IP multicast packet to a list of outgoing interfaces (OIF) and corresponding NHLFEs.
3. The head-end node replicates the received IP multicast packet to each NHLFE. Replication is performed at ingress toward the fabric and/or at egress forwarding engine depending on the location of the OIF.
4. At ingress, the head-end node performs a PUSH operation on each of the replicated packets.

Procedures at LSR Node

The following procedures occur at an LSR node that is not a branch node:

- The LSR performs a label swapping operation on a leaf of the P2MP LSP. This is a conventional operation of an LSR in a P2P LSP. An ILM entry is programmed at the ingress of the LSR to map an incoming label to a NHLFE.

The following is an exception handling procedure for control packets received on an ILM in an LSR.

- Packets that arrive with the TTL in the outer label expiring are sent to the CPM for further processing and are not forwarded to the egress NHLFE.
-

Procedures at Branch LSR Node

The following procedures occur at an LSR node that is a branch node:

- The LSR performs a replication and a label swapping for each leaf of the P2MP LSP. An ILM entry is programmed at the ingress of the LSR to map an incoming label to a list of OIF and corresponding NHLFEs.
- There is a limit of 127 OIF/NHLFEs per ILM entry.

The following is an exception handling procedure for control packets received on an ILM in a branch LSR:

- Packets that arrive with the TTL in the outer label expiring are sent to the CPM for further processing and not copied to the LSP branches.

Procedures at Egress LER Node

The following procedures occur at the leaf node of the P2MP LSP (egress LER):

- The egress LER performs a pop operation. An ILM entry is programmed at the ingress of the egress LER to map an incoming label to a list of next-hop/OIF.

The following is an exception handling procedure for control packets received on an ILM in an egress LER.

- The packet is sent to the CPM for further processing if there is any of the IP header exception handling conditions set after the label is popped: 127/8 destination address, router alert option set, or any other options set.
-

Procedures at BUD LSR Node

The following are procedures at an LSR node which is both a branch node and an egress leaf node (bud node):

- The bud LSR performs a pop operation on one or many replications of the received packet and a swap operation of the remaining replications. An ILM entry is programmed at ingress of the LSR to map the incoming label to list of NHLFE/OIF and next-hop/OIF.

Note however, the exact same packets are replicated to an LSP leaf and to a local interface.

The following are the exception handling procedures for control packets received on an ILM in a bud LSR:

- Packets which arrive with the TTL in the outer label expiring are sent to the CPM and are not copied to the LSP branches.
- Packets whose TTL does not expire are copied to all branches of the LSP. The local copy of the packet is sent to the CPM for further processing if there is any of the IP header exception handling conditions set after the label is popped: 127/8 destination address, router alert option set, or any other options set.

Ingress Path Management for P2MP LSP Packets

The SR OS provides the ingress multicast path management (IMPM) capability that allows users to manage the way IP multicast streams are forwarded over the router's fabric and to maximize the use of the fabric multicast path capacity.

IMPM consists of two components, a bandwidth policy and a multicast information policy. The bandwidth policy configures the parameters of the multicast paths to the fabric. This includes the rate limit and the multicast queue parameters of each path. The multicast information policy configures the bandwidth and preference parameters of individual multicast flows corresponding to a channel, for example, a $\langle *,G \rangle$ or a $\langle S,G \rangle$, or a bundle of channels.

By default both, the IOM-2 and IOM-3/IMM ingress data paths provide two multicast paths through the fabric referred to as high-priority path and low-priority path respectively. When a multicast packet is received on an ingress network or access interface or on a VPLS SAP, the packet's classification will determine its forwarding class and priority or profile as per the ingress QoS policy. This then determines which of the SAP or interface multicast queues it must be stored in. By default SAP and interface expedited forwarding class queues forward over the high-priority multicast path and the non expedited forwarding class queues forward over the low-priority multicast path.

When IMPM on the ingress MDA is enabled, one or more multicast paths are enabled depending on the IOM type. In addition, multicast flows managed by IMPM will be stored in a separate shared multicast queue for each multicast path. These queues are configured in the bandwidth policy.

IMPM maps a packet to one of the paths dynamically based on monitoring the bandwidth usage of each packet flow matching a $\langle *,G \rangle$ or $\langle S,G \rangle$ record. The multicast bandwidth manager assigns multicast flows to a primary path, and ancillary path for IOM-2, based on the flow preference until the rate limits of each path is reached. At that point in time, a multicast flow is mapped to the secondary flow. If a path congests, the bandwidth manager will remove and black-hole lower preference flows to guarantee bandwidth to higher preference flows. The preference of a multicast flow is configured in the multicast info policy.

A packet received on a P2MP LSP ILM is managed by IMPM when IMPM is enabled on the ingress MDA and the packet matches a specific multicast record. When IMPM is enabled but the packet does not match a multicast record, or when IMPM is disabled, a packet received on a P2MP LSP ILM is mapped to a multicast path differently depending if the ingress IOM is an IOM-2 or IOM-3.

Ingress P2MP Path Management on IOM-3/IMMs

On an ingress IOM-3/IMM, there are multiple multicast paths available to forward multicast packets, depending on the hardware being used. Each path has a set of multicast queues and

associated with it. Two paths are enabled by default, a primary path and a secondary path, and represent the high-priority and low-priority paths respectively. Each VPLS SAP, access interface, and network interface will have a set of per forwarding class multicast and/or broadcast queues which are defined in the ingress QoS policy associated with them. The expedited queues will be attached to the primary path while the non-expedited queues will be attached to secondary path.

When IMPM is enabled and/or when a P2MP LSP ILM exists on the ingress IOM-3/IMM, the remaining multicast paths are also enabled. 16 multicast paths are supported by default with 28 on 7950 XRS systems and 7750 SR12-e systems, with the latter having the **tools perform system set-fabric-speed fabric-speed-b**. One path remains as a secondary path and the rest are primary paths.

A separate pair of shared multicast queues is created on each of the primary paths, one for IMPM managed packets and one for P2MP LSP packets not managed by IMPM. The secondary path does not forward IMPM managed packets or P2MP LSP packets. These queues have default rate (PIR=CIR) and CBS/MBS/Hi-Priority-Only thresholds but can be changed away from default under the bandwidth policy.

A VPLS snooped packet, a PIM routed packet, or a P2MP LSP packet is managed by IMPM if it matches a $\langle *,G \rangle$ or a $\langle S,G \rangle$ multicast record in the ingress forwarding table and IMPM is enabled on the ingress MDA where the packet is received. The user enables IMPM on the ingress MDA data path using the **config>card>mda>ingress>mcst-path-management** command.

A packet received on an IP interface and to be forwarded to a P2MP LSP NHLFE or a packet received on a P2MP LSP ILM is not managed by IMPM when IMPM is disabled on the ingress MDA where the packet is received or when IMPM is enabled but the packet does not match any multicast record. A P2MP LSP packet duplicated at a branch LSR node is an example of a packet not managed by IMPM even when IMPM is enabled on the ingress MDA where the P2MP LSP ILM exists. A packet forwarded over a P2MP LSP at an ingress LER and which matches a $\langle *,G \rangle$ or a $\langle S,G \rangle$ is an example of a packet which is not managed by IMPM if IMPM is disabled on the ingress MDA where the packet is received.

When a P2MP LSP packet is not managed by IMPM, it is stored in the unmanaged P2MP shared queue of one of the primary multicast paths.

By default, non-managed P2MP LSP traffic is distributed across the IMPM primary paths using hash mechanisms. This can be optimized by enabling IMPM on any forwarding complex, which allows the system to redistributed this traffic on all forwarding complexes across the IMPM paths to achieve a more even capacity distribution. Be aware that enabling IMPM will cause routed and VPLS (IGMP and PIM) snooped IP multicast groups to be managed by IMPM.

The above ingress data path procedures apply to packets of a P2MP LSP at ingress LER, LSR, branch LSR, bud LSR, and egress LER. Note that in the presence of both IMPM managed traffic and unmanaged P2MP LSP traffic on the same ingress forwarding plane, the user must account for the presence of the unmanaged traffic on the same path when setting the rate limit for an IMPM path in the bandwidth policy.

Ingress P2MP Path Management on IOM-2

The following procedures apply at the ingress data path for packets received from or to be forwarded to a P2MP LSP at ingress LER, LSR, branch LSR, bud LSR, and egress LER.

On ingress IOM-2, there are 3 multicast paths which are available for forwarding multicast packets. Each path has a set of multicast queues and a multicast VoQ associated with it. Two paths are enabled by default, a primary path and a secondary path, and represent the high-priority and low-priority paths respectively. Each VPLS SAP, access interface, and network interface will have a set of per forwarding class multicast and/or broadcast queues which are defined in the ingress QoS policy associated with them. The expedited queues will be attached to the primary path while the non-expedited queues will be attached to the secondary path.

When IMPM is disabled, packets of P2MP LSP arriving on a network interface will be queued in that interface queue corresponding to the forwarding class of the packet.

When the user enables IMPM on the ingress MDA, a third multicast path, referred to as ancillary path, is added on the ingress IOM-2. This path reuses unused capacity from the unicast paths. The high-priority and low-priority paths are renamed as primary and secondary paths respectively.

A VPLS snooped packet or a PIM routed packet is managed by IMPM if it matches a <*,G> or a <S,G> multicast record in the ingress IOM-2 forwarding table and IMPM is enabled on the ingress MDA where the packet is received. The user enables IMPM on the ingress MDA data path using the **config>card>mda>ingress>mcast-path-management** command.

A P2MP LSP packet which matches a multicast record is also managed by IMPM on ingress IOM-2 and is thus distributed to one of the primary, ancillary, or secondary path according to the congestion level of the paths and the preference of the packet's multicast flow as configured in the multicast info policy 2.

RSVP Control Plane in a P2MP LSP

P2MP RSVP LSP is specified in RFC 4875, *Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)*.

A P2MP LSP is modeled as a set of root-to-leaf (S2L) sub-LSPs. The root, for example the head-end node, triggers signaling using one or multiple path messages. A path message can contain the signaling information for one or more S2L sub-LSPs. The leaf sub-LSP paths are merged at branching points.

A P2MP LSP is identified by the combination of <P2MP ID, tunnel ID, extended tunnel ID> part of the P2MP session object, and <tunnel sender address, LSP ID> fields in the P2MP sender_template object.

A specific sub-LSP is identified by the <S2L sub-LSP destination address> part of the S2L_SUB_LSP object and an ERO and secondary ERO (SERO) objects.

The following are characteristics of this feature:

1. Supports the de-aggregated method for signaling the P2MP RSVP LSP. Each root to leaf is modeled as a P2P LSP in the RSVP control plane. Only data plane merges the paths of the packets.
2. Each S2L sub-LSP is signaled in a separate path message. Each leaf node responds with its own resv message. A branch LSR node will forward the path message of each S2L sub-LSP to the downstream LSR without replicating it. It will also forward the resv message of each S2L sub-LSP to the upstream LSR without merging it with the resv messages of other S2L sub-LSPs of the same P2MP LSP. The same is done for subsequent refreshes of the path and resv states.
3. The node will drop aggregated RSVP messages on the receive side if originated by another vendor's implementation.
4. The user configures a P2MP LSP by specifying the optional create-time parameter **p2mp-lsp** following the LSP name. Next, the user creates a primary P2MP instance using the keyword **primary-p2mp-instance**. Then a path name of each S2L sub-LSP must be added to the P2MP instance using the keyword **s2l-path**. The paths can be empty paths or can specify a list of explicit hops. The path name must exist and must have been defined in the **config>router>mpls>path** context.
5. The same path name can be re-used by more than one S2L of the primary P2MP instance. However the to keyword must have a unique argument per S2L as it corresponds to the address of the egress LER node.
6. The user can configure a secondary instance of the P2MP LSP to backup the primary one. In this case, the user enters the name of the secondary P2MP LSP instance under the same LSP name. One or more secondary instances can be created. The trigger for the head-end

node to switch the path of the LSP from the primary P2MP instance to the secondary P2MP instance is to be determined. This could be based on the number of leaf LSPs which went down at any given time.

7. The following parameters can be used with a P2MP LSP: adaptive, cspf, exclude, fast-reroute, from, hop-limit, include, metric, retry-limit, retry-timer, resignal-timer.
8. The following parameters cannot be used with a P2MP LSP: adspec, primary, secondary, to.
9. The node ingress LER will not inset an adspec object in the path message of an S2L sub-LSP. If received in the resv message, it will be dropped. The operational MTU of an S2L path is derived from the MTU of the outgoing interface of that S2L path.
10. The **to** parameter is not available at the LSP level but at the path level of each S2L sub-LSP of the primary or secondary instance of this P2MP LSP.
11. The hold-timer configured in the **config>router>mpls>hold-timer** context applies when signaling or re-signaling an individual S2L sub-LSP path. It does not apply when the entire tree is signaled or re-signaled.
12. The head-end node can add and/or remove a S2L sub-LSP of a specific leaf node without impacting forwarding over the already established S2L sub-LSPs of this P2MP LSP and without re-signaling them.
13. The head-end node performs a make-before break (MBB) on an individual S2L path of a primary P2MP instance whenever it applies the FRR global revertive procedures to this path. If CSPF finds a new path, RSVP signals this S2L path with the same LSP-ID as the existing path.
14. All other configuration changes, such as adaptive/no-adaptive, use-te-metric, no-frr, cspf/no-cspf, result in the tear-down and re-try of all affected S2L paths as is the case for P2P LSP paths.
15. MPLS requests CSPF to re-compute the whole set of S2L paths of a given active P2MP instance each time the P2MP re-signal timer expires. The P2MP re-signal timer is configured separately from the P2P LSP. MPLS performs a global MBB and moves each S2L sub-LSP in the instance into its new path using a new P2MP LSP ID if the global MBB is successful. This is regardless of the cost of the new S2L path.
16. MPLS will request CSPF to re-compute the whole set of S2L paths of a given active P2MP instance each time the user performs a manual re-signal of the P2MP instance. MPLS then always performs a global MBB and moves each S2L sub-LSP in the instance into its new path using a new P2MP LSP ID if the global MBB is successful. This is regardless of the cost of the new S2L path. The user executes a manual re-signal of the P2MP LSP instance using the command: **tools>perform>router>mpls>resignal p2mp-lsp *lsp-name* p2mp-instance *instance-name***.
17. When performing global MBB, MPLS runs a separate MBB on each S2L in the P2MP LSP instance. If an S2L MBB does not succeed the first time, MPLS will re-try the S2L using the re-try timer and re-try count values inherited from P2MP LSP configuration.

However, there will be a global MBB timer set to 600 seconds and which is not configurable. If the global MBB succeeds, for example, all S2L MBBs have succeeded, before the global timer expires, MPLS moves the all S2L sub-LSPs into their new path. Otherwise when this timer expires, MPLS checks if all S2L paths have at least tried once. If so, it then aborts the global MBB. If not, it will continue until all S2Ls have re-tried once and then aborts the global MBB. Once global MBB is aborted, MPLS will move all S2L sub-LSPs into the new paths only if the set of S2Ls with a new path found is a superset of the S2Ls which have a current path which is up.

18. While make-before break is being performed on individual S2L sub-LSP paths, the P2MP LSP will continue forwarding packets on S2L sub-LSP paths which are not being re-optimized and on the older S2L sub-LSP paths for which make-before-break operation was not successful. MBB will thus result in duplication of packets until the old path is torn down.
19. The MPLS data path of an LSR node, branch LSR node, and bud LSR node will be able to re-merge S2L sub-LSP paths of the same P2MP LSP in case their ILM is on different incoming interfaces and their NHLFE is on the same or different outgoing interfaces. This could occur anytime there are equal cost paths through this node for the S2L sub-LSPs of this P2MP LSP.
20. Link-protect FRR bypass using P2P LSPs is supported. In link protect, the PLR protecting an interface to a branch LSR will only make use of a single P2P bypass LSP to protect all S2L sub-LSPs traversing the protected interface.
21. Refresh reduction on RSVP interface and on P2P bypass LSP protecting one or more S2L sub-LSPs.
22. A manual bypass LSP cannot be used for protecting S2L paths of a P2MP LSP.
23. The following MPLS features do operate with P2MP LSP:
 - BFD on RSVP interface.
 - MD5 on RSVP interface.
 - IGP metric and TE metric for computing the path of the P2MP LSP with CSPF.
 - SRLG constraint for computing the path of the P2MP LSP with CSPF. SRLG is supported on FRR backup path only.
 - TE graceful shutdown.
 - Admin group constraint.
24. The following MPLS features are not operable with P2MP LSP:
 - Class based forwarding over P2MP RSVP LSP.
 - LDP-over-RSVP where the RSVP LSP is a P2MP LSP.
 - Diff-Serv TE.
 - Soft pre-emption of RSVP P2MP LSP.

Forwarding Multicast Packets over RSVP P2MP LSP in the Base Router

Multicast packets are forwarded over the P2MP LSP at the ingress LER based on a static join configuration of the multicast group against the tunnel interface associated with the originating P2MP LSP. At the egress LER, packets of a multicast group are received from the P2MP LSP via a static assignment of the specific <S,G> to the tunnel interface associated with a terminating LSP.

Procedures at Ingress LER Node

The forwarding of multicast packets over a P2MP LSP follows the following procedures:

1. The user creates a tunnel interface associated with the P2MP LSP: **configure>router>tunnel-interface rsvp-p2mp lsp-name**. The **configure>router>pim>tunnel-interface** command has been discontinued.
2. The user adds static multicast group joins to the PIM interface, either as a specific <S,G> or as a <*,G>: **configure>router>igmp>tunnel-interface>static>group>source ip-address** and **configure>router>igmp>tunnel-interface>static>group>starg**.

The tunnel interface identifier consists of a string of characters representing the LSP name for the RSVP P2MP LSP. Note that MPLS will actually pass to PIM a more structured tunnel interface identifier. The structure will follow the one BGP uses to distribute the PMSI tunnel information in BGP multicast VPN as specified in draft-ietf-l3vpn-2547bis-mcast-bgp, *Multicast in MPLS/BGP IP VPNs*. The format is: <extended tunnel ID, reserved, tunnel ID, P2MP ID> as encoded in the RSVP-TE P2MP LSP session_attribute object in RFC 4875.

The user can create one or more tunnel interfaces in PIM and associate each to a different RSVP P2MP LSP. The user can then assign static multicast group joins to each tunnel interface. Note however that a given <*,G> or <S,G> can only be associated with a single tunnel interface.

A multicast packet which is received on an interface and which succeeds the RPF check for the source address will be replicated and forwarded to all OIFs which correspond to the branches of the P2MP LSP. The packet is sent on each OIF with the label stack indicated in the NHLFE of this OIF. The packets will also be replicated and forwarded natively on all OIFs which have received IGMP or PIM joins for this <S,G>.

The multicast packet can be received over a PIM or IGMP interface which can be an IES interface, a spoke SDP-terminated IES interface, or a network interface.

In order to duplicate a packet for a multicast group over the OIF of both P2MP LSP branches and the regular PIM or IGMP interfaces, the tap mask for the P2MP LSP and that of the PIM based interfaces will need to be combined into a superset MCID.

Procedures at Egress LER Node

Procedures with a Primary Tunnel Interface

The user configures a tunnel interface and associates it with a terminating P2MP LSP leaf using the command: **config>router>tunnel-interface rsvp-p2mp lsp-name sender sender-address**. The **configure>router>pim>tunnel-interface** command has been discontinued.

The tunnel interface identifier consists of a couple of string of characters representing the LSP name for the RSVP P2MP LSP followed by the system address of the ingress LER. The LSP name must correspond to a P2MP LSP name configured by the user at the ingress LER and must not contain the special character “:” Note that MPLS will actually pass to PIM a more structured tunnel interface identifier. The structure will follow the one BGP uses to distribute the PMSI tunnel information in BGP multicast VPN as specified in draft-ietf-l3vpn-2547bis-mcast-bgp. The format is: <extended tunnel ID, reserved, tunnel ID, P2MP ID> as encoded in the RSVP-TE P2MP LSP session_attribute object in RFC 4875.

The egress LER accepts multicast packets the following methods:

1. The regular RPF check on unlabeled IP multicast packets, which is based on routing table lookup.
2. The static assignment which specifies the receiving of a multicast group <*,G> or a specific <S,G> from a primary tunnel-interface associated with an RSVP P2MP LSP.

One or more primary tunnel interfaces in the base router instance can be configured. In other words, the user will be able to receive different multicast groups, <*,G> or specific <S,G>, from different P2MP LSPs. This assumes that the user configured static joins for the same multicast groups at the ingress LER to forward over a tunnel interface associated with the same P2MP LSP.

A multicast info policy CLI option allows the user to define a bundle and specify channels in the bundle that must be received from the primary tunnel interface. The user can apply the defined multicast info policy to the base router instance.

At any given time, packets of the same multicast group can be accepted from either the primary tunnel interface associated with a P2MP LSP or from a PIM interface. These are mutually exclusive options. As soon as a multicast group is configured against a primary tunnel interface in the multicast info policy, it is blocked from other PIM interfaces.

However, if the user configured a multicast group to be received from a given primary tunnel interface, there is nothing preventing packets of the same multicast group from being received and accepted from another primary tunnel interface. However, an ingress LER will not allow the same multicast group to be forwarded over two different P2MP LSPs. The only possible case is that of two ingress LERs forwarding the same multicast group over two P2MP LSPs towards the same egress LER.

A multicast packet received on a tunnel interface associated with a P2MP LSP can be forwarded over a PIM or IGMP interface which can be an IES interface, a spoke SDP terminated IES interface, or a network interface.

Note that packets received from a primary tunnel-interface associated with a terminating P2MP LSP cannot be forwarded over a tunnel interface associated with an originating P2MP LSP.

MPLS Service Usage

TiMetra Systems routers enable service providers to deliver virtual private networks (VPNs) and Internet access using Generic Routing Encapsulation (GRE) and/or MPLS tunnels, with Ethernet and/or SONET/SDH interfaces.

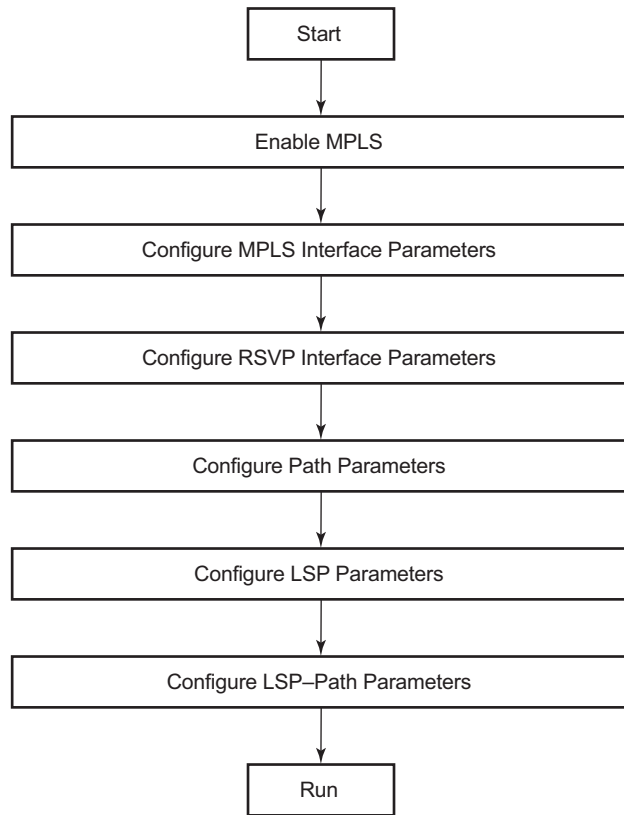
Service Distribution Paths

A service distribution path (SDP) acts as a logical way of directing traffic from one router to another through a uni-directional (one-way) service tunnel. The SDP terminates at the far-end router which directs packets to the correct service egress service access point (SAP) on that device. All services mapped to an SDP use the same transport encapsulation type defined for the SDP (either GRE or MPLS).

For information about service transport tunnels, refer to the Service Distribution Paths (SDPs) section in the OS Services Guide. They can support up to eight forwarding classes and can be used by multiple services. Multiple LSPs with the same destination can be used to load-balance traffic.

MPLS/RSVP Configuration Process Overview

Figure 34 displays the process to configure MPLS and RSVP parameters.



al_0212

Figure 34: MPLS and RSVP Configuration and Implementation Flow

Configuration Notes

This section describes MPLS and RSVP caveats.

- Interfaces must already be configured in the `config>router>interface` context before they can be specified in MPLS and RSVP.
- A router interface must be specified in the `config>router>mpls` context in order to apply it or modify parameters in the `config>router>rsvp` context.
- A system interface must be configured and specified in the `config>router>mpls` context.
- Paths must be created before they can be applied to an LSP.

Configuring MPLS and RSVP with CLI

This section provides information to configure MPLS and RSVP using the command line interface.

Topics in this section include:

- [MPLS Configuration Overview on page 176](#)
 - [LSPs on page 176](#)
 - [Paths on page 176](#)
 - [Router Interface on page 177](#)
 - [Choosing the Signaling Protocol on page 177](#)
- [Basic MPLS Configuration on page 178](#)
- [Common Configuration Tasks on page 179](#)
 - [Configuring MPLS Components on page 180](#)
 - [Configuring Global MPLS Parameters on page 180](#)
 - [Configuring an MPLS Interface on page 181](#)
 - [Configuring MPLS Paths on page 182](#)
 - [Configuring an MPLS LSP on page 183](#)
 - [Configuring Manual Bypass Tunnels on page 185](#)
- [Configuring RSVP Parameters on page 187](#)
 - [Configuring RSVP Message Pacing Parameters on page 188](#)
 - [Configuring Graceful Shutdown on page 189](#)
- [MPLS Configuration Management Tasks on page 190](#)
- [RSVP Configuration Management Tasks on page 195](#)

MPLS Configuration Overview

Multiprotocol Label Switching (MPLS) enables routers to forward traffic based on a simple label embedded into the packet header. A router examines the label to determine the next hop for the packet, saving time for router address lookups to the next node when forwarding packets. MPLS is not enabled by default and must be explicitly enabled.

In order to implement MPLS, the following entities must be configured:

- [LSPs on page 176](#)
 - [Paths on page 176](#)
 - [Router Interface on page 177](#)
-

LSPs

To configure MPLS-signaled label-switched paths (LSPs), an LSP must run from an ingress router to an egress router. Configure only the ingress router and configure LSPs to allow the software to make the forwarding decisions or statically configure some or all routers in the path. The LSP is set up by Resource Reservation Protocol (RSVP), through RSVP signaling messages. The router automatically manages label values. Labels that are automatically assigned have values ranging from 1,024 through 1,048,575 (see [Label Values on page 24](#)).

A static LSP is a manually set up LSP where the nexthop IP address and the outgoing label are explicitly specified.

Paths

To configure signaled LSPs, you must first create one or more named paths on the ingress router. For each path, the transit routers (hops) in the path are specified.

Router Interface

At least one router interface and one system interface must be defined in the **config>router>interface** context in order to configure MPLS on an interface.

Choosing the Signaling Protocol

In order to configure a static or a RSVP signaled LSP, you must enable MPLS on the router which automatically enables RSVP and adds automatically the system interface into both contexts. Any other network IP interface, other than loopbacks, added to MPLS is also automatically enabled in RSVP and becomes a TE link. When the interface is enabled in RSVP, the IGP instance will advertise the Traffic Engineering (TE) information for the link to other routers in the network in order to build their TE database and compute CSPF paths. Operators must enable the traffic-engineering option in the ISIS or OSPF instance for this. Operators can also configure under the RSVP context of the interface the RSVP protocol parameters for that interface.

If only static label switched paths are used in your configurations, then operators must manually define the paths through the MPLS network. Label mappings and actions configured at each hop must be specified. Operators can disable RSVP on the interface if it is used only for incoming or outgoing static LSP label by shutting down the interface in the RSVP context. The latter causes IGP to withdraw the TE link from its advertisement which removes it from its local and neighbors TE database.

If dynamic LSP signaling is implemented in an operator's network then they must keep RSVP enabled on the interfaces they want to use for explicitly defined or CSPF calculated LSP path.

Basic MPLS Configuration

This section provides information to configure MPLS and configuration examples of common configuration tasks. To enable MPLS, you must configure at least one MPLS interface. The other MPLS configuration parameters are optional. This follow displays an example of an MPLS configuration.

```
A:ALA-1>config>router>mpls# info
-----
admin-group "green" 15
  admin-group "yellow" 20
  admin-group "red" 25
  interface "system"
  exit
  interface "StaticLabelPop"
    admin-group "green"
    label-map 50
    pop
    no shutdown
  exit
  interface "StaticLabelPop"
    label-map 35
    swap 36 nexthop 10.10.10.91
    no shutdown
  exit
  path "secondary-path"
    no shutdown
  exit
  path "to-NYC"
    hop 1 10.10.10.104 strict
    no shutdown
  exit
  lsp "lsp-to-eastcoast"
    to 10.10.10.104
    from 10.10.10.103
    fast-reroute one-to-one
    exit
    primary "to-NYC"
    exit
    secondary "secondary-path"
    exit
    no shutdown
  exit
  static-lsp "StaticLabelPush"
    to 10.10.11.105
    push 60 nexthop 10.10.11.105
    no shutdown
  exit
  no shutdown
-----
A:ALA-1>config>router>mpls#
```

Common Configuration Tasks

This section provides a brief overview of the tasks to configure MPLS and provides the CLI commands.

The following protocols must be enabled on each participating router.

- MPLS
- RSVP (for RSVP-signaled MPLS only), which is automatically enabled when MPLS is enabled.

In order for MPLS to run, you must configure at least one MPLS interface in the **config>router>mpls** context.

- An interface must be created in the **config>router>interface** context before it can be applied to MPLS.
- In the **config>router>mpls** context, configure path parameters for configuring LSP parameters. A path specifies some or all hops from ingress to egress. A path can be used by multiple LSPs.
- When an LSP is created, the egress router must be specified in the **to** command and at least one primary or secondary path must be specified. All other statements under the LSP hierarchy are optional.

Configuring MPLS Components

Use the MPLS and RSVP CLI syntax displayed below for:

- [Configuring Global MPLS Parameters on page 180](#)
 - [Configuring an MPLS Interface on page 181](#)
 - [Configuring MPLS Paths on page 182](#)
 - [Configuring an MPLS LSP on page 183](#)
 - [Configuring a Static LSP on page 184](#)
 - [Configuring Manual Bypass Tunnels on page 185](#)
 - [Configuring RSVP Parameters on page 187](#)
 - [Configuring RSVP Message Pacing Parameters on page 188](#)
 - [Configuring Graceful Shutdown on page 189](#)
-

Configuring Global MPLS Parameters

Admin groups can signify link colors, such as red, yellow, or green. MPLS interfaces advertise the link colors it supports. CSPF uses the information when paths are computed for constrained-based LSPs. CSPF must be enabled in order for admin groups to be relevant.

To configure MPLS admin-group parameters, enter the following commands:

CLI Syntax: `mpls`
 `admin-group group-name group-value`
 `frr-object`
 `resignal-timer minutes`

The following displays an admin group configuration example:

```
A:ALA-1>config>router>mpls# info
-----
    resignal-timer 500
    admin-group "green" 15
    admin-group "yellow" 20
    admin-group "red" 25
    ...
-----
A:ALA-1>config>router>mpls#
```

Configuring an MPLS Interface

Configure the **label-map** parameters if the interface is used in a static LSP.
To configure an MPLS interface on a router, enter the following commands:

```
CLI Syntax: config>router>mpls
                interface
                no shutdown
                admin-group group-name [group-name...(up to 32 max)]
                label-map
                pop
                swap
                no shutdown
                srlg-group group-name [group-name...(up to 5 max)]
                te-metric value
```

The following displays an interface configuration example:

```
A:ALA-1>config>router>mpls# info
-----
...
        interface "to-104"
            admin-group "green"
            admin-group "red"
            admin-group "yellow"
            label-map 35
                swap 36 nexthop 10.10.10.91
            no shutdown
        exit
    exit
    no shutdown
...
-----
A:ALA-1>config>router>mpls#
```

Configuring MPLS Paths

Configure an LSP path to use in MPLS. When configuring an LSP, the IP address of the hops that the LSP should traverse on its way to the egress router must be specified. The intermediate hops must be configured as either **strict** or **loose** meaning that the LSP must take either a direct path from the previous hop router to this router (**strict**) or can traverse through other routers (**loose**).

Use the following CLI syntax to configure a path:

```
CLI Syntax: config>router> mpls
                path path-name
                  hop hop-index ip-address {strict|loose}
                no shutdown
```

The following displays a path configuration example:

```
A:ALA-1>config>router>mpls# info
-----
        interface "system"
        exit
        path "secondary-path"
            hop 1 10.10.0.121 strict
            hop 2 10.10.0.145 strict
            hop 3 10.10.0.1 strict
            no shutdown
        exit
        path "to-NYC"
            hop 1 10.10.10.103 strict
            hop 2 10.10.0.210 strict
            hop 3 10.10.0.215 loose
        exit
-----
A:ALA-1>config>router>mpls#
```

Configuring an MPLS LSP

Configure an LSP path for MPLS. When configuring an LSP, you must specify the IP address of the egress router in the **to** statement. Specify the primary path to be used. Secondary paths can be explicitly configured or signaled upon the failure of the primary path. All other statements are optional.

The following displays an MPLS LSP configuration:

```
A:ALA-1>config>router>mplp# info
-----
...
    lsp "lsp-to-eastcoast"
      to 192.168.200.41
      rsvp-resv-style ff
      cspf
      include "red"
      exclude "green"
      adspec
      fast-reroute one-to-one
      exit
      primary "to-NYC"
        hop-limit 10
      exit
      secondary "secondary-path"
        bandwidth 50000
      exit
      no shutdown
    exit
    no shutdown
-----
A:ALA-1>config>router>mpls#
```

Configuring a Static LSP

An LSP can be explicitly (statically) configured. Static LSPs are configured on every node along the path. The label's forwarding information includes the address of the next hop router.

Use the following CLI syntax to configure a static LSP:


```
CLI Syntax: config>router>mpls
                static-lsp lsp-name
                    to ip-address
                    push out-label nexthop ip-addr
                    no shutdown
```

The following displays a static LSP configuration example:

```
A:ALA-1>config>router>mpls# info
-----
...
    static-lsp "static-LSP"
        to 10.10.10.124
        push 60 nexthop 10.10.42.3
        no shutdown
    exit
...
-----
A:ALA-1>config>router>mpls#
```


Configuring Manual Bypass Tunnels

Consider the following network setup.



The user first configures the option to disable the dynamic bypass tunnels on node B if required. The CLI for this configuration is:

```
config>router>mpls>dynamic-bypass [disable | enable]
```

By default, dynamic bypass tunnels are enabled.

Next, the user configures an LSP on node B, such as B-E-F-C to be used only as bypass. The user specifies each hop in the path, for example, the bypass LSP has a strict path.

Note that including the bypass-only keyword disables the following options under the LSP configuration:

- bandwidth
- fast-reroute
- secondary

The following LSP configuration options are allowed:

- adaptive
- adspec
- cspf
- exclude
- hop-limit
- include
- metric

The following example displays a bypass tunnel configuration:

```
A:ALA-48>config>router>mpls>path# info
-----
...
    path "BEFC"
        hop 10 10.10.10.11  strict
        hop 20 10.10.10.12  strict
        hop 30 10.10.10.13  strict
        no shutdown
    exit

    lsp "bypass-BC"
        to 10.10.10.15
        primary "BEFC"
        exit
        no shutdown
...
-----
A:ALA-48>config>router>mpls>path#
```

Next, the configures an LSP from A to D and indicates fast-reroute bypass protection by selecting facility as the FRR method (**config>router>mpls>lsp>fast-reroute facility**). If the LSP goes through B, and bypass is requested, and the next hop is C, and there is a manually configured bypass-only tunnel from B to C, excluding link BC, then node B uses that.

Configuring RSVP Parameters

RSVP is used to set up LSPs. RSVP must be enabled on the router interfaces that are participating in signaled LSPs. The **keep-multiplier** and **refresh-time** default values can be modified in the RSVP context.

Initially, interfaces are configured in the **config>router>mpls>interface** context. Only these existing (MPLS) interfaces are available to modify in the **config>router> rsvp** context. Interfaces cannot be directly added in the RSVP context.

The following example displays an RSVP configuration example:

```
A:ALA-1>config>router>rsvp# info
-----
      interface "system"
        no shutdown
      exit
      interface to-104
        hello-interval 4000
        no shutdown
      exit
      no shutdown
-----
A:ALA-1>config>router>rsvp#
```

Configuring RSVP Message Pacing Parameters

RSVP message pacing maintains a count of the messages that were dropped because the output queue for the egress interface was full.

Use the following CLI syntax to configure RSVP parameters:

CLI Syntax: config>router>rsvp
no shutdown
msg-pacing
 period *milli-seconds*
 max-burst *number*

The following example displays a RSVP message pacing configuration example:

```
A:ALA-1>config>router>rsvp# info
-----
keep-multiplier 5
refresh-time 60
msg-pacing
    period 400
    max-burst 400
exit
interface "system"
    no shutdown
exit
interface to-104
    hello-interval 4000
    no shutdown
exit
no shutdown
-----
A:ALA-1>config>router>rsvp#
```

Configuring Graceful Shutdown

TE graceful shutdown can be enabled on a specific interface using the **config>router>rsvp>interface>graceful-shutdown** command. This interface is referred to as the maintenance interface.

Graceful shutdown can be disabled by executing the **no** form of the command at the RSVP interface level or at the RSVP level. In this case, the user configured TE parameters of the maintenance links are restored and the maintenance node floods them.

MPLS Configuration Management Tasks

This section discusses the following MPLS configuration management tasks:

- [Modifying MPLS Parameters on page 190](#)
 - [Modifying MPLS Path Parameters on page 192](#)
 - [Modifying MPLS Static LSP Parameters on page 193](#)
 - [Deleting an MPLS Interface on page 194](#)
-

Deleting MPLS

NOTE: In order to remove the MPLS instance, MPLS must be disabled (shutdown) and all SDP bindings to LSPs removed. If MPLS is not shutdown first, when the **no mpls** command is executed, a warning message on the console displays indicating that MPLS is still administratively up.

When MPLS is shut down, the **no mpls** command deletes the protocol instance and removes all configuration parameters for the MPLS instance.

To disable MPLS, use the **shutdown** command.

To remove MPLS on a router, enter the following command:

CLI Syntax: `config>router# no mpls`

Modifying MPLS Parameters

NOTE: You must shut down MPLS entities in order to modify parameters. Re-enable (**no shutdown**) the entity for the change to take effect.

Modifying an MPLS LSP

Some MPLS LSP parameters such as primary and secondary, must be shut down before they can be edited or deleted from the configuration.

The following displays a MPLS LSP configuration example. Refer to the LSP configuration on [page 183](#).

```
A:ALA-1>>config>router>mpls>lsp# info
-----
      shutdown
      to 10.10.10.104
      from 10.10.10.103
      rsvp-resv-style ff
      include "red"
      exclude "green"
      fast-reroute one-to-one
      exit
      primary "to-NYC"
        hop-limit 50
      exit
      secondary "secondary-path"
      exit
-----
A:ALA-1>config>router>mpls#
```

Modifying MPLS Path Parameters

In order to modify path parameters, the **config>router>mpls>path** context must be shut down first.

The following displays a path configuration example. Refer to the LSP configuration on [page 182](#).

```
A:ALA-1>config>router>mpls# info
#-----
echo "MPLS"
#-----
...
    path "secondary-path"
        hop 1 10.10.0.111  strict
        hop 2 10.10.0.222  strict
        hop 3 10.10.0.123  strict
        no shutdown
    exit
    path "to-NYC"
        hop 1 10.10.10.104  strict
        hop 2 10.10.0.210  strict
        no shutdown
    exit
...
#-----
A:ALA-1>config>router>mpls#
```


Modifying MPLS Static LSP Parameters

In order to modify static LSP parameters, the **config>router>mpls>path** context must be shut down first.

The following displays a static LSP configuration example. Refer to the static LSP configuration on [page 184](#).

```
A:ALA-1>config>router>mpls# info
-----
...
    static-lsp "static-LSP"
        to 10.10.10.234
        push 102704 nexthop 10.10.8.114
        no shutdown
    exit
    no shutdown
-----
A:ALA-1>config>router>mpls#
```

Deleting an MPLS Interface

In order to delete an interface from the MPLS configuration, the interface must be shut down first.

Use the following CLI syntax to delete an interface from the MPLS configuration:

CLI Syntax: `mpls`
 `[no] interface ip-int-name`
 `shutdown`

```
A:ALA-1>config>router>mpls# info
-----
...
  admin-group "green" 15
  admin-group "red" 25
  admin-group "yellow" 20
  interface "system"
  exit
  no shutdown
-----
A:ALA-1>config>router>mpls#
```

RSVP Configuration Management Tasks

This section discusses the following RSVP configuration management tasks:

- [Modifying RSVP Parameters on page 195](#)
 - [Modifying RSVP Message Pacing Parameters on page 196](#)
 - [Deleting an Interface from RSVP on page 196](#)
-

Modifying RSVP Parameters

Only interfaces configured in the MPLS context can be modified in the RSVP context.

The **no rsvp** command deletes this RSVP protocol instance and removes all configuration parameters for this RSVP instance.

The **shutdown** command suspends the execution and maintains the existing configuration.

The following example displays a modified RSVP configuration example:

```
A:ALA-1>config>router>rsvp# info
-----
      keep-multiplier 5
      refresh-time 60
      msg-pacing
        period 400
        max-burst 400
      exit
      interface "system"
      exit
      interface "test1"
        hello-interval 5000
      exit
      no shutdown
-----
A:ALA-1>config>router>rsvp#
```

Modifying RSVP Message Pacing Parameters

RSVP message pacing maintains a count of the messages that were dropped because the output queue for the egress interface was full.

The following example displays command usage to modify RSVP parameters:

The following example displays a modified RSVP message pacing configuration example. Refer to the RSVP message pacing configuration on [page 187](#).

```
A:ALA-1>config>router>rsvp# info
```

```
-----  
      keep-multiplier 5  
      refresh-time 60  
      msg-pacing  
        period 200  
        max-burst 200  
      exit  
      interface "system"  
      exit  
      interface "to-104"  
      exit  
      no shutdown  
-----
```

```
A:ALA-1>config>router>rsvp#
```

Deleting an Interface from RSVP

Interfaces cannot be deleted directly from the RSVP configuration. An interface must have been configured in the MPLS context, which enables it automatically in the RSVP context. The interface must first be deleted from the MPLS context. This removes the association from RSVP.

See [Deleting an MPLS Interface on page 194](#) for information on deleting an MPLS interface.

MPLS/RSVP Command Reference

Command Hierarchies

- [MPLS Commands on page 197](#)
- [MPLS Path Commands on page 204](#)
- [LSP Commands on page 201](#)
- [RSVP Commands on page 205](#)
- [Show Commands on page 207](#)
- [Tools Commands on page 209](#)
- [Clear Commands on page 210](#)
- [Debug Commands on page 211](#)

MPLS Commands

```

config
  — router
    — [no] mpls
      — [no] admin-group-frr
      — auto-bandwidth-multipliers sample-multiplier number1 adjust-multiplier
         number2
      — no auto-bandwidth-multipliers
      — auto-lsp lsp-template template-name { policy peer-prefix-policy [peer-prefix-pol-
         icy...(upto 5 max)] | one-hop }
      — no auto-lsp lsp-template template-name
      — bypass-resignal-timer minutes
      — no bypass-resignal-timer
      — [no] cspf-on-loose-hop
      — dynamic-bypass [enable | disable]
      — exponential-backoff-retry
      — [no] frr-object
      — hold-timer seconds
      — no hold-timer
      — ingress-statistics
        — [no] lsp lsp-name sender ip-address
          — accounting-policy policy-id
          — no accounting-policy
          — [no] collect-stats
          — [no] shutdown
        — [no] p2p-template-lsp rsvp-session-name SessionNameString sender
           sender-address
          — accounting-policy policy-id
          — no accounting-policy
          — [no] collect-stats
  
```

- [no] **max-stats**
- [no] **shutdown**
- [no] **p2mp-template-lsp** *rsvp-session-name* *SessionNameString* *sender* *sender-address*
 - **accounting-policy** *policy-id*
 - **no accounting-policy**
 - [no] **collect-stats**
 - [no] **max-stats**
 - [no] **shutdown**
- [no] **interface** *ip-int-name*
 - [no] **admin-group** *group-name* [*group-name...*(up to 5 max)]
 - **no admin-group**
 - [no] **label-map** *in-label*
 - [no] **pop**
 - [no] **shutdown**
 - **swap** {*out-label* | **implicit-null-label**} *nexthop* *ip-addr*
 - **no swap** {*out-label* | **implicit-null-label**}
 - [no] **mpls-tp-mep**
 - **if-num** *if-num*
 - **no if-num**
 - **if-num-validation** {**enable**|**disable**}
 - [no] **shutdown**
 - [no] **srlg-group** *group-name* [*group-name...*(up to 5 max)]
 - **no srlg-group**
 - **te-metric** *metric*
 - **no te-metric**
- [no] **ldp-over-rsvp** [**include** | **exclude**]
- **least-fill-min-thd** *percent*
- **no least-fill-min-thd**
- **least-fill-reoptim-thd** *percent*
- **no least-fill-reoptim-thd**
- **lsp-init-retry-timeout** *seconds*
- **no lsp-init-retry-timeout**
- [no] **logger-event-bundling**
- **max-bypass-associations** *integer*
- **no max-bypass-associations**
- [no] **mbb-prefer-current-hops**
- **p2p-active-path-fast-retry** *seconds* [*1..10*] *seconds*
- **no p2p-active-path-fast-retry**
- **p2mp-s21-fast-retry** *seconds* [*1..10*] *seconds*
- **no p2mp-s21-fast-retry**
- **preemption-timer** *seconds*
- **no preemption-timer**
- **p2mp-resignal-timer** *minutes*
- **no p2mp-resignal-timer**
- **resignal-timer** *minutes*
- **no resignal-timer**
- [no] **retry-on-igp-overload**
- **secondary-fast-retry-timer** *seconds*
- **no secondary-fast-retry-timer**
- [no] **shutdown**
- [no] **srlg-database**
 - [no] **router-id** *ip*
 - [no] **interface** *ip-addr* **srlg-group** *group-name* [*group-name..*(up to 5 max)]

- [no] **shutdown**
- [no] **srlg-fr** [strict]
- [no] **static-lsp** *lsp-name*
 - **push** {*label* | **implicit-null-label**} **nexthop** *ip-address*
 - **no push** {*out-label* | **implicit-null-label**}
 - [no] **shutdown**
 - **toip-address**
- **static-lsp-fast-retry** *seconds*
- [no] **static-lsp-fast-retry**
- **user-srlg-db** [enable | disable]

MPLS-TP Commands

```
config
  — router
    — [no] mpls
      — [no] mpls-tp
        — global-id global-id
        — no global-id
        — node-id node-id
        — no node-id
        — [no] oam-template name
          — hold-time-down timer
          — no hold-time-down
          — hold-time-up timer
          — no hold-time-up
          — bfd-template name
          — no bfd-template
        — protection-template name
        — no protection-template
          — [no] revertive
          — wait-to-restore interval
          — no wait-to-restore
          — rapid-psc-timer interval
          — no rapid-psc-timer
          — slow-psc-timer interval
          — no slow-psc-timer
        — [no] shutdown
        — tp-tunnel-id-range start-id end-id
        — no tp-tunnel-id-range
        — transit-path path-name
        — no transit-path
          — [no] forward-path
          — in-label in-label out-label out-link if-name [next-hop next-hop]
          — no in-label
          — [no] mep
            — dsmap if-num
            — no dsmap
          — path-id {lsp-num lsp-num | working-path | protect-path [src-global-id src-global-id] src-node-id src-node-id src-tunnel-num src-tunnel-num [dest-global-id dest-global-id] dest-node-id dest-node-id [dest-tunnel-num dest-tunnel-num]}
          — no path-id
          — [no] reverse-path
          — [no] shutdown
```


LSP Commands

```

config
  — router
    — [no] mpls
      — [no] lsp lsp-name [bypass-only | p2mp-lsp | mpls-tp src-tunnel-num]
        — [no] adaptive
        — [no] adspec
        — [no] auto-bandwidth
          — adjust-down percent [bw mbps]
          — no adjust-down
          — adjust-up percent [bw mbps]
          — no adjust-up
          — max-bandwidth mbps
          — no max-bandwidth
          — min-bandwidth mbps
          — no min-bandwidth
          — [no] monitor-bandwidth
          — multipliers sample-multiplier num1 adjust-multiplier num2
          — no multipliers
          — overflow-limit number threshold percent [bw mbps]
          — no overflow-limit
          — underflow-limit number threshold percent [bw mbps]
          — no underflow-limit
        — [no] bgp-shortcut
        — bgp-transport-tunnel include / exclude
        — class-type ct-number
        — no class-type
        — [no] cspf [use-te-metric]
        — dest-global-id dest-global-id
        — no dest-global-id
        — dest-tunnel-number dest-tunnel-number
        — no dest-tunnel-number
        — [no] egress-statistics
          — accounting-policy policy-id
          — no accounting-policy
          — [no] collect-stats
          — [no] shutdown
        — [no] exclude group-name [group-name...(up to 5 max)]
        — [no] exclude-node ip-address
        — fast-reroute frr-method
        — no fast-reroute
          — [no] propagate-admin-group
          — bandwidth rate-in-mbps
          — no bandwidth
          — hop-limit number
          — no hop-limit
          — [no] node-protect
        — from ip-address
        — hop-limit number
        — no hop-limit
        — igp-shortcut [lfa-protect | lfa-only] [relative-metric [offset]]
        — [no] igp-shortcut
        — [no] include group-name [group-name...(up to 5 max)]

```

- **ldp-over-rsvp** [**include** | **exclude**]
- [**no**] **least-fill**
- [**no**] **ldp-over-rsvp** [**include** | **exclude**]
- **main-ct-retry-limit** *number*
- **no main-ct-retry-limit**
- [**no**] **metric** *metric*
- **p2mp-id** *id*
- [**no**] **primary** *path-name*
 - [**no**] **adaptive**
 - **backup-class-type** *ct-number*
 - **no backup-class-type**
 - **bandwidth** *rate-in-mpbs*
 - **no bandwidth**
 - **class-type** *ct-number*
 - **no class-type**
 - [**no**] **exclude** *group-name* [*group-name...*(up to 5 max)]
 - **hop-limit** *number*
 - **no hop-limit**
 - [**no**] **include** *group-name* [*group-name...*(up to 5 max)]
 - **priority** *setup-priority hold-priority*
 - **no priority**
 - [**no**] **record**
 - [**no**] **record-label**
 - [**no**] **shutdown**
- [**no**] **primary-p2mp-instance** *instance-name*
 - [**no**] **adaptive**
 - **bandwidth** *rate-in-mbps*
 - **no bandwidth**
 - [**no**] **exclude** *group-name* [*group-name...*(up to 5 max)]
 - [**no**] **hop-limit**
 - **hop-limit** *number*
 - **no hop-limit**
 - [**no**] **include** *group-name* [*group-name...*(up to 5 max)]
 - [**no**] **record**
 - [**no**] **record-label**
 - [**no**] **s2l-path** *path-name* **to** *ip-address*
 - [**no**] **shutdown**
 - [**no**] **shutdown**
- [**no**] **propagate-admin-group**
- [**no**] **protect-tp-path**
 - **bfd-enable** [**cc** | **cc_cv**]
 - **no bfd-enable**
 - [**no**] **mep**
 - **dsmap** *if-num*
 - **no dsmap**
 - **in-label** *in-label*
 - **no in-label**
 - **lsp-num** *lsp-num*
 - **no lsp-num**
 - [**no**] **mep**
 - **oam-template** *name*
 - **no oam-template**
 - **out-label** *out-label* **out-link** *if-name* [**next-hop** *ip-address*]
 - **no out-label**
 - **protection-template** *name*

- **no protection-template**
- **[no] shutdown**
- **retry-limit** *number*
- **no retry-limit**
- **retry-timer** *seconds*
- **no retry-timer**
- **rsvp-resv-style** [*se* | *ff*]
- **[no] secondary** *path-name*
 - **[no] adaptive**
 - **bandwidth** *rate-in-mbps*
 - **no bandwidth**
 - **class-type** *ct-number*
 - **no class-type**
 - **[no] exclude** *group-name* [*group-name...*(up to 5 max)]
 - **hop-limit** *number*
 - **no hop-limit**
 - **[no] include** *group-name* [*group-name...*(up to 5 max)]
 - **[no] path-preference** *preference-number*
 - **priority** *setup-priority hold-priority*
 - **no priority**
 - **[no] record**
 - **[no] record-label**
 - **[no] shutdown**
 - **[no] srlg**
 - **[no] standby**
- **[no] shutdown**
- **to** [*ip-address* | **node-id** [*a.b.c.d.* | *1...4,294,967,295*]]
- **vprn-auto-bind** [**include** | **exclude**]
- **[no] working-tp-path**
 - **bfd-enable** [*cc* | *cc_cv*]
 - **no bfd-enable**
 - **in-label** *in-label*
 - **no in-label**
 - **lsp-num** *lsp-num*
 - **no lsp-num**
 - **[no] mep**
 - **dsmap** *if-num*
 - **no dsmap**
 - **oam-template** *name*
 - **no oam-template**
 - **out-label** *out-label out-link if-name* [**next-hop** *ip-address*]
 - **no out-label**
 - **[no] shutdown**
- **lsp-template** *template-name* [**p2mp** | **one-hop-p2p** | **mesh-p2p**]
- **no lsp-template** *template-name*
 - **[no] adspec**
 - **[no] auto-bandwidth**
 - **adjust-down** *percent* [*bw mbps*]
 - **no adjust-down**
 - **adjust-up** *percent* [*bw mbps*]
 - **no adjust-up**
 - **fc** *fc-name sampling-weight* *sampling-weight*
 - **no fc**
 - **max-bandwidth** *mbps*
 - **no max-bandwidth**

- **min-bandwidth** *mbps*
- **no min-bandwidth**
- **[no] monitor-bandwidth**
- **multipliers** **sample-multiplier** *num1* **adjust-multiplier** *num2*
- **no multipliers**
- **overflow-limit** *number* **threshold** percent [**bw** *mbps*]
- **no overflow-limit**
- **underflow-limit** *number* **threshold** percent [**bw** *mbps*]
- **no underflow-limit**
- **[no] bandwidth** *rate-in-mbps*
- **[no] cspf** [**use-te-metric**]
- **[no] default-path** *path-name*
- **[no] egress-statistics**
 - **accounting-policy** *policy-id*
 - **no accounting-policy**
 - **[no] collect-stats**
- **[no] exclude-node** *ip-address*
- **fast-reroute** *frr-method*
- **no fast-reroute**
 - **[no] propagate-admin-group**
 - **bandwidth** *rate-in-mbps*
 - **no bandwidth**
 - **hop-limit** *number*
 - **no hop-limit**
 - **[no] node-protect**
- **from** *ip-address*
- **hop-limit** *number*
- **no hop-limit**
- **igp-shortcut** [**lfa-protect** | **lfa-only**] [**relative-metric** [*offset*]]
- **[no] igp-shortcut**
- **[no] include** *group-name* [*group-name...*(up to 5 max)]
- **ldp-over-rsvp** [**include** | **exclude**]
- **[no] least-fill**
- **[no] metric** *metric*
- **[no] propagate-admin-group**
- **[no] record**
- **[no] record-label**
- **retry-limit** *number*
- **no retry-limit**
- **retry-timer** *seconds*
- **no retry-timer**
- **rsvp-resv-style** [**se** | **ff**]
- **vprn-auto-bind** [**include** | **exclude**]

MPLS Path Commands

- ```

config
 — router
 — [no] mpls
 — [no] path path-name
 — hop hop-index ip-address {strict | loose}
 — no hop hop-index
 — [no] shutdown
 — [no] static-lsp lsp-name
 — push label nexthop ip-address

```

- **no push** *out-label*
- **to** *ip-addr*
- **[no] shutdown**

## RSVP Commands

- ```

config
  — router
    — [no] rsvp
      — diffserv-te [mam | rdm]
      — no diffserv-te
        — class-type-bw ct0 %link-bandwidth ct1 %link-bandwidth ct2 %link-
          bandwidth ct3 %link-bandwidth ct4 %link-bandwidth ct5 %link-band-
          width ct6 %link-bandwidth ct7 %link-bandwidth
        — no class-type-bw
        — fc fc-name class-type ct-number
        — no fc fc-name
        — te-class te-class-number class-type ct-number priority priority
        — no te-class te-class-number
      — gr-helper-time max-recovery recovery-interval [1..1800] seconds max-restart
        restart-interval [1..300] seconds
      — no gr-helper-time
      — [no] graceful-shutdown
      — [no] implicit-null-label
      — [no] interface ip-int-name
        — authentication-key [authentication-key | hash-key] [hash | hash2]
        — no authentication-key
        — [no] bfd-enable
        — [no] graceful-shutdown
        — class-type-bw ct0 %link-bandwidth ct1 %link-bandwidth ct2 %link-
          bandwidth ct3 %link-bandwidth ct4 %link-bandwidth ct5 %link-band-
          width ct6 %link-bandwidth ct7 %link-bandwidth
        — no class-type-bw
        — gr-helper [enable | disable]
        — [no] graceful-shutdown
        — hello-interval milli-seconds
        — no hello-interval
        — implicit-null-label [enable | disable]
        — no implicit-null-label
        — [no] refresh-reduction
          — [no] reliable-delivery
        — [no] shutdown
        — subscription percentage
        — no subscription
        — te-up-threshold threshold-level [threshold-level,...(up to 16 max)]
        — no te-up-threshold
        — te-down-threshold threshold-level [threshold-level,...(up to 16 max)]
        — no te-down-threshold
      — keep-multiplier number
      — no keep-multiplier
      — [no] msg-pacing
        — max-burst number
        — no max-burst
        — period milli-seconds
        — no period

```

- **node-id-in-rro** <include|exclude>
- **p2p-merge-point-abort-timer** [*1..65535*] *seconds*
- **no p2p-merge-point-abort-timer**
- **p2mp-merge-point-abort-timer** [*1..65535*] *seconds*
- **no p2mp-merge-point-abort-timer**
- **preemption-timer** *seconds*
- **no preemption-timer**
- **rapid-retransmit-time** *hundred-milliseconds*
- **no rapid-retransmit-time**
- **rapid-retry-limit** *number*
- **no rapid-retry-limit**
- **refresh-reduction-over-bypass** [**enable** | **disable**]
- **refresh-time** *seconds*
- **no refresh-time**
- [**no**] **graceful-shutdown**
- [**no**] **shutdown**
- [**no**] **te-threshold-update**
 - [**no**] **on-cac-failure**
 - **update-timer** *seconds*
 - **no update-timer**
- **te-up-threshold** *threshold-level* [*threshold-level...(up to 16 max)*]
- **no te-up-threshold**
- **te-down-threshold** *threshold-level* [*threshold-level...(up to 16 max)*]
- **no te-down-threshold**

Show Commands

```

show
  — router
    — mpls
      — admin-group group-name
      — bypass-tunnel [to ip-address] [protected-lsp name] [dynamic | manual| p2mp]
        [detail]
      — interface [ip-int-name|ip-address] [label-map label]
      — interface [ip-int-name|ip-address] statistics
      — label start-label [end-label | in-use / owner]
      — label-range
      — lsp [lsp-name] [status {up|down}] [from ip-address| to ip-address] [detail]
      — lsp {transit | terminate} [status {up|down}] [from ip-address | to ip-address | lsp-
name name] [detail]
      — lsp count
      — lsp lsp-name activepath
      — lsp [lsp-name] path [path-name] [status {up | down}] [detail]
      — lsp [lsp-name] path [path-name] mbb
      — lsp-egress-stats
      — lsp-egress-stats lsp-name
      — lsp-ingress-stats ip-address lsp lsp-name
      — lsp-ingress-stats sender-address:lsp-name
      — lsp-template [lsp-template-name] [detail]
      — mpls-tp
        — oam-template
        — protection-template
        — status
        — transit-path [path-name] [detail]
      — p2mp-info [type {originate|transit|terminate}] [s2l-endpoint ip-address]
      — p2mp-lsp [lsp-name] [detail]
      — p2mp-lsp [lsp-name] p2mp-instance [p2mp-instance-name] [mbb]
      — p2mp-lsp [lsp-name] p2mp-instance [p2mp-instance-name] s2l [s2l-name] [to s2l-
to-address]] [status {up | down}] [detail]
      — p2mp-lsp [lsp-name] p2mp-instance [p2mp-instance-name] s2l [s2l-name] [to s2l-
to-address]] mbb
      — srlg-database [router-id ip-address] [interface ip-address]
      — srlg-group [group-name]
      — static-lsp [lsp-name]
      — static-lsp {transit | terminate}
      — static-lsp count
      — statistics-summary
      — status
      — tp-lsp [lsp-name] [status {up | down}] [from ip-address | to ip-address] [detail]
      — tp-lsp [lsp-name] path [protecting | working] [detail]
      — tp-lsp [lsp-name] protection

show
  — router
    — rsvp
      — interface [interface [ip-int-name]] statistics [detail]
      — neighbor [ip-address] [detail]

```

- **session** [session-type] [from ip-address| to ip-address] lsp-name name] [status {up|down}][detail]
- **statistics**
- **status**

Tools Commands

```

tools
  — dump
    — router
      — mpls
        — bypass-tunnel
        — ftn
        — ilm
        — lspinfo
        — memory-usage
        — te-lspinfo [endpoint ip-address] [sender ip-address] [lspid lsp-id] [detail]
          [p2p | p2p-tid tunnel-id]
        — te-lspinfo [endpoint ip-address] [sender ip-address] [lspid lsp-id] [detail]
          [p2p | p2p-tid tunnel-id] { [phops] [nhops] [s2l ip-address] }
      — rsvp
        — psb
        — rsb
    — perform
      — router
        — mpls
          — adjust-autobandwidth [lsp lsp-name [force [bandwidth mbps]]]
          — cspf to ip-address
          — force-switch-path [lsp lsp-name] [path path-name]
          — [no] force-switch-path [lsp lsp-name]
          — plr to ip-addr [from ip-addr] [bandwidth bandwidth] [include-bitmap
            bitmap] [exclude-bitmap bitmap] [hop-limit limit] [exclude-address excl-addr
            [excl-addr...(up to 8 max)]] [use-te-metric] [strict-srlg] [srlg-group
            grp-id...(up to 8 max)] [exclude-node excl-node-id [excl-node-id...(up to 8
            max)]] [skip-interface interface-name] [ds-class-type class-type] [cspf-
            reqtype req-type] [least-fill-min-thd thd] [setup-priority val] [hold-pri-
            ority val]
          — resignal {lsp lsp-name path path-name / delay minutes}
          — resignal {p2mp-lsp p2mp-lsp-name p2mp-instance p2mp-instance-name |
            p2mp-delay p2mp-minutes}
          — resignal-bypass {lsp bypass-lsp-name [force] | delay minutes}
          — switch-path [lsp lsp-name] [path path-name]
          — trap-suppress number-of-traps time-interval
          — update-path {lsp lsp-name path current-path-name new-path new-path-
            name}

```

Router Commands

```

config
  — router
    — [no] igmp
      — [no] tunnel-interface rsvp-p2mp lsp-name [sender ip-address]
    — [no] tunnel-interface rsvp-p2mp lsp-name [sender ip-address]

```

Clear Commands

```
clear
  — router
    — mpls
      — interface [ip-int-name] [statistics]
      — lsp lsp-name
      — lsp-autobandwidth [lsp-name]
      — lsp-egress-stats
      — lsp-egress-stats lsp-name
      — lsp-ingress-stats
      — lsp-ingress-stats ip-address lsp lsp-name
      — lsp-ingress-stats sender-address:lsp-name
    — rsvp
      — interface
      — statistics
```

Debug Commands

```

debug
  — router
    — mpls [lsp lsp-name] [sender source-address] [endpoint endpoint-address] [tunnel-id tunnel-id] [lsp-id lsp-id] [interface ip-int-name]
    — no mpls
      — [no] event
        — all [detail]
        — no all
        — frr [detail]
        — no frr
        — iom [detail]
        — no iom
        — lsp-setup [detail]
        — no lsp-setup
        — mbb [detail]
        — no mbb
        — misc [detail]
        — no misc
        — xc [detail]
        — no xc
      — rsvp [lsp lsp-name] [sender source-address] [endpoint endpoint-address] [tunnel-id tunnel-id] [lsp-id lsp-id] [interface ip-int-name]
      — no rsvp
        — [no] event
          — all [detail]
          — no all
          — misc [detail]
          — no misc
          — nbr [detail]
          — no nbr
          — path [detail]
          — no path
          — resv [detail]
          — no resv
          — te-threshold-update
          — no te-threshold-update
        — [no] packet
          — all [detail]
          — no all
          — hello [detail]
          — no hello
          — path [detail]
          — no path
          — patherr [detail]
          — no patherr
          — pathtear [detail]
          — no pathtear
          — resv [detail]
          — no resv
          — resvrr [detail]
          — no resvrr
          — resvtear [detail]
          — no resvtear

```

MPLS Configuration Commands

Generic Commands

shutdown

Syntax	[no] shutdown
Context	config>router>mpls config>router>mpls>interface config>router>mpls>lsp>primary config>router>mpls>lsp>secondary
Description	<p>This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.</p> <p>MPLS is not enabled by default and must be explicitly enabled (no shutdown).</p> <p>The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.</p> <p>The no form of this command places the entity into an administratively enabled state.</p>
Default	no shutdown

MPLS Commands

mpls

Syntax	[no] mpls
Context	config>router
Description	<p>This command enables the context to configure MPLS parameters. MPLS is not enabled by default and must be explicitly enabled (no shutdown). The shutdown command administratively disables MPLS.</p> <p>The no form of this command deletes this MPLS protocol instance; this will remove all configuration parameters for this MPLS instance.</p> <p>MPLS must be shutdown and all SDP bindings to LSPs removed before the MPLS instance can be deleted. If MPLS is not shutdown, when the no mpls command is executed, a warning message on the console displays indicating that MPLS is still administratively up.</p>

accounting-policy

Syntax	accounting-policy <i>acct-policy-id</i> no accounting-policy
Context	config>router>mpls>ingr-stats config>router>mpls>lsp>egr-stats config>router>mpls>lsp>template>egr-stats
Description	<p>This command associates an accounting policy to the MPLS instance.</p> <p>An accounting policy must be defined before it can be associated else an error message is generated.</p> <p>The no form of this command removes the accounting policy association.</p>
Default	none
Parameters	<i>acct-policy-id</i> — Enter the accounting <i>policy-id</i> as configured in the config>log>accounting-policy context.
Values	1 — 99

collect-stats

Syntax	[no] collect-stats
Context	config>router>mpls>ingr-stats config>router>mpls>lsp>egr-stats

```
config>router>mpls>lsp-template>egr-stats
```

Description	This command enables accounting and statistical data collection. When applying accounting policies the data, by default, is collected in the appropriate records and written to the designated billing file. When the no collect-stats command is issued the statistics are still accumulated by the forwarding engine. However, the CPU will not obtain the results and write them to the billing file. If a subsequent collect-stats command is issued then the counters written to the billing file include all the traffic while the no collect-stats command was in effect.
Default	collect-stats

max-stats

Syntax	[no] max-stats
Context	config>router>mpls>ingr-stats config>router>mpls>lsp>egr-stats config>router>mpls>lsp-template>egr-stats
Description	This command enables accounting and statistical data collection. When applying accounting policies the data, by default, is collected in the appropriate records and written to the designated billing file. When the no max-stats command is issued the statistics are still accumulated by the forwarding engine. However, the CPU will not obtain the results and write them to the billing file. If a subsequent max-stats command is issued then the counters written to the billing file include all the traffic while the no max-stats command was in effect.
Default	max-stats

dynamic-bypass

Syntax	dynamic-bypass [enable disable] no dynamic-bypass
Context	config>router>mpls
Description	This command disables the creation of dynamic bypass LSPs in FRR. One or more manual bypass LSPs must be configured to protect the primary LSP path at the PLR nodes.
Default	enable

egress-statistics

Syntax	[no] egress-statistics
Context	config>router>mpls>lsp config>router>mpls>lsp-template
Description	<p>This command configures statistics in the egress data path of an originating LSP at a head-end node. The user must execute the no shutdown for this command to effectively enable statistics.</p> <p>The same set of counters is updated for packets forwarded over any path of the LSP and over the lifetime of the LSP. In steady state, the counters are updated for packets forwarded over the active path of the LSP. The active path can be the primary path, one of the secondary paths, the FRR detour path, or the FRR bypass path when the head-end node is also the PLR.</p> <p>LSP statistics are not collected on a dynamic or a static bypass tunnel itself. LSP egress statistics are also not collected if the head-end node is also the Penultimate-Popping Hop (PHP) node for a single-hop LSP using an implicit null label.</p> <p>When a hierarchy of LSPs is in use, statistics collection on the outermost label corresponding to the tunneling LSP and on the inner labels, corresponding to the tunneled LSPs, are mutually exclusive. A consequence of this is that when the user enables statistics collection on an RSVP LSP which is also used for tunneling LDP FECs with the LDP over RSVP feature, then statistics will be collected on the RSVP LSP only. There will be no statistics collected from an LDP FEC tunneled over this RSVP LSP regardless if the user enabled statistics collection on this FEC. When, the user disables statistics collection on the RSVP LSP, then statistics collection, if enabled, will be performed on a tunneled LDP FEC.</p> <p>The no form of this command disables the statistics in the egress data path and removes the accounting policy association from the RSVP LSP.</p>
Default	no egress-statistics

exponential-backoff-retry

Syntax	exponential-backoff-retry no exponential-backoff-retry
Context	configure>router>mpls
Description	<p>This command enables the use of an exponential back-off timer when re-trying an LSP. When an LSP path establishment attempt fails, the path is put into retry procedures and a new attempt will be performed at the expiry of the user-configurable retry timer (config>router>mpls>lsp>retry-timer). By default, the retry time is constant for every attempt. The exponential back-off timer procedures will double the value of the user configured retry timer value at every failure of the attempt to adjust to the potential network congestion that caused the failure. An LSP establishment fails if no Resv message was received and the Path message retry timer expired or a PathErr message was received before the timer expired.</p>

admin-group-frr

Syntax	[no] admin-group-frr
Context	config>router>mpls
Description	<p>This command enables the use of the admin-group constraints in the association of a manual or dynamic bypass LSP with the primary LSP path at a Point-of-Local Repair (PLR) node.</p> <p>When this command is enabled, each PLR node reads the admin-group constraints in the FAST_REROUTE object in the Path message of the LSP primary path. If the FAST_REROUTE object is not included in the Path message, then the PLR will read the admin-group constraints from the Session Attribute object in the Path message.</p> <p>If the PLR is also the ingress LER for the LSP primary path, then it just uses the admin-group constraint from the LSP and/or path level configurations.</p> <p>The PLR node then uses the admin-group constraints along with other constraints, such as hop-limit and SRLG, to select a manual or dynamic bypass among those that are already in use.</p> <p>If none of the manual or dynamic bypass LSP satisfies the admin-group constraints, and/or the other constraints, the PLR node will request CSPF for a path that merges the closest to the protected link or node and that includes or excludes the specified admin-group IDs.</p> <p>If the user changes the configuration of the above command, it will not have any effect on existing bypass associations. The change will only apply to new attempts to find a valid bypass.</p> <p>The no form of this command disables the use of administrative group constraints on a FRR backup LSP at a PLR node.</p>
Default	no frr-admin-group

frr-object

Syntax	[no] frr-object
Context	config>router>mpls
Description	<p>This command specifies whether fast reroute for LSPs using the facility bypass method is signalled with or without the fast reroute object using the one-to-one keyword. The value is ignored if fast reroute is disabled for the LSP or if the LSP is using one-to-one Backup.</p>
Default	frr-object — The value is by default inherited by all LSPs.

hold-timer

Syntax	hold-timer <i>seconds</i> no hold-timer
Context	config>router>mpls
Description	This command specifies the amount of time that the ingress node holds before programming its data plane and declaring the LSP up to the service module. This occurs anytime the ingress node brings up an LSP path or switches traffic from a working path to another working path of the same LSP. The no form of the command reverts the hold-timer to the default value.
Parameters	<i>seconds</i> — Specifies the time, in seconds, for which the ingress node holds before programming its data plane and declaring the LSP up to the service module. Values 0 — 10
Default	1 second

ingress-statistics

Syntax	ingress-statistics
Context	config>router>mpls
Description	This command provides the context for the user to enter the LSP names for the purpose of enabling ingress data path statistics at the terminating node of the LSP, for example, egress LER.
Default	none

least-fill-min-thd

Syntax	least-fill-min-thd <i>percent</i> no least-fill-min-thd
Context	config>router>mpls
Description	This parameter is used in the least-fill path selection process. When comparing the percentage of least available link bandwidth across the sorted paths, whenever two percentages differ by less than the value configured as the least-fill-min-thresh, CSPF will consider them equal and will apply a random number generator to select the path among these paths The no form of the command resets this parameter to its default value.
Default	5
Parameters	<i>percentage</i> — Specifies the least fill minimum threshold value as a percentage. Values 1 — 100%

least-fill-reoptim-thd

Syntax	least-fill-reoptim-thd <i>percent</i> no least-fill-reoptim-thd
Context	config>router>mpls
Description	<p>This parameter is used in the least-fill path selection method. During a timer-based re-signaling of an LSP path which has the least-fill option enabled, CSPF will first update the least-available bandwidth figure for the current path of this LSP. It then applies the least-fill path selection method to select a new path for this LSP. If the new computed path has the same cost as the current path, it will compare the least-available bandwidth figures of the two paths and if the difference exceeds the user configured optimization threshold, MPLS will generate a trap to indicate that a better least-fill path is available for this LSP. This trap can be used by an external SNMP based device to trigger a manual re-signaling of the LSP path since the timer-based re-signaling will not re-signal the path in this case. MPLS will generate a path update trap at the first MBB event which results in the re-signaling of the LSP path. This should clear the eligibility status of the path at the SNMP device.</p> <p>The no form of this command resets this parameter to its default value.</p>
Default	10
Parameters	<i>percentage</i> — Specifies the least fill reoptimization threshold value as a percentage.
	Values 1 — 100%

lsp

Syntax	[no] lsp <i>lsp-name sender sender-address</i>
Context	config>router>mpls>ingress-statistics
Description	<p>This command configures statistics in the ingress data path of a terminating RSVP LSP at an egress LER. The LSP name must correspond to the name configured by the operator at the ingress LER. It must not contain the special character “:” which is used as a field separator by the ingress LER for encoding the LSP and path names into the RSVP session name field in the session_ attribute object. The operator must execute the no shutdown for this command to effectively enable statistics.</p> <p>The same set of counters is updated for packets received over any path of this LSP and over the lifetime of the LSP. In steady-state, the counters are updated for packets received over the active path of the LSP. The active path can be the primary path, one of the secondary paths, the FRR detour path, or the FRR bypass path when the tail-end node is also the MP.</p> <p>When a hierarchy of LSPs is in use, statistics collection on the outermost label corresponding to the tunneling LSP and on the inner labels, corresponding to the tunneled LSPs are mutually exclusive. A consequence of this is that when the operator enables statistics collection on an RSVP LSP which is also used for tunneling LDP FECs with the LDP over RSVP feature, then statistics will be collected on the RSVP LSP only. There will be no statistics collected for an LDP FEC tunneled over this RSVP LSP and also egressing on the same node regardless if the operator enabled statistics collection on this FEC. When, the operator disables statistics collection on the RSVP LSP, then statistics collection, if enabled, will be performed on a tunneled LDP FEC.</p>

The operator can enable statistics collection on a manual bypass terminating on the egress LER. However all LSPs which primary path is protected by the manual bypass will not collect statistics when they activate forwarding over the manual bypass. When, the operator disables statistics collection on the manual bypass LSP, then statistics collection on the protected LSP, if enabled, will continue when the bypass LSP is activated.

The **no** form of this command disables statistics for this RSVP LSP in the ingress data path and removes the accounting policy association from the LSP.

Default none

Parameters **sender-address** *ip-address* — A string of 15 characters representing the IP address of the ingress LER for the LSP.

lsp-name — A string of up to 32 characters identifying the LSP name as configured at the ingress LER.

p2p-template-lsp

Syntax [**no**] **p2p-template-lsp** **rsvp-session-name** *SessionNameString* **sender** *sender-address*

Context config>router>mpls>ingress-stats

Description This command configures an ingress statistics context matching on the RSVP session name for a RSVP P2P auto-LSP at the egress LER.

When the ingress LER signals the path of a template based **one-hop-p2p** or **mesh-p2p auto-lsp**, it includes the name of the LSP and that of the path in the Session Name field of the Session Attribute object in the Path message. The encoding is as follows:

Session Name: *lsp-name::path-name*, where *lsp-name* component is encoded as follows:

1. P2MP LSP via user configuration for L3 multicast in global routing instance: “LspNameFrom-Config”
2. P2MP LSP as I-PMSI or S-PMSI in L3 mVPN: *templateName-SvcId-mTTmIndex*
3. P2MP LSP as I-PMSI in VPLS/B-VPLS: *templateName-SvcId-mTTmIndex*.

The ingress statistics CLI configuration allows the user to match either on the exact name of the P2P auto-LSP or on a context that matches on the template name and the destination of the LSP at the ingress LER.

When the matching is performed on a context, the user must enter the RSVP session name string in the format “*templateName-svcId*” to include the LSP template name as well as the mVPN VPLS/B-VPLS service ID as configured at the ingress LER. In this case, one or more P2MP LSP instances signaled by the same ingress LER could be associated with the ingress statistics configuration. In this case, the user is provided with CLI parameter **max-stats** to limit the maximum number of stat indices which can be assigned to this context. If the context matches more than this value, the additional request for stat indices from this context will be rejected.

Note the following rules when configuring an ingress statistics context based on template matching:

- **max-stats**, once allocated, can be increased but not decreased unless the entire ingress statistics context matching a template name is deleted.
- In order to delete ingress statistics context matching a template name, a shutdown is required.

- An accounting policy cannot be configured or de-configured until the ingress statistics context matching a template name is shut down.
- After deleting an accounting policy from an ingress statistics context matching a template name, the policy is not removed from the log until a **no shut** is performed on the ingress statistics context.

If there are no stat indices available at the time the session of the P2P LSP matching a template context is signaled and the session state installed by the egress LER, no stats are allocated to the session.

Furthermore, the assignment of stat indices to the LSP names that match the context will also be not deterministic. The latter is due to the fact that a stat index is assigned and released following the dynamics of the LSP creation or deletion by the ingress LER. For example, a multicast stream crosses the rate threshold and is moved to a newly signaled S-PMSI dedicated to this stream. Later on, the same stream crosses the threshold downwards and is moved back to the shared I-PMSI and the P2MP LSP corresponding to the S-PMSI is deleted by the ingress LER.

The **no** form deletes the ingress statistics context matching on the RSVP session name.

- Parameters**
- rsvp-session-name** *SessionNameString* — Specifies the name of the RSVP P2MP session in the format of “templateName-svcId”. This field can hold up to 64 characters.
- sender** *sender-address* — Specifies a string of 15 characters representing the IP address of the ingress LER for the LSP.

p2mp-template-lsp

- Syntax** [no] **p2mp-template-lsp rsvp-session-name** *SessionNameString* **sender** *sender-address*
- Context** config>router>mpls>ingress-stats
- Description** This command configures an ingress statistics context matching on the RSVP session name for a RSVP P2MP LSP at the egress LER.

When the ingress LER signals the path of the S2L sub-LSP, it includes the name of the LSP and that of the path in the Session Name field of the Session Attribute object in the Path message. The encoding is as follows:

Session Name: <lsp-name::path-name>, where lsp-name component is encoded as follows:

- P2MP LSP via user configuration for L3 multicast in global routing instance: “LspNameFrom-Config”
- P2MP LSP as I-PMSI or S-PMSI in L3 mVPN: templateName-SvcId-mTTmIndex
- P2MP LSP as I-PMSI in VPLS/B-VPLS: templateName-SvcId-mTTmIndex

The ingress statistics CLI configuration allows the user to match either on the exact name of the P2MP LSP as configured at the ingress LER or on a context that matches on the template name and the service-id as configured at the ingress LER.

When the matching is performed on a context, the user must enter the RSVP session name string in the format “templateName-svcId” to include the LSP template name as well as the mVPN VPLS/B-VPLS service ID as configured at the ingress LER. In this case, one or more P2MP LSP instances signaled by the same ingress LER could be associated with the ingress statistics configuration and the user is provided with CLI parameter max-stats to limit the maximum number of stat indices that can

be assigned to this context. If the context matches more than this value, the additional request for stat indices from this context will be rejected. A background tasks monitors the ingress statistics templates which have one or more matching LSP instances with unassigned stat index and assigns one to them as they are freed.

Note the following rules when configuring an ingress statistics context based on template matching:

- max-stats, once allocated, can be increased but not decreased unless the entire ingress statistics context matching a template name is deleted.
- In order to delete ingress statistics context matching a template name, a shutdown is required.
- An accounting policy cannot be configured or de-configured until the ingress statistics context matching a template name is shut down.
- After deleting an accounting policy from an ingress statistics context matching a template name, the policy is not removed from the log until a “no shut” is performed on the ingress statistics context.

If there are no stat indices available at the time the session of the P2MP LSP matching a template context is signaled and the session state installed by the egress LER, no stats are allocated to the session.

Furthermore, the assignment of stat indices to the LSP names that match the context will also be not deterministic. The latter is due to the fact that a stat index is assigned and released following the dynamics of the LSP creation or deletion by the ingress LER. For example, a multicast stream crosses the rate threshold and is moved to a newly signaled S-PMSI dedicated to this stream. Later on, the same steam crosses the threshold downwards and is moved back to the shared I-PMSI and the P2MP LSP corresponding to the S-PMSI is deleted by the ingress LER.

The **no** form deletes the ingress statistics context matching on the RSVP session name.

- Parameters**
- rsvp-session-name** *SessionNameString* — Specifies the name of the RSVP P2MP session in the format of “templateName-svcId”. This field can hold up to 64 characters.
 - sender** *sender-address* — Specifies a string of 15 characters representing the IP address of the ingress LER for the LSP.

logger-event-bundling

- Syntax** **[no] logger-event-bundling**
- Context** configure>router>mpls
- Description** This feature merges two of the most commonly generated MPLS traps, vRtrMplsXCCreate and vRtrMplsXCDelete, which can be generated at both LER and LSR into a new specific trap vRtrMplsSessionsModified. In addition, this feature will perform bundling of traps of multiple RSVP sessions, i.e., LSPs, into this new specific trap.

The intent is to provide a tool for the user to minimize trap generation in an MPLS network. Note that the MPLS trap throttling will not be applied to this new trap.

The **no** version of this command disables the merging and bundling of the above MPLS traps.

lsp-template

Syntax	lsp-template template-name [p2mp one-hop-p2p mesh-p2p] no lsp-template template-name
Context	config>router>mpls
Description	This command creates a template construct that can be referenced by client application where dynamic LSP creation is required. The LSP template type p2mp , one-hop-p2p , or mesh-p2p is mandatory. The no form of command deletes LSP template. LSP template cannot be deleted if a client application is using it.
Parameters	<i>lsp-template-name</i> — Specifies the name of the LSP template. Any LSP template name and LSP name must not be the same. p2mp one-hop-p2p mesh-p2p — Identifies the type of the LSP this template will signal.

lsp-init-retry-timeout

Syntax	lsp-init-retry-timeout seconds no lsp-init-retry-timeout				
Context	config>router>mpls				
Description	This command configures the initial LSP path retry-timer. The new LSP path initial retry-timer is used instead of the retry-timer to abort the retry cycle when no RESV is received. The retry-timer will govern exclusively the time between two retry cycles and to handle retrying of an LSP path in a failure case with PATH errors or RESVTear. The intent is that the user can now control how many refreshes of the pending PATH state can be performed before starting a new retry-cycle with a new LSP-id. This is all done without affecting the ability to react faster to failures of the LSP path, which will continue to be governed by the retry-timer. The no form of this command returns the timer to the default value.				
Parameters	<i>seconds</i> — Specifies the value, in seconds, used as the fast retry timer for a secondary path. <table> <tr> <td>Values</td> <td>10—600</td> </tr> <tr> <td>Default</td> <td>30</td> </tr> </table>	Values	10—600	Default	30
Values	10—600				
Default	30				

lsp-template

Syntax	lsp-template <i>template-name</i> [p2mp one-hop-p2p mesh-p2p] no lsp-template <i>template-name</i>
Context	config>router>mpls
Description	<p>This command creates a template construct that can be referenced by client application where dynamic LSP creation is required. The LSP template type p2mp, one-hop-p2p, or mesh-p2p is mandatory.</p> <p>The no form of command deletes LSP template. LSP template cannot be deleted if a client application is using it.</p>
Parameters	<p><i>lsp-template-name</i> — Specifies the name to identify LSP template. ANY LSP template name and LSP name must not be the same.</p> <p>p2mp one-hop-p2p mesh-p2p — Identifies the type of the LSP this template will signal.</p>

propagate-admin-group

Syntax	[no] propagate-admin-group
Context	config>router>mpls>lsp>fast-reroute config>router>mpls>lsp-template>fast-reroute
Description	<p>The command enables the signaling of the primary LSP path admin-group constraints in the FRR object at the ingress.</p> <p>When this command is executed, the admin-group constraints configured in the context of the P2P LSP primary path, or the ones configured in the context of the LSP and inherited by the primary path, are copied into the FAST_REROUTE object. The admin-group constraints are copied into the 'include-any' or 'exclude-any' fields.</p> <p>The ingress LER thus propagates these constraints to the downstream nodes during the signaling of the LSP to allow them to include the admin-group constraints in the selection of the FRR backup LSP for protecting the LSP primary path.</p> <p>The ingress LER will insert the FAST_REROUTE object by default in a primary LSP path message. If the user disables the object using the following command, the admin-group constraints will not be propagated: configure>router>mpls>no frr-object .</p> <p>Note that the same admin-group constraints can be copied into the Session Attribute object. They are intended for the use of an LSR, typically an ABR, to expand the ERO of an inter-area LSP path. They are also used by any LSR node in the path of a CSPF or non-CSPF LSP to check the admin-group constraints against the ERO regardless if the hop is strict or loose. These are governed strictly by the command:</p> <p>configure>router>mpls>lsp>propagate-admin-group</p> <p>In other words, the user may decide to copy the primary path admin-group constraints into the FAST_REROUTE object only, or into the Session Attribute object only, or into both. Note, however, that the PLR rules for processing the admin-group constraints can make use of either of the two object admin-group constraints.</p>

This feature is supported with the following LSP types and in both intra-area and inter-area TE where applicable:

- Primary path of a RSVP P2P LSP.
- S2L path of an RSVP P2MP LSP instance
- LSP template for an S2L path of an RSVP P2MP LSP instance.

The **no** form of this command disables the signaling of administrative group constraints in the FRR object.

Default no propagate-admin-group

max-bypass-associations

Syntax **max-bypass-associations** *integer*
no max-bypass-associations

Context config>router>mpls

Description This command allows the user to set a maximum number of LSP primary path associations with each manual or dynamic bypass LSP that is created in the system.

By default, a Point of Local Repair (PLR) node will associate a maximum of 1000 primary LSP paths with a given bypass before using the next available manual bypass or signaling a new dynamic bypass.

Note that a new bypass LSP may need to be signaled if the constraint of a given primary LSP path is not met by an existing bypass LSP even if the max-bypass-associations for this bypass LSP has not been reached.

The **no** form of the command re-instates the default value of this parameter.

Default no max-bypass-associations

Values 1—131,072

mbb-prefer-current-hops

Syntax [**no**] **mbb-prefer-current-hops**

Context config>router>mpls

Description This command implements a new option in the CSPF path computation during a Make-Before-Break (MBB) procedure of an RSVP LSP.

When MPLS performs an MBB for the primary or secondary path of a P2P LSP, or the S2L path of a P2MP LSP, and the new **mbb-prefer-current-hops** option is enabled in MPLS context, CSPF will select a path, among equal-cost candidate paths, with the most overlapping links with the current path. Normally, CSPF selects the path randomly.

The procedures of the new MBB CSPF path selection apply to LSP without the least-fill option enabled. If the least-fill rule results in a different path, the LSP path will be moved though. Users can still favor stability over least-fill condition by applying a larger value to the parameter **least-fill-min-**

thd under the MPLS context such that a path will only be moved when the difference of the least-available bandwidth becomes significant enough between the most used links in the equal cost paths. If that difference is not significant enough, CSPF will select the path with the most overlapping links instead of selecting a path randomly.

The procedures when the new **mbb-prefer-current-hops** option is enabled apply to all MBB types. Thus, it applies to the auto-bandwidth MBB, the configuration change MBB, the soft pre-emption MBB, the TE graceful shutdown MBB, the delayed retry MBB (for SRLG secondary LSP path), the path change MBB, the timer resignal MBB, and the manual resignal MBB.

During the FRR global revertive MBB, CSPF selects a random link among the ones available between the PLR node and the Merge Point node, including the failed link if it has restored in the meantime. These links cannot be checked for overlap with the current path.

The TE graceful shutdown MBB will still avoid the link or node that is in maintenance and the soft pre-emption MBB will still avoid the link that is overbooked.

For an inter-area LSP, this feature applies to the subset of the path from the ingress LER to the exit ABR.

The procedures of this feature are not applied to a zero bandwidth CSFP LSP, including an auto-bandwidth CSPF LSP while its operational bandwidth is zero, and to a non-CSPF LSP.

resignal-timer

Syntax	resignal-timer <i>minutes</i> no resignal-timer
Context	config>router>mpls
Description	<p>This command specifies the value for the LSP resignal timer. The resignal timer is the time, in minutes, the software waits before attempting to resignal the LSPs.</p> <p>When the resignal timer expires, if the new computed path for an LSP has a better metric than the current recorded hop list, an attempt is made to resignal that LSP using the make-before-break mechanism. If the attempt to resignal an LSP fails, the LSP will continue to use the existing path and a resignal will be attempted the next time the timer expires.</p> <p>The no form of the command disables timer-based LSP resignalling.</p>
Default	no resignal-timer
Parameters	<i>minutes</i> — The time the software waits before attempting to resignal the LSPs.
Values	30 — 10080

retry-on-igp-overload

Syntax	[no] retry-on-igp-overload
Context	config>router>mpls
Description	This command enables tearing down LSPs when IGP is in overload state.

secondary-fast-retry-timer

Syntax	secondary-fast-retry-timer <i>seconds</i> no secondary-fast-retry-timer
Context	config>router>mpls
Description	<p>This command specifies the value used as the fast retry timer for a secondary path. If the first attempt to set up a secondary path fails due to a path error, the fast retry timer will be started for the secondary path so that the path can be retried sooner. If the next attempt also fails, further retries for the path will use the configured value for LSP retry timer.</p> <p>If retry-timer for the LSP is configured to be less than the MPLS secondary-fast-retry-timer, all retries for the secondary path will use the LSP retry-timer.</p> <p>The no form of the command reverts to the default.</p>
Default	no secondary-fast-retry-timer
Parameters	<i>seconds</i> — specifies the value, in seconds, used as the fast retry timer for a secondary path
Values	1 — 10

srlg-frr

Syntax	srlg-frr [strict] no srlg-frr
Context	config>router>mpls
Description	<p>This command enables the use of the Shared Risk Loss Group (SRLG) constraint in the computation of FRR bypass or detour to be associated with any primary LSP path on this system.</p> <p>When this option is enabled, CSPF includes the SRLG constraint in the computation of a FRR detour or bypass for protecting the primary LSP path.</p> <p>CSPF prunes all links with interfaces which belong to the same SRLG as the interface which is being protected, i.e., the outgoing interface at the PLR the primary path is using. If one or more paths are found, the MPLS/RSVP task will select one based on best cost and will signal the bypass/detour. If not and the user included the strict option, the bypass/detour is not setup and the MPLS/RSVP task will keep retrying the request to CSPF. Otherwise, if a path exists which meets the other TE constraints, other than the SRLG one, the bypass/detour is setup.</p> <p>A bypass or a detour LSP path is not guaranteed to be SRLG disjoint from the primary path. This is because only the SRLG constraint of the outgoing interface at the PLR the primary path is using is checked.</p> <p>When the MPLS/RSVP task is searching for a SRLG bypass tunnel to associate with the primary path of the protected LSP, it will first check if any configured manual bypass LSP with CSPF enabled satisfies the SLRG constraints. The MPLS/RSVP skips any non-CSPF bypass LSP in the search as there is no ERO returned to check the SLRG constraint. If no path is found, it will check if an existing dynamic bypass LSP satisfies the SLRG and other primary path constraints. If not, then it will make a request to CSPF.</p>

Once the primary path of the LSP is set up and is operationally up, any subsequent changes to the SRLG group membership of an interface the primary path is using would not be considered by the MPLS/RSVP task at the PLR for bypass/detour association until the next opportunity the primary path is re-signaled. The path may be re-signaled due to a failure or to a make-before break operation. Make-before break occurs as a result of a global revertive operation, a timer based or manual re-optimization of the LSP path, or a user change to any of the path constraints.

Once the bypass or detour path is setup and is operationally UP, any subsequent changes to the SRLG group membership of an interface the bypass/detour path is using would not be considered by the MPLS/RSVP task at the PLR until the next opportunity the association with the primary LSP path is re-checked. The association is re-checked if the bypass path is re-optimized. Detour paths are not re-optimized and are re-signaled if the primary path is down.

Enabling or disabling srlg-frr only takes effect at the next opportunity the LSP paths are resignaled. The user can wait for the resignal timer to expire or can cause the paths to be resignaled immediately by executing at the ingress LER the **tools perform router mpls resignal** command. Note that in order to force the dynamic bypass LSP to be resignaled using the SRLG constraint of the primary paths it is associated with, it is recommend to first disable dynamic bypass LSPs on the system using the “configure router mpls dynamicbypass” command, then manually resignal the LSP paths using the above tools perform command finally re-enable dynamic bypass LSPs on the system. Before performing this procedure, the user must ensure that no dynamic bypass LSP on the node is active to avoid causing the primary LSP path to go down.

An RSVP interface can belong to a maximum of 64 SRLG groups. The user configures the SRLG groups using the command **config>router>mpls>srlg-group**. The user configures the SRLG groups an RSVP interface belongs to using the **srlg-group** command in the **config>router>mpls>interface** context.

The **no** form of the command reverts to the default value.

Default	no srlg-frr
Parameters	strict — Specifies the name of the SRLG group within a virtual router instance.
	Values no slr-frr (default) srlg-frr (non-strict) srlg-frr strict (strict)

srlg-group

Syntax	[no] srlg-group <i>group-name</i> [<i>group-name...</i> (up to 5 max)] no srlg-group
Context	config>router>interface>if-attribute config>service>ies>interface>if-attribute config>service>vprn>interface>if-attribute config>router>mpls>interface
Description	This command configures the SRLG membership of an interface. The user can apply SRLGs to an IES, VPRN, network IP, or MPLS interface. An interface can belong to up to 64 SRLG groups. However, each single operation of the srlg-group command allows a maximum of five (5) groups to be specified at a time. Once an SRLG group is bound to one or more interface, its value cannot be changed until all bindings are removed.

The configured SRLG membership will be applied in all levels/areas the interface is participating in. The same interface cannot have different memberships in different levels/areas.

It should be noted that only the SRLGs bound to an MPLS interface are advertised in TE link TLVs and sub-TLVs when the traffic-engineering option is enabled in IS-IS or OSPF. IES and VPRN interfaces do not have their attributes advertised in TE TLVs.

The no form of this command deletes one or more of the SRLG memberships of an interface.

The user can also delete all memberships of an interface by not specifying a group name.

Default no srlg-group

Parameters *group-name* — Specifies the name of the group, up to 32 characters. The association of group name and value should be unique within an IP/MPLS domain.

user-srlg-db

Syntax **user-srlg-db [enable | disable]**

Context config>router>mpls

Description This command enables the use of CSPF by the user SRLG database. When the MPLS module makes a request to CSPF for the computation of an SRLG secondary path, CSPF will query the local SRLG and compute a path after pruning links that are members of the SRLG IDs of the associated primary path. When MPLS makes a request to CSPF for an FRR bypass or detour path to associate with the primary path, CSPF queries the user SRLG database and computes a path after pruning links that are members of the SRLG IDs of the PLR outgoing interface.

If an interface was not entered into the user SRLG database, it is assumed that it does not have any SRLG membership. CSPF will not query the TE database for IGP advertised interface SRLG information.

The disable keyword disables the use of the user SRLG database. CSPF will then resume queries into the TE database for SRLG membership information. The user SRLG database is maintained.

Default user-srlg-db disable

srlg-database

Syntax **[no] srlg-database**

Context config>router>mpls

Description This command provides the context for the user to enter manually the link members of SRLG groups for the entire network at any node that needs to signal LSP paths (for example, a head-end node).

The **no** form of the command deletes the entire SRLG database. CSPF will assume all interfaces have no SRLG membership association if the database was not disabled with the command **config>router>mpls>user-srlg-db disable**.

router-id

Syntax	[no] router-id <i>ip</i>
Context	config>router>mpls>srlg-database
Description	<p>This command provides the context for the user to manually enter the link members of SRLG groups for a specific router in the network. The user must also use this command to enter the local interface SRLG membership into the user SRLG database. Use by CSPF of all interface SRLG membership information of a specific router ID may be temporarily disabled by shutting down the node. If this occurs, CSPF will assume these interfaces have no SRLG membership association.</p> <p>The no form of this command will delete all interface entries under the router ID.</p>
Parameters	<i>ip-address</i> — Specifies the router ID for this system. This must be the router ID configured under the base router instance, the base OSPF instance or the base IS-IS instance.

interface

Syntax	interface <i>ip-address</i> srlg-group <i>group-name</i> [<i>group-name...</i> (up to 5 max)] no interface <i>ip-address</i> [srlg-group <i>group-name...</i> (up to 5 max)]
Context	config>router>mpls>srlg-database>router-id
Description	<p>This command allows the operator to manually enter the SRLG membership information for any link in the network, including links on this node, into the user SRLG database.</p> <p>An interface can be associated with up to 5 SRLG groups for each execution of this command. The operator can associate an interface with up to 64 SRLG groups by executing the command multiple times.</p> <p>CSPF will not use entered SRLG membership if an interface is not validated as part of a router ID in the routing table.</p> <p>The no form of the command deletes a specific interface entry in this user SRLG database. The group-name must already exist in the config>router>mpls>srlg-group context.</p>
Default	none
Parameters	<i>ip-int-name</i> — The name of the network IP interface. An interface name cannot be in the form of an IP address. srlg-group <i>group-name</i> — Specifies the SRLG group name. Up to 1024 group names can be defined in the config>router>mpls context. The SRLG group names must be identical across all routers in a single domain.

MPLS Interface Commands

interface

Syntax	[no] interface <i>ip-int-name</i>
Context	config>router>mpls
Description	<p>This command specifies MPLS protocol support on an IP interface. No MPLS commands are executed on an IP interface where MPLS is not enabled. An MPLS interface must be explicitly enabled (no shutdown).</p> <p>The no form of this command deletes all MPLS commands such as label-map which are defined under the interface. The MPLS interface must be shutdown first in order to delete the interface definition. If the interface is not shutdown, the no interface <i>ip-int-name</i> command does nothing except issue a warning message on the console indicating that the interface is administratively up.</p>
Default	shutdown
Parameters	<p><i>ip-int-name</i> — The name of the network IP interface. An interface name cannot be in the form of an IP address. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.</p> <p>Values 1 to 32 alphanumeric characters.</p>

admin-group

Syntax	[no] admin-group <i>group-name</i> [<i>group-name...</i> (up to 5 max)] no admin-group
Context	<pre>config>router>interface>if-attribute config>service>ies>interface>if-attribute config>service>vprn>interface>if-attribute config>router>mpls>interface</pre>
Description	<p>This command configures the admin group membership of an interface. The user can apply admin groups to an IES, VPRN, network IP, or MPLS interface. Each single operation of the admin-group command allows a maximum of five (5) groups to be specified at a time. However, a maximum of 32 groups can be added to a given interface through multiple operations. Once an admin group is bound to one or more interface, its value cannot be changed until all bindings are removed. The configured admin-group membership will be applied in all levels/areas the interface is participating in. The same interface cannot have different memberships in different levels/areas. It should be noted that only the admin groups bound to an MPLS interface are advertised in TE link TLVs and sub-TLVs when the traffic-engineering option is enabled in IS-IS or OSPF. IES and VPRN interfaces do not have their attributes advertised in TE TLVs.</p> <p>The no form of this command deletes one or more of the admin-group memberships of an interface. The user can also delete all memberships of an interface by not specifying a group name.</p>
Default	no admin-group

Parameters *group-name* — Specifies the name of the group with up to 32 characters. The association of group name and value should be unique within an IP/MPLS domain.

auto-bandwidth-multipliers

Syntax **auto-bandwidth-multipliers sample-multiplier** *number1* **adjust-multiplier** *number2*
no auto-bandwidth-multipliers

Context config>router>mpls

Description This command specifies the number of collection intervals in the adjust interval.

Parameters **sample-multiplier** *number1* — Specifies the multiplier for collection intervals in a sample interval.

Values 1 — 511

Default 1

adjust-multiplier *number2* — Specifies the number of collection intervals in the adjust interval.

Values 1 — 16383

Default 288

auto-lsp

Syntax **auto-lsp lsp-template** *template-name* {**policy** *peer-prefix-policy* [**peer-prefix-policy**... (upto 5 max)] | **one-hop**}
no auto-lsp lsp-template *template-name*

Context config>router>mpls

Description This command enables the automatic creation of an RSVP point-to-point LSP to a destination node whose router-id matches a prefix in the specified peer prefix policy. This LSP type is referred to as auto-LSP of type mesh.

The user can associate multiple templates with same or different peer prefix policies. Each application of an LSP template with a given prefix in the prefix list will result in the instantiation of a single CSPF computed LSP primary path using the LSP template parameters as long as the prefix corresponds to a router-id for a node in the TE database. This feature does not support the automatic signaling of a secondary path for an LSP. If the user requires the signaling of multiple LSPs to the same destination node, s/he must apply a separate LSP template to the same or different prefix list that contains the same destination node. Each instantiated LSP will have a unique LSP-id and a unique tunnel-ID. This feature also does not support the signaling of a non-CSPF LSP. The selection of the **no cspf** option in the LSP template is thus blocked.

Up to five (5) peer prefix policies can be associated with a given LSP template at all times. Each time the user executes the above command with the same or different prefix policy associations, or the user changes a prefix policy associated with an LSP template, the system re-evaluates the prefix policy. The outcome of the re-evaluation will tell MPLS if an existing LSP needs to be torn down or if a new LSP needs to be signaled to a destination address that is already in the TE database.

If a /32 prefix is added to (removed from) or if a prefix range is expanded (shrunk) in a prefix list associated with a LSP template, the same prefix policy re-evaluation described above is performed.

The user must perform a **no shutdown** of the template before it takes effect. Once a template is in use, the user must shutdown the template before effecting any changes to the parameters except for those LSP parameters for which the change can be handled with the Make-Before-Break (MBB) procedures. These parameters are **bandwidth** and enabling **fast-reroute** without the **hop-limit** or **node-protect** options. For all other parameters, the user shuts down the template and once it is added, removed or modified, the existing instances of the LSP using this template are torn down and re-signaled.

The trigger to signal the LSP is when the router with a router-id matching a prefix in the prefix list appears in the Traffic Engineering database. The signaled LSP is installed in the Tunnel Table Manager (TTM) and is available to applications such as LDP-over-RSVP, resolution of BGP label routes, resolution of BGP, IGP, and static routes. It can also be used for provisioning an SDP is however not available to be used as a provisioned SDP for explicit binding or auto-binding by services.

Except for the MBB limitations to the configuration parameter change in the LSP template, MBB procedures for manual and timer based re-signaling of the LSP, for TE Graceful Shutdown and for soft pre-emption are supported.

The **one-to-one** option under **fast-reroute**, the LSP Diff-Serv **class-type** and **backup-class-type** parameters are not supported. If **diffserv-te** is enabled under RSVP, the auto-created LSP will still be signaled but with the default LSP class type.

If the **one-hop** option is specified instead of a prefix list, this command enables the automatic signaling of one-hop point-to-point LSPs using the specified template to all directly connected neighbors. This LSP type is referred to as auto-LSP of type one-hop. Although the provisioning model and CLI syntax differ from that of a mesh LSP only by the absence of a prefix list, the actual behavior is quite different. When the above command is executed, the TE database will keep track of each TE link that comes up to a directly connected IGP neighbor which router-id is discovered. It then instructs MPLS to signal an LSP with a destination address matching the router-id of the neighbor and with a strict hop consisting of the address of the interface used by the TE link. Thus, the **auto-lsp** command with the **one-hop** option will result in one or more LSPs signaled to the neighboring router.

An auto-created mesh or one-hop LSP can have egress statistics collected at the ingress LER by adding the **egress-statistics** node configuration into the LSP template. The user can also have **ingress statistics** collected at the egress LER using the same ingress-statistics node in CLI used with a provisioned LSP. The user must specify the full LSP name as signaled by the ingress LER in the RSVP session name field of the Session Attribute object in the received Path message.

The **no** form of this command deletes all LSP signaled using the specified template and prefix policy. When the **one-hop** option is used, it deletes all one-hop LSPs signaled using the specified template to all directly connected neighbors.

Parameters **lsp-template** *template-name* — Specifies an LSP template name up to 32 characters in length.
policy *peer-prefix-policy* — Specifies an peer prefix policy name up to 32 characters in length.

bypass-resignal-timer

Syntax	bypass-resignal-timer <i>minutes</i> no bypass-resignal-timer
Context	config>router>mpls
Description	<p>This command triggers the periodic global re-optimization of all dynamic bypass LSP paths associated with RSVP P2P LSP. The operation is performed at each expiry of the user configurable bypass LSP re-signal timer.</p> <p>When this command is enabled, MPLS makes a request to CSPF for the best path for each dynamic bypass LSP originated on this node. The constraints of the first associated LSP primary path and which originally triggered the signaling of the bypass LSP must be satisfied. In order to do this, MPLS saves the original Path State Block (PSB) of that LSP primary path even if the latter is torn down.</p> <p>If CSPF returns no path or returns a new path that is equal in terms of cost to the current path, the PSB associations are not updated. If CSPF returns a new path with a different cost from the current one, MPLS will signal it.</p> <p>Once the new path is successfully signaled, MPLS will evaluate each PSB of each PLR (i.e., each unique avoid-node or avoid-link constraint) associated with the older bypass LSP path to check if the corresponding LSP primary path constraints are still satisfied by the new bypass LSP path. If so, the PSB association is moved to the new bypass LSP.</p> <p>Each PSB whose constraints are not satisfied remains associated with the older bypass LSP and will be checked at the next background PSB re-evaluation, or at the next timer or manual bypass re-optimization. Furthermore, if the older bypass LSP is SRLG disjoint with a primary path that has the non-strict SRLG constraint while the new bypass LSP is not SRLG disjoint, the PSB association is not moved.</p> <p>If a specific PLR associated with a bypass LSP is active, the corresponding PSBs remain associated with the older bypass LSP until the Global Revertive Make-Before-Break (MBB) tears down all corresponding primary paths, which will also cause the older bypass LSP to be torn down. Note that while in that state, the older bypass LSP will not get any new PSB association until it is torn down.</p> <p>This feature also implements a background PSB re-evaluation task which audits in the background each RSVP session and determines if an existing manual or dynamic bypass is more optimal for that session. If so, it moves the PSB association to this bypass. If the PLR for this session is active, no action is taken and the PSB will be re-examined at the next re-evaluation.</p> <p>The periodic bypass re-optimization feature evaluates only the PSBs of the PLRs associated with that bypass LSP and only against the new bypass LSP path. The background re-evaluation task will, however, audit all PSBs on the system against all existing manual and dynamic bypass LSPs.</p> <p>Furthermore, PSBs that have not been moved by the dynamic or manual re-optimization of a bypass LSP, due to the PSB constraints not being met by the new signaled bypass LSP path, will be re-evaluated by the background task against all existing manual and dynamic bypass LSPs.</p> <p>Finally, the background re-evaluation task will check for PSBs that have requested node-protect bypass LSP but are currently associated with a link-protect bypass LSP, as well as PSBs that requested FRR protection and that have no association. This is in addition to the attempt made at the receipt of a Resv on the protected LSP path such that the association is speed up.</p> <p>This feature is not supported with inter-area dynamic bypass LSP and bypass LSP protecting S2L paths of a P2MP LSP.</p>

The **no** form of this command disables the periodic global re-optimization of dynamic bypass LSP paths.

Default no bypass-resignal timer. The periodic global re-optimization of dynamic bypass LSP paths is disabled.

Parameters *minutes* — Specifies the time, in minutes, MPLS waits before attempting to re-signal dynamic bypass LSP paths originated on the system.

Values 30 — 10080

cspf-on-loose-hop

Syntax **[no] cspf-on-loose-hop**

Context config>router>mpls

Description This command enables the option to do CSPF calculations until the next loose hop or the final destination of LSP on LSR. On receiving a PATH message on LSR and processing of all local hops in the received ERO, if the next hop is loose, then the LSR node will first do a CSPF calculation until the next loose hop. On successful completion of CSPF calculation, ERO in PATH message is modified to include newly calculated intermediate hops and propagate it forward to the next hop. This allows setting up inter-area LSPs based on ERO expansion method.

NOTE: The LSP may fail to set up if this option is enabled on an LSR that is not an area border router and receives a PATH message without proper next loose hop in ERO. The 'cspf-on-loose-hop' configuration is allowed to change dynamically and applied to new LSP setup after change.

Default no cspf-on-loose-hop

srlg-group

Syntax **[no] srlg-group** *group-name* [*group-name...*(up to 5 max)]

Context config>router>mpls>interface

Description This command defines the association of RSVP interface to an SRLG group. An interface can belong to up to 64 SRLG groups. However, each single operation of the **srlg-group** command allows a maximum of 5 groups to be specified at a time.

The **no** form of this command deletes the association of the interface to the SRLG group.

Default none

Parameters *group-name* — Specifies the name of the SRLG group within a virtual router instance up to 32 characters.

te-metric

Syntax	te-metric <i>value</i> no te-metric
Context	config>router>mpls>interface
Description	<p>This command configures the traffic engineering metric used on the interface. This metric is in addition to the interface metric used by IGP for the shortest path computation.</p> <p>This metric is flooded as part of the TE parameters for the interface using an opaque LSA or an LSP. The IS-IS TE metric is encoded as sub-TLV 18 as part of the extended IS reachability TLV. The metric value is encoded as a 24-bit unsigned integer. The OSPF TE metric is encoded as a sub-TLV Type 5 in the Link TLV. The metric value is encoded as a 32-bit unsigned integer.</p> <p>When the use of the TE metric is enabled for an LSP, CSPF will first prune all links in the network topology which do not meet the constraints specified for the LSP path. Such constraints include bandwidth, admin-groups, and hop limit. Then, CSPF will run an SPF on the remaining links. The shortest path among the all SPF paths will be selected based on the TE metric instead of the IGP metric which is used by default.</p> <p>The TE metric in CSPF LSP path computation can be configured by entering the command config>router>mpls>lsp>cspf>use-te-metric.</p> <p>Note that the TE metric is only used in CSPF computations for MPLS paths and not in the regular SPF computation for IP reachability.</p> <p>The no form of the command reverts to the default value.</p>
Default	no te-metric
	The value of the IGP metric is advertised in the TE metric sub-TLV by IS-IS and OSPF.
Parameters	<i>value</i> — Specifies the metric value.
	Values 1 — 16777215

node-id-in-rro

Syntax	[no] node-id-in-rro <include exclude>
Context	config>router>rsvp>
Description	<p>This command enables the option to include node-id sub-object in RRO. Node-ID sub-object propagation is required to provide fast reroute protection for LSP that spans across multiple area domains.</p> <p>If this option is disabled, then node-id is not included in RRO object.</p>
Default	node-id-in-rro exclude

p2p-merge-point-abort-timer

Syntax	p2p-merge-point-abort-timer [1.. 65535] <i>seconds</i> no p2p-merge-point-abort-timer
Context	config>router>rsvp
Description	
Default	0 (disabled)

p2mp-merge-point-abort-timer

Syntax	p2mp-merge-point-abort-timer [1.. 65535] <i>seconds</i> no p2mp-merge-point-abort-timer
Context	config>router>rsvp
Description	
Default	0 (disabled)

p2p-active-path-fast-retry

Syntax	p2p-active-path-fast-retry <i>seconds</i> [1..10] <i>seconds</i> no p2p-active-path-fast-retry
Context	config>router>rsvp
Description	
Default	0 (disabled)

p2mp-s21-fast-retry

Syntax	p2mp-s21-fast-retry <i>seconds</i> [1..10] <i>seconds</i> no p2mp-s21-fast-retry
Context	config>router>rsvp
Description	
Default	0 (disabled)

preemption-timer

Syntax	preemption-timer <i>seconds</i> no preemption-timer
Context	config>router>rsvp
Description	<p>This parameter configures the time in seconds a node holds to a reservation for which it triggered the soft pre-emption procedure.</p> <p>The pre-empting node starts a separate preemption timer for each pre-empted LSP path. While this timer is on, the node should continue to refresh the Path and Resv for the pre-empted LSP paths. When the preemption timer expires, the node tears down the reservation if the head-end node has not already done so.</p> <p>A value of zero means the LSP should be pre-empted immediately; hard pre-empted.</p> <p>The no form of this command reverts to the default value.</p>
Default	300
Parameters	<i>seconds</i> — Specifies the time, in seconds, of the preemption timer. Values 0 — 1800 seconds

label-map

Syntax	[no] label-map <i>in-label</i>
Context	config>router>mpls>interface
Description	<p>This command is used on transit routers when a static LSP is defined. The static LSP on the ingress router is initiated using the config router mpls static-lsp <i>lsp-name</i> command. An <i>in-label</i> can be associated with either a pop or a swap action, but not both. If both actions are specified, the last action specified takes effect.</p> <p>The no form of this command deletes the static LSP configuration associated with the <i>in-label</i>.</p>
Parameters	<i>in-label</i> — Specifies the incoming MPLS label on which to match. Values 32 — 1023

pop

Syntax	[no] pop
Context	config>router>mpls>if>label-map
Description	<p>This command specifies that the incoming label must be popped (removed). No label stacking is supported for a static LSP. The service header follows the top label. Once the label is popped, the packet is forwarded based on the service header.</p> <p>The no form of this command removes the pop action for the <i>in-label</i>.</p>
Default	none

shutdown

Syntax	[no] shutdown
Context	config>router>mpls>if>label-map
Description	<p>This command disables the label map definition. This drops all packets that match the specified <i>in-label</i> specified in the label-map in-label command.</p> <p>The no form of this command administratively enables the defined label map action.</p>
Default	no shutdown

swap

Syntax	swap { <i>out-label</i> implicit-null-label } nexthop <i>ip-address</i> no swap { <i>out-label</i> implicit-null-label }
Context	config>router>mpls>interface>label-map
Description	This command swaps the incoming label and specifies the outgoing label and next hop IP address on an LSR for a static LSP. The no form of this command removes the swap action associated with the <i>in-label</i> .
Default	none
Parameters	implicit-null-label — Specifies the use of the implicit label value for the outgoing label of the swap operation. <i>out-label</i> — Specifies the label value to be swapped with the in-label. Label values 16 through 1,048,575 are defined as follows: Label values 16 through 31 are reserved. Label values 32 through 1,023 are available for static assignment. Label values 1,024 through 2,047 are reserved for future use. Label values 2,048 through 18,431 are statically assigned for services. Label values 28,672 through 131,071 are dynamically assigned for both MPLS and services. Label values 131,072 through 1,048,575 are reserved for future use. Values 16 — 1048575 nexthop <i>ip-address</i> — The IP address to forward to. If an ARP entry for the next hop exists, then the static LSP will be marked operational. If ARP entry does not exist, software will set the operational status of the static LSP to down and continue to ARP for the configured nexthop. Software will continuously try to ARP for the configured nexthop at a fixed interval.

mpls-tp-mep

Syntax	[no] mpls-tp-mep
Context	config>router>mpls>interface
Description	This command enables the context for a section layer MEP for MPLS-TP on an MPLS interface.
Default	none

if-num

Syntax	if-num <i>if-num</i> no if-num
Context	config>router>mpls>interface>mpls-tp-mep
Description	This command configures the MPLS-TP interface number for the MPLS interface. This is a 32-bit unsigned integer that is node-wide unique.
Parameters	<i>if-num</i> — This is a 32-bit value that is unique to the node.
	Values 1 — 4,294,967,295

if-num-validation

Syntax	if-num-validation { enable disable } no if-num-validation
Context	config>router>mpls>interface>mpls-tp-mep
Description	<p>The if-num-validation command is used to enable or disable validation of the if-num in LSP Trace packet against the locally configured if-num for the interface over which the LSP Trace packet was received at the egress LER. This is because some 3rd-party implementations may not perform interface validation for unnumbered MPLS-TP interfaces and instead set the if-num in the dsmap TLV to 0. If the value is 'enabled', the node performs the validation of the ingress and egress if-nums received in the LSP echo request messages that ingress on this MPLS-interface. It validates that the message arrives on the interface as identified by the ingress if-num, and is forwarded on the interface as identified by the egress if-num.</p> <p>If the value is 'disabled', no validation is performed for the ingress and egress if-nums received in the LSP echo request messages that ingress on this MPLS-interface."</p>
Default	enable
Parameters	<p>enable — Enables interface number validation.</p> <p>disable — Disables interface number validation.</p>

MPLS-TP Commands

mpls-tp

Syntax	[no] mpls-tp
Context	config>router>mpls
Description	Generic MPLS-TP parameters and MPLS-TP trabsit paths are configured under this context. If a user configures no mpls , normally the entire mpls configuration is deleted. However, in the case of mpls-tp, a check is made that there is no other mpls-tp configuration (e.g., services or LSPs using mpls-tp on the node). The mpls-tp context cannot be deleted if MPLS-TP LSPs or SDPs exist on the system. A shutdown of mpls-tp will bring down all MPLS-TP LSPs on the system.
Default	no mpls-tp

tp-tunnel-id-range

Syntax	tp-tunnel-id-range <i>start-id end-id</i> no tp-tunnel-id-range
Context	config>router>mpls>mpls-tp
Description	This command configures the range of MPLS tunnel IDs reserved for MPLS-TP LSPs. The maximum difference between the start-id and end-id is 4K. The tunnel ID referred to here is the RSVP-TE tunnel ID. This maps to the MPLS-TP Tunnel Number. There are some cases where the dynamic LSPs may have caused fragmentation to the number space such that contiguous range [<i>end-id</i> – <i>start-id</i>] is not available. In these cases, the command will fail. There are no default values for the <i>start-id</i> and <i>end-id</i> of the tunnel id range, and they must be configured to enable MPLS-TP.
Default	no tunnel-id-range
Parameters	<i>start-id</i> — Specifies the start ID. Values 1 — 61440 <i>end-id</i> — Specifies the end ID. Values 1 — 61440

oam-template

Syntax	[no] oam-template <i>name</i>
Context	config>router>mpls>mpls-tp
Description	This command creates or edits an OAM template. Generally applicable proactive OAM parameters are configured using templates. The top-level template is the OAM template. Generic MPLS-TP OAM and fault management parameters are configured in the OAM Template. Proactive CC/CV uses BFD and parameters such as Tx/Rx timer intervals, multiplier and other session/fault management parameters specific to BFD are configured using a BFD Template, which is referenced from the OAM template.
Default	no oam-template
Parameters	<i>name</i> — Specifies a text string name for the template of up to 32 characters in printable 7-bit ASCII, enclosed in double quotes. Named OAM templates are referenced from the MPLS-TP path MEP configuration.

hold-time-down

Syntax	hold-time-down <i>timer</i> no hold-time-down
Context	config>router>mpls>mpls-tp>oam-template
Description	This command configures the hold-down dampening timer. It is equivalent to a hold-off timer.
Default	no hold-time-down
Parameters	<i>interval</i> — Specifies the hold-down dampening timer interval. Values 0 — 5000 deciseconds in 10 ms increments

hold-time-up

Syntax	hold-time-up <i>timer</i> no hold-time-up
Context	config>router>mpls>mpls-tp>oam-template
Description	This command configures the hold-up dampening timer. This can be used to provide additional dampening to the state of proactive CC BFD sessions.
Default	no hold-time-up
Parameters	<i>interval</i> — Specifies the hold-up dampening timer interval. Values 0 — 500 deciseconds, in 100 ms increments Default 2 seconds

bfd-template

Syntax	bfd-template <i>name</i> no bfd-template
Context	config>router>mpls>mpls-tp>oam-template
Description	This command configures a named BFD template to be referenced by an OAM template.
Default	no bfd-template
Parameters	<i>name</i> — Specifies the BFD template name as a text string up to 32 characters in printable 7-bit ASCII, enclosed in double quotes.
	Values

protection-template

Syntax	protection-template <i>name</i> no protection-template
Context	config>router>mpls>mpls-tp
Description	Protection templates are used to define generally applicable protection parameters for MPLS-TP tunnels. Only linear protection is supported, and so the application of a named template to an MPLS-TP LSP implies that linear protection is used. A protection template is applied under the MEP context of the protect-path of an MPLS-TP LSP. The protection-template command creates or edits a named protection template.
Default	no protection-template
Parameters	<i>name</i> — Specifies the protection template name as a text string of up to 32 characters in printable 7-bit ASCII, enclosed in double quotes.

revertive

Syntax	[no] revertive
Context	config>router>mpls>mpls-tp>protection-template
Description	This command configured revertive behavior for MPLS-TP linear protection. The protect-tp-path MEP must be in the shutdown state for of the MPLS-TP LSPs referencing this protection template in order to change the revertive parameter.
Default	revertive

wait-to-restore

Syntax	wait-to-restore <i>interval</i> no wait-to-restore
Context	config>router>mpls>mpls-tp>protection-template
Description	This command configures the WTR timer. It determines how long to wait until the active path of an MPLS-TP LSP is restored to the working path following the clearing of a defect on the working path. It is applicable to revertive mode, only.
Default	no wait-to-restore
Parameters	<i>interval</i> — Specifies the WTR timer interval. Values 0 — 720 seconds in 1 second increments

rapid-psc-timer

Syntax	rapid-psc-timer <i>interval</i> no rapid-psc-timer
Context	config>router>mpls>mpls-tp>protection-template
Description	This command configures the rapid timer value to be used for protection switching coordination (PSC) packets for MPLS-TP linear protection (RFC 6378).
Default	no rapid-psc-timer
Parameters	<i>interval</i> — Specifies the rapid timer interval. Values [10, 100, 1000 ms] Default 10 ms

slow-psc-timer

Syntax	slow-psc-timer <i>interval</i> no slow-psc-timer
Context	config>router>mpls>mpls-tp>protection-template
Description	This command configures the slow timer value to be used for protection switching coordination (PSC) packets for MPLS-TP linear protection (RFC 6378).
Default	no rapid-psc-timer
Parameters	<i>interval</i> — Specifies the slow timer interval. Values [10, 100, 1000 ms]

global-id

Syntax	global-id <i>global-id</i> no global-id
Context	config>router>mpls>mpls-tp
Description	<p>This command configures the MPLS-TP Global ID for the node. This is used as the ‘from’ Global ID used by MPLS-TP LSPs originating at this node. If a value is not entered, the Global ID is taken to be Zero. This is used if the global-id is not configured. If an operator expects that inter domain LSPs will be configured, then it is recommended that the global ID should be set to the local ASN of the node, as configured under config>system. If two-byte ASNs are used, then the most significant two bytes of the global-id are padded with zeros.</p> <p>In order to change the value of the global-id, config>router>mpls>mpls-tp must be in the shutdown state. This will bring down all of the MPLS-TP LSPs on the node. New values a propagated to the system when a no shutdown is performed.</p>
Default	no global-id
Parameters	<i>global-id</i> — Specifies the global ID for the node. Values 0 — 4294967295

node-id

Syntax	node-id <i>node-id</i> no node-id
Context	config>router>mpls>mpls-tp
Description	<p>This command configures the MPLS-TP Node ID for the node. This is used as the ‘from’ Node ID used by MPLS-TP LSPs originating at this node. The default value of the node-id is the system interface IPv4 address. The Node ID may be entered in 4-octed IPv4 address format, <a.b.c.d>, or as an unsigned 32 bit integer. Note that it is not treated as a routable IP address from the perspective of IP routing, and is not advertised in any IP routing protocols.</p> <p>The MPLS-TP context cannot be administratively enabled unless at least a system interface IPv4 address is configured because MPLS requires that this value is configured.</p>
Default	no node-id
Parameters	<i>node-id</i> — Specifies the MPLS-TP node ID for the node. Values <a.b.c.d> or [1— 4294967295] Default System interface IPv4 address

transit-path

Syntax	transit-path <i>path-name</i> no transit-path
Context	config>router>mpls>mpls-tp
Description	This command enables the configuration or editing of an MPLS-TP transit path at an LSR.
Default	no transit-path
Parameters	<i>path-name</i> — Specifies the template of up to 32 characters in printable 7-bit ASCII, enclosed in double quotes.

path-id

Syntax	path-id { lsp-num <i>lsp-num</i> working-path protect-path [src-global-id <i>src-global-id</i>] src-node-id <i>src-node-id</i> src-tunnel-num <i>src-tunnel-num</i> [dest-global-id <i>dest-global-id</i>] dest-node-id <i>dest-node-id</i> [dest-tunnel-num <i>dest-tunnel-num</i>]} no path-id
Context	config>router>mpls>mpls-tp>transit-path
Description	<p>This command configures path ID for an MPLS-TP transit path at an LSR. The path ID is equivalent to the MPLS-TP LSP ID and is used to generate the maintenance entity group intermediate point (MIP) identifier for the LSP at the LSR. A path-id must be configured for on-demand OAM to verify an LSP at the LSR.</p> <p>The path-id must contain at least the following parameters: lsp-num, src-node-id, src-global-id, src-tunnel-num, dest-node-id.</p> <p>The path-id must be unique on a node. It is recommended that this is also configured to be a globally unique value.</p> <p>The no form of the command removes the path ID from the configuration.</p>
Default	no path-id
Parameters	<p><i>lsp-num</i> — Specifies the LSP number.</p> <p>Values 1 — 65535, or working path, or protect-path. A working-path is equivalent to a lsp-num of 1, and a protect-path is an lsp-num of 2.</p> <p><i>src-global-id</i> — Specifies the source global ID.</p> <p>Values 0 — 4294967295</p> <p><i>src-node-id</i> — Specifies the source node ID.</p> <p>Values a.b.c.d or 1 — 4294967295</p> <p><i>src-tunnel-num</i> — Specifies the source tunnel number.</p> <p>Values 1 — 61440</p>

dest-global-id — Specifies the destination global ID. If the destination global ID is not entered, then it is set to the same value as the source global ID.

Values 0 — 4294967295

dest-node-id — Specifies the destination node ID.

Values a.b.c.d or 1 — 4294967295

dest-tunnel-num — Specifies the destination tunnel number. If the destination tunnel number is not entered, then it is set to the same value as the source tunnel number.

Values 1 — 61440

forward-path

Syntax	[no] forward-path
Context	config>router>mpls>mpls-tp>transit-path
Description	This command enables the forward path of an MPLS-TP transit path to be created or edited. The forward path must be created before the reverse path. The no form of this command removes the forward path. The forward path cannot be removed if a reverse exists.
Default	no forward-path

reverse-path

Syntax	[no] reverse-path
Context	config>router>mpls>mpls-tp>transit-path
Description	This command enables the reverse path of an MPLS-TP reverse path to be created or edited. The reverse path must be created after the forward path. The no form of this command removes the reverse path. The reverse path must be removed before the forward path.
Default	no reverse-path

in-label

Syntax	in-label <i>in-label</i> out-label <i>out-label</i> out-link <i>if-name</i> [next-hop <i>next-hop</i>] no in-label
Context	config>router>mpls>mpls-tp>transit-path>forward-path config>router>mpls>mpls-tp>transit-path>reverse-path
Description	This command configures the label mapping associated with a forward path or reverse path of an MPLS-TP transit path to be configured. The incoming label, outgoing label and outgoing interface must be configured, using the in-label , out-label and out-link parameters. If the out-link refers to a numbered IP interface, the user may optionally configure the next-hop parameter and the system will determine the interface to use to reach the configured next-hop, but will check that the user-entered value for the <i>out-link</i> corresponds to the link returned by the system. If they do not correspond, then the path will not come up.
Default	no in-label
Parameters	<i>in-label</i> — Specifies the in label. Values 32 — 16415 <i>out-label</i> — Specifies the out label. Values 32 — 16415 <i>if-name</i> — Specifies the name of the outgoing interface use for the path. <i>next-hop</i> — Specifies the next-hop. Values a.b.c.d

shutdown

Syntax	[no] shutdown
Context	config>router>mpls>mpls-tp>transit-path
Description	This command administratively enables or disables an MPLS-TP transit path.
Default	no shutdown

LSP Commands

lsp

Syntax	[no] lsp <i>lsp-name</i> [bypass-only p2mp-lsp mpls-tp <i>src-tunnel-num</i>]
Context	config>router>mpls
Description	<p>This command creates an LSP that is either signaled dynamically by the router, or a statically provisioned MPLS-TP LSP.</p> <p>When the LSP is created, the egress router must be specified using the to command and at least one primary or secondary path must be specified for signaled LSPs, or at least one working path for MPLS-TP LSPs. All other statements under the LSP hierarchy are optional.</p> <p>LSPs are created in the administratively down (shutdown) state.</p> <p>The no form of this command deletes the LSP. All configuration information associated with this LSP is lost. The LSP must be administratively shutdown and unbound from all SDPs before it can be deleted.</p>
Default	none
Parameters	<p><i>lsp-name</i> — Name that identifies the LSP. The LSP name can be up to 32 characters long and must be unique.</p> <p>bypass-only — Defines an LSP as a manual bypass LSP exclusively. When a path message for a new LSP requests bypass protection, the PLR first checks if a manual bypass tunnel satisfying the path constraints exists. If one is found, the router selects it. If no manual bypass tunnel is found, the router dynamically signals a bypass LSP in the default behavior. The CLI for this feature includes a knob that provides the user with the option to disable dynamic bypass creation on a per node basis.</p> <p>p2mp-lsp — Defines an LSP as a point-to-multipoint LSP. The following parameters can be used with a P2MP LSP: <i>adaptive</i>, <i>adspec</i>, <i>cspf</i>, <i>exclude</i>, <i>fast-reroute</i>, <i>from</i>, <i>hop-limit</i>, <i>include</i>, <i>metric</i>, <i>retry-limit</i>, <i>retry-timer</i>, <i>resignal-timer</i>. The following parameters cannot be used with a P2MP LSP: <i>primary</i>, <i>secondary</i>, <i>to</i>, <i>dest-global-id</i>, <i>dest-tunnel-number</i>, <i>working-tp-path</i>, <i>protect-tp-path</i>.</p> <p>mpls-tp <i>src-tunnel-num</i> — Defines an LSP as an MPLS-TP LSP. The <i>src-tunnel-num</i> is a mandatory create time parameter for mpls-tp LSPs, and has to be assigned by the user based on the configured range of tunnel IDs. The following parameters can only be used with an MPLS-TP LSP: <i>to</i>, <i>dest-global-id</i>, <i>dest-tunnel-number</i>, <i>working-tp-path</i>, <i>protect-tp-path</i>. Other parameters defined for the above LSP types cannot be used.</p>

adaptive

Syntax	[no] adaptive
Context	config>router>mpls>lsp config>router>mpls>lsp-template
Description	This command enables the make-before-break functionality for an LSP or LSP path. When enabled for the LSP, make-before-break will be performed for primary path and all the secondary paths of the LSP.
Default	adaptive

adspec

Syntax	[no] adspec
Context	config>router>mpls>lsp config>router>mpls>lsp-template
Description	<p>When enabled, the ADSPEC object will be included in RSVP messages for this LSP. The ADSPEC object is used by the ingress LER to discover the minimum value of the MTU for links in the path of the LSP. By default, the ingress LER derives the LSP MTU from that of the outgoing interface of the LSP path.</p> <p>Note that a bypass LSP always signals the ADSPEC object since it protects both primary paths which signal the ADSPEC object and primary paths which do not. This means that MTU of LSP at ingress LER may change to a different value from that derived from the outgoing interface even if the primary path has ADSPEC disabled.</p>
Default	no adspec — No ADSPEC objects are included in RSVP messages.

auto-bandwidth

Syntax	[no] auto-bandwidth
Context	config>router>mpls>lsp config>router>mpls>lsp-template
Description	<p>This command enables (and the no form disables) automatic adjustments of LSP bandwidth.</p> <p>Auto-bandwidth at the LSP level cannot be executed unless adaptive is configured in the config>router>mpls>lsp context.</p>
Default	no auto-bandwidth

adjust-down

Syntax	adjust-down <i>percent</i> [bw <i>mbps</i>] no adjust-down
Context	config>router>mpls>lsp>auto-bandwidth config>router>mpls>lsp-template>auto-bandwidth
Description	This command configures the minimum threshold for decreasing the bandwidth of an LSP based on active measurement of LSP bandwidth. The no form of this command is equivalent to adjust-down 5.
Default	no adjust-down
Parameters	<i>percent</i> — Specifies the minimum difference between the current bandwidth reservation of the LSP and the (measured) maximum average data rate, expressed as a percentage of the current bandwidth, for decreasing the bandwidth of the LSP. Values 1 — 100 Default 5 <i>mbps</i> — Specifies the minimum difference between the current bandwidth reservation of the LSP and the (measured) maximum average data rate, expressed as an absolute bandwidth (mbps), for decreasing the bandwidth of the LSP. Values 0 — 100000 Default 0

adjust-up

Syntax	adjust-up <i>percent</i> [bw <i>mbps</i>] no adjust-up
Context	config>router>mpls>lsp>auto-bandwidth config>router>mpls>lsp-template>auto-bandwidth
Description	This command configures the minimum threshold for increasing the bandwidth of an LSP based on active measurement of LSP bandwidth. The no form of this command is equivalent to adjust-up 5.
Default	no adjust-up
Parameters	<i>percent</i> — Specifies the minimum difference between the current bandwidth reservation of the LSP and the (measured) maximum average data rate, expressed as a percentage of the current bandwidth, for increasing the bandwidth of the LSP. Values 1-100 Default 5

mbps — Specifies the minimum difference between the current bandwidth reservation of the LSP and the (measured) maximum average data rate, expressed as an absolute bandwidth (mbps), for increasing the bandwidth of the LSP

Values 0 — 100000

Default 0

fc

Syntax **fc** *fc-name* **sampling-weight** *sampling-weight*
no fc

Context config>router>mpls>lsp-template>auto-bandwidth

Description This command configures the sampling weight.

max-bandwidth

Syntax **max-bandwidth** *mbps*
no max-bandwidth

Context config>router>mpls>lsp>auto-bandwidth
config>router>mpls>lsp-template>auto-bandwidth

Description This command configures the maximum bandwidth that auto-bandwidth allocation is allowed to request for an LSP.

The LSP maximum applies whether the bandwidth adjustment is triggered by normal adjust-interval expiry, the overflow limit having been reached, or manual request.

The **no** form of the command means max-bandwidth is 100 Gbps.

The max-bandwidth must be greater than the min-bandwidth.

Default no max-bandwidth

Parameters *mbps* — Specifies the maximum bandwidth in mbps.

Values 0 — 100000

Default 0

min-bandwidth

Syntax	min-bandwidth <i>mbps</i> no min-bandwidth
Context	config>router>mpls>lsp>auto-bandwidth config>router>mpls>lsp-template>auto-bandwidth
Description	<p>This command configures the minimum bandwidth that auto-bandwidth allocation is allowed to request for an LSP.</p> <p>The LSP minimum applies whether the bandwidth adjustment is triggered by normal adjust-timer expiry or manual request.</p> <p>The no form of the command means min-bandwidth is zero.</p>
Default	no min-bandwidth
Parameters	<i>mbps</i> — Specifies the minimum bandwidth in mbps.
Values	0 — 100000
Default	0

monitor-bandwidth

Syntax	[no] monitor-bandwidth
Context	config>router>mpls>lsp>auto-bandwidth config>router>mpls>lsp-template>auto-bandwidth
Description	<p>This command enables the collection and display of auto-bandwidth measurements, but prevents any automatic bandwidth adjustments from taking place.</p> <p>This command is mutually exclusive with the overflow-limit command.</p> <p>The no form of the command the collection and display of auto-bandwidth measurements.</p>

multipliers

Syntax	multipliers sample-multiplier <i>num1</i> adjust-multiplier <i>num2</i> no multipliers
Context	config>router>mpls>lsp>auto-bandwidth config>router>mpls>lsp-template>auto-bandwidth
Description	<p>This command configures the sample-multiplier and adjust-multiplier applicable to one particular LSP.</p> <p>The sample-multiplier configures the number of collection intervals between measurements of the number of bytes that have been transmitted on the LSP. The byte counts include the layer 2 encapsulation of MPLS packets and represent traffic of all forwarding classes and priorities (in-profile vs, out-of-profile) belonging to the LSP. The router calculates the average data rate in each</p>

sample interval. The maximum of this average data rate over multiple sample intervals is the measured bandwidth input to the auto-bandwidth adjustment algorithms.

The adjust-multiplier is the number of collection intervals between periodic evaluations by the ingress LER about whether to adjust the LSP bandwidth. The router keeps track of the maximum average data rate of each LSP since the last reset of the adjust-count.

The adjust-multiplier is not allowed to be set to a value less than the sample-multiplier. It is recommended that the adjust-multiplier be a multiple of the sample-multiplier.

The **no** form of this command instructs the system to take the value from the auto-bandwidth-defaults command.

Default **no multipliers**

Parameters *number1* — The number of collection intervals in a sample interval.

Values 1 — 511

Default inherited

number2 — The number of collection intervals in an adjust interval.

Values 1 — 16383

Default inherited

overflow-limit

Syntax **overflow-limit** *number* **threshold** *percent* [**bw** *mbps*]
no overflow-limit

Context config>router>mpls>lsp>auto-bandwidth
config>router>mpls>lsp-template>auto-bandwidth

Description This command configures overflow-triggered auto-bandwidth adjustment. It sets the threshold at which bandwidth adjustment is initiated due to the configured number of overflow samples having been reached, regardless of how much time remains until the adjust interval ends.

A sample interval is counted as an overflow if the average data rate during the sample interval is higher than the currently reserved bandwidth by at least the thresholds configured as part of this command.

If overflow-triggered auto-bandwidth adjustment is successful the overflow count, maximum average data rate and adjust count are reset. If overflow-triggered auto-bandwidth adjustment fails then the overflow count is reset but the maximum average data rate and adjust count maintain current values.

This command is mutually exclusive with the monitor-bandwidth command.

The **no** form of this command disables overflow-triggered automatic bandwidth adjustment.

Default **no overflow-limit**

Parameters *number* — The number of overflow samples that triggers an overflow auto-bandwidth adjustment attempt.

Values 1 — 10

Default 0 (disabled)

percent — The minimum difference between the current bandwidth of the LSP and the sampled data rate, expressed as a percentage of the current bandwidth, for counting an overflow sample.

Values 1 — 100

Default 0 (disabled)

mbps — The minimum difference between the current bandwidth of the LSP and the sampled data rate, expressed as an absolute bandwidth (Mbps) relative to the current bandwidth, for counting an overflow sample.

Values 1— 100000

Default 0 (disabled)

underflow-limit

Syntax **underflow-limit** *number* **threshold** *percent* [**bw** *mbps*]
no underflow-limit

Context config>router>mpls>lsp>auto-bandwidth
config>router>mpls>lsp-template>auto-bandwidth

Description This command configures underflow-triggered auto-bandwidth adjustment. An underflow auto-bandwidth adjustment can occur any time during the adjust-interval; it is triggered when the number of consecutive underflow samples reaches the threshold N configured as part of this command. The new bandwidth of the LSP after a successful underflow adjustment is the maximum data rate observed in the last N consecutive underflow samples.

A sample interval is counted as an underflow if the average data rate during the sample interval is lower than the currently reserved bandwidth by at least the thresholds configured as part of this command.

This command is mutually exclusive with the **monitor-bandwidth** command.

The **no** form of this command disables underflow-triggered automatic bandwidth adjustment.

Default no underflow-limit

Parameters *number* — The number of consecutive underflow samples that triggers an underflow auto-bandwidth adjustment attempt.

Values 0 — 10

Default 0 (disabled)

percent — The minimum difference between the current bandwidth of the LSP and the sampled data rate, expressed as a percentage of the current bandwidth, for counting an underflow sample.

Values 0 —100

Default 0 (disabled)

mbps — The minimum difference between the current bandwidth of the LSP and the sampled data rate, expressed as an absolute bandwidth (Mbps) relative to the current bandwidth, for counting an underflow sample.

Values 0 —100,000

Default 0 (disabled)

bgp-shortcut

Syntax	[no] bgp-shortcut
Context	config>router>mpls>lsp
Description	This command enables the use of RSVP LSP for IPv4 BGP routes.

bgp-transport-tunnel

Syntax	bgp-transport-tunnel <i>include exclude</i>
Context	config>router>mpls>lsp
Description	This command allows or blocks RSVP-TE LSP to be used as a transport LSP for BGP tunnel routes.
Default	bgp-transport-tunnel include
Parameters	<p><i>include</i> — Allows RSVP-TE LSP to be used as transport LSP from the ASBR to local PE router, from ingress PE to ASBR in the local AS or between multi-hop eBGP peers with ASBR to ASBR adjacency.</p> <p><i>exclude</i> — Blocks RSVP-TE LSP to be used as transport LSP from the ASBR to local PE router, from ingress PE to ASBR in the local AS or between multi-hop eBGP peers with ASBR to ASBR adjacency.</p>

class-type

Syntax	class-type <i>ct-number</i> no class-type
Context	config>router>mpls>lsp config>router>mpls>lsp>primary config>router>mpls>lsp>secondary
Description	<p>This command configures the Diff-Serv Class Type (CT) for an LSP, the LSP primary path, or the LSP secondary path. The path level configuration overrides the LSP level configuration. However, only one CT per LSP path will be allowed as per RFC 4124.</p> <p>The signaled CT of a dynamic bypass is always be CT0 regardless of the CT of the primary LSP path. The setup and hold priorities must be set to default values, i.e., 7 and 0 respectively. This assumes that the operator configured a couple of TE classes, one which combines CT0 and a priority of 7 and the</p>

other which combines CTO and a priority of 0. If not, the bypass LSP will not be signaled and will go into the down state.

The operator cannot configure the CT, setup priority, and hold priority of a manual bypass. They are always signaled with CT0 and the default setup and holding priorities.

The signaled CT and setup priority of a detour LSP must match those of the primary LSP path it is associated with.

If the operator changes the CT of an LSP or of an LSP path, or changes the setup and holding priorities of an LSP path, the path will be torn down and retried.

An LSP which does not have the CT explicitly configured will behave like a CT0 LSP when Diff-Serv is enabled.

If the operator configured a combination of a CT and a setup priority and/or a combination of a CT and a holding priority for an LSP path that are not supported by the user-defined TE classes, the LSP path will be kept in a down state and an error code will be displayed in the show command output for the LSP path.

The **no** form of this command reverts to the default value.

Default	no class-type.
Parameters	<i>ct-number</i> — The Diff-Serv Class Type number.
Values	0 – 7
Default	0

bandwidth

Syntax	bandwidth <i>rate-in-mbps</i>
Context	config>router>mpls>lsp>primary-p2mp-instance config>router>mpls>lsp-template
Description	This command specifies the amount of bandwidth to be reserved for the P2MP instance.
Parameters	<i>rate-in-mbps</i> — specifies the bandwidth, in Mbps.
Values	0 — 100000

cspf

Syntax	[no] cspf [<i>use-te-metric</i>]
Context	config>router>mpls>lsp config>router>mpls>lsp-template
Description	This command enables Constrained Shortest Path First (CSPF) computation for constrained-path LSPs. Constrained-path LSPs are the ones that take configuration constraints into account. CSPF is also used to calculate the detour routes when fast-reroute is enabled.

Explicitly configured LSPs where each hop from ingress to egress is specified do not use CSPF. The LSP will be set up using RSVP signaling from ingress to egress.

If an LSP is configured with **fast-reroute** *frr-method* specified but does not enable CSPF, then neither global revertive nor local revertive will be available for the LSP to recover.

Default no cspf

Parameters *use-te-metric* — Specifies to use the use of the TE metric for the purpose of the LSP path computation by CSPF.

dest-global-id

Syntax **dest-global-id** *dest-global-id*
no dest-global-id

Context config>router>mpls>lsp

Description This optional command configures the MPLS-TP Global ID of the far end node of the MPLS-TP LSP. This command is only allowed for MPLS-TP LSPs. Global ID values of 0 indicate that the local node's configured global ID is used. If the local global-id is 0, then the dest-global-id must also be 0. The dest-global-id cannot be changed if an LSP is in use by an SDP.

Default 0

Parameters *dest-global-id* — Specifies the destination global ID.

Values 0 — 4294967295

Default 0

dest-tunnel-number

Syntax **dest-tunnel-number** *dest-tunnel-number*
no dest-tunnel-number

Context config>router>mpls>lsp

Description This optional command configures the MPLS-TP tunnel number of the LSP at the far end node of the MPLS-TP LSP. This command is only allowed for MPLS-TP LSPs. If it is not entered, then the system will take the dest-tunnel-number to be the same as the src-tunnel-num for the LSP.

Default The default value is the configured *src-tunnel-num*.

Parameters *dest-tunnel-number* — Specifies the destination tunnel number.

Values 1 — 61440

Default *src-tunnel-number*

working-tp-path

Syntax	[no] working-tp-path
Context	config>router>mpls>lsp
Description	<p>This command creates or edits the working path for an MPLS-TP LSP. At least one working path (but not more than one working path) must be created for an MPLS-TP LSP. If MPLS-TP linear protection is also configured, then this is the path that is used as the default working path for the LSP, and it must be created prior to the protect path. The working-tp-path can only be deleted if no protect-tp-path exists for the LSP.</p> <p>The following commands are applicable to the working-tp-path: lsp-num, in-label, out-label, mep, shutdown.</p>
Default	no working-tp-path

protect-tp-path

Syntax	[no] protect-tp-path
Context	config>router>mpls>lsp
Description	<p>This command creates or edits the protect path for an MPLS-TP LSP. At least one working path must exist before a protect path can be created for an MPLS-TP LSP. If MPLS-TP linear protection is also configured, then this is the path that is used as the default protect path for the LSP. The protect path must be deleted before the working path. Only one protect path can be created for each MPLS-TP LSP.</p> <p>The following commands are applicable to the working-tp-path: lsp-num, in-label, out-label, mep, shutdown.</p>

lsp-num

Syntax	lsp-num <i>lsp-num</i> no lsp-num
Context	config>mpls>lsp>working-tp-path config>mpls>lsp>protect-tp-path
Description	This command configures the MPLS-TP LSP Number for the working TP path or the Protect TP Path.
Default	no lsp-num
Parameters	<i>lsp-num</i> — Specifies the LSP number. Values 1 — 65535 Default 1 for a working path, 2 for a protect path

in-label

Syntax	in-label <i>in-label</i> no in-label
Context	config>mpls>lsp>working-tp-path config>mpls>lsp>protect-tp-path
Description	This command configures the incoming label for the reverse path or the working path or the protect path of an MPLS-TP LSP. MPLS-TP LSPs are bidirectional, and so an incoming label value must be specified for each path.
Default	no in-label
Parameters	<i>in-label</i> — Specifies the in label.
	Values 32 — 16415

out-label

Syntax	out-label <i>out-label</i> out-link <i>if-name</i> [next-hop <i>ip-address</i>] no out-label
Context	config>mpls>lsp>working-tp-path config>mpls>lsp>protect-tp-path
Description	This command configures the outgoing label value to use for an MPLS-TP working or protect path. The out-link is the outgoing interface on the node that this path will use, and must be specified. If the out-link refers to a numbered IP interface, the user may optionally configure the next-hop parameter and the system will determine the interface to use to reach the configured next-hop, but will check that the user-entered value for the <i>out-link</i> corresponds to the link returned by the system. If they do not correspond, then the path will not come up.
Default	no out-label
Parameters	<i>out-label</i> — Specifies the out label.
	Values 32 — 16415
	<i>if-name</i> — Specifies the interface name.
	<i>ip-address</i> — Specifies the IPv4 address in a.b.c.d

mep

Syntax	[no] mep
Context	config>mpls>lsp>working-tp-path config>mpls>lsp>protect-tp-path
Description	<p>This command creates or edits an MPLS-TP maintenance entity group (MEG) endpoint (MEP) on and MPLS-TP path. MEPs represent the termination point for OAM flowing on the path, as well as linear protection for the LSP. Only one MEP can be configured at each end of the path.</p> <p>The following commands are applicable to a MEP on an MPLS-TP working or protect path: oam-template, bfd-enable, and shutdown. In addition, a protection-template may be configured on a protect path.</p> <p>The no form of the command removes a MEP from an MPLS-TP path.</p>

mip

Syntax	[no] mip
Context	config>router>mpls>lsp>transit-path>forward-path config>router>mpls>lsp>transit-path>reverse-path
Description	This command creates a context for maintenance entity group intermediate point (MIP) parameters for the forward path and the reverse path of an MPLS-TP LSP at an LSR.
Default	none

dsmap

Syntax	dsmap if-num no dsmap				
Context	config>router>mpls>lsp>working-tp-path>mep config>router>mpls>lsp>protect-tp-path>mep config>router>mpls>lsp>transit-path>forward-path>mip config>router>mpls>lsp>transit-path>reverse-path>mip				
Description	This command is used to configure the values to use in the DSMAP TLV sent by a node in an LSP Trace echo request for a static MPLS-TP LSP. A node sending a DSMAP TLV will include the in-if-num and out-if-num values. Additionally, it will include the out-label for the LSP in the Label TLV for the DSMAP in the echo request message.				
Parameters	<p><i>if-num</i> — This is a 32-bit value corresponding to the expected ingress interface if-num used by an MPLS-TP LSP for the next hop downstream. A value of zero means that no interface validation will be performed.</p> <table><tr><td>Values</td><td>0 — 4,294,967,295</td></tr><tr><td>Default</td><td>0</td></tr></table>	Values	0 — 4,294,967,295	Default	0
Values	0 — 4,294,967,295				
Default	0				

oam-template

Syntax	oam-template <i>name</i> no oam-template
Context	config>mpls>lsp>working-tp-path config>mpls>lsp>protect-tp-path
Description	This command applies a OAM template to an MPLS-TP working or protect path. It contains configuration parameters for proactive OAM mechanisms that can be enabled on the path e.g. BFD. Configuration of an OAM template is optional. The no form of the command removes the OAM template from the path.
Default	no oam-template
Parameters	<i>name</i> — Specifies a text string name for the template up to 32 characters in printable 7-bit ASCII, enclosed in double quotes.

bfd-enable

Syntax	bfd-enable [cc cc_cv] no bfd-enable
Context	config>mpls>lsp>working-tp-path config>mpls>lsp>protect-tp-path
Description	The command associates the operational state of an MPLS-TP path with a BFD session whose control packets flow on the path. The BFD packets are encapsulated in a generic associated channel (G-ACh) on the path. The timer parameters of the BFD session are taken from the the OAM template of the MEP. A value of cc means that the BFD session is only used for continuity check of the the MPLS-TP path. In this case, the cc timer parameters of the OAM template apply. A value of cv means that the BFD session is used for both continuity checking and connectivity verification, and the cv timers of the OAM template apply. This form of the bfd-enable command is only applicable when it is configured under a MEP used on an MPLS-TP working or protect path.
Default	no bfd-enable
Parameters	cc cc_cv — cc indicates that BFD runs in CC only mode. This mode uses GACH channel type 0x07. cc_cv indicates that BFD runs in combined CC and CV mode. This mode uses channel type 0x22 for MPLS-TP CC packets, and 0x23 for MPLS-TP CV packets.

protection-template

Syntax	protection-template <i>name</i> no protection-template
Context	config>mpls>lsp>protect-tp-path
Description	<p>This command applies a protection template name to an MPLS-TP LSP that the protect path is configured under. If the template is applied, then MPLS-TP 1:1 linear protection is enabled on the LSP, using the parameters specified in the named template.</p> <p>A named protection template can only be applied to the protect path context of an MPLS-TP LSP.</p> <p>The no form of the command removes the template and thus disables mpls-tp linear protection on the LSP.</p>
Default	no protection-template
Parameters	<i>name</i> — Specifies a text string for the template up to 32 characters in printable 7-bit ASCII, enclosed in double quotes.

exclude

Syntax	[no] exclude <i>group-name</i> [<i>group-name...</i> (up to 5 max)]
Context	config>router>mpls>lsp config>router>mpls>lsp-template
Description	<p>This command specifies the admin groups to be excluded when an LSP is set up in the primary or secondary contexts. Each single operation of the exclude command allows a maximum of 5 groups to be specified at a time. However, a maximum of 32 groups can be specified per LSP through multiple operations. The admin groups are defined in the config>router>mpls>admin-group context.</p> <p>Use the no form of the command to remove the exclude command.</p>
Default	no exclude
Parameters	<i>group-name</i> — Specify the existing group-name to be excluded when an LSP is set up.

exclude-node

Syntax	[no] exclude-node <i>ip-address</i>
Context	config>router>mpls>lsp
Description	<p>This command enables the option to include XRO object in the bypass LSP PATH message object. The exclude-node option is required for manual bypass LSP with XRO to FRR protect ABR node in a multi-vendor network deployment. This command must be configured on the PLR node that protects the ABR node. The ABR node IP address must be configured as exclude-node.</p>
Default	no exclude-node

fast-reroute

Syntax	fast-reroute <i>frr-method</i> no fast-reroute
Context	config>router>mpls>lsp config>router>mpls>lsp-template
Description	<p>This command creates a pre-computed detour LSP from each node in the path of the LSP. In case of failure of a link or LSP between two nodes, traffic is immediately rerouted on the pre-computed detour LSP, thus avoiding packet-loss.</p> <p>When fast-reroute is enabled, each node along the path of the LSP tries to establish a detour LSP as follows:</p> <ul style="list-style-type: none"> • Each upstream node sets up a detour LSP that avoids only the immediate downstream node, and merges back on to the actual path of the LSP as soon as possible. <p>If it is not possible to set up a detour LSP that avoids the immediate downstream node, a detour can be set up to the downstream node on a different interface.</p> <ul style="list-style-type: none"> • The detour LSP may take one or more hops (see hop-limit) before merging back on to the main LSP path. • When the upstream node detects a downstream link or node failure, the ingress router switches traffic to a standby path if one was set up for the LSP. <p>Fast reroute is available only for the primary path. No configuration is required on the transit hops of the LSP. The ingress router will signal all intermediate routers using RSVP to set up their detours. TE must be enabled for fast-reroute to work.</p> <p>If an LSP is configured with fast-reroute <i>frr-method</i> specified but does not enable CSPF, then neither global revertive nor local revertive will be available for the LSP to recover.</p> <p>The no form of the fast-reroute command removes the detour LSP from each node on the primary path. This command will also remove configuration information about the hop-limit and the bandwidth for the detour routes.</p> <p>The no form of fast-reroute hop-limit command reverts to the default value.</p>
Default	no fast-reroute — When fast-reroute is specified, the default fast-reroute method is one-to-one.
Parameters	<p>one-to-one — In the one-to-one technique, a label switched path is established which intersects the original LSP somewhere downstream of the point of link or node failure. For each LSP which is backed up, a separate backup LSP is facility — This option, sometimes called many-to-one, takes advantage of the MPLS label stack. Instead of creating a separate LSP for every backed-up LSP, a single LSP is created which serves to backup up a set of LSPs. This LSP tunnel is called a bypass tunnel.</p> <p>The bypass tunnel must intersect the path of the original LSP(s) somewhere downstream of the point of local repair (PLR). Naturally, this constrains the set of LSPs being backed-up via that bypass tunnel to those that pass through a common downstream node. All LSPs which pass through the PLR and through this common node which do not also use the facilities involved in the bypass tunnel are candidates for this set of LSPs.</p>

bandwidth

Syntax	bandwidth <i>rate-in-mbps</i> no bandwidth
Context	config>router>mpls>lsp>fast-reroute config>router>mpls>lsp-template>fast-reroute
Description	This command is used to request reserved bandwidth on the detour path. When configuring an LSP, specify the traffic rate associated with the LSP. When configuring fast reroute, allocate bandwidth for the rerouted path. The bandwidth rate does not need to be the same as the bandwidth allocated for the LSP.
Default	no bandwidth — Bandwidth is not reserved for a rerouted path.
Parameters	<i>rate-in-mbps</i> — Specifies the amount of bandwidth in Mbps to be reserved for the LSP path.

hop-limit

Syntax	hop-limit <i>limit</i> no hop-limit
Context	config>router>mpls>lsp>fast-reroute config>router>mpls>lsp-template>fast-reroute
Description	For fast reroute, how many more routers a detour is allowed to traverse compared to the LSP itself. For example, if an LSP traverses four routers, any detour for the LSP can be no more than ten router hops, including the ingress and egress routers.
Default	16
Parameters	<i>limit</i> — Specify the maximum number of hops.

Values 0 — 255

node-protect

Syntax	[no] node-protect
Context	config>router>mpls>lsp>fast-reroute config>router>mpls>lsp-template>fast-reroute
Description	This command enables or disables node and link protection on the specified LSP. Node protection ensures that traffic from an LSP traversing a neighboring router will reach its destination even if the neighboring router fails.
Default	node-protect

from

Syntax	from <i>ip-address</i>
Context	config>router>mpls>lsp config>router>mpls>lsp-template
Description	<p>This optional command specifies the IP address of the ingress router for the LSP. When this command is not specified, the system IP address is used. IP addresses that are not defined in the system are allowed. If an invalid IP address is entered, LSP bring-up fails and an error is logged.</p> <p>If an interface IP address is specified as the from address, and the egress interface of the nexthop IP address is a different interface, the LSP is not signaled. As the egress interface changes due to changes in the routing topology, an LSP recovers if the from IP address is the system IP address and not a specific interface IP address.</p> <p>Only one from address can be configured.</p>
Default	The system IP address
Parameters	<p><i>ip-address</i> — This is the IP address of the ingress router. This can be either the interface or the system IP address. If the IP address is local, the LSP must egress through that local interface which ensures local strictness.</p> <p>Default System IP address</p> <p>Values System IP or network interface IP addresses</p>

hop-limit

Syntax	hop-limit <i>number</i> no hop-limit
Context	config>router>mpls>lsp config>router>mpls>lsp>fast-reroute config>router>mpls>lsp-template>fast-reroute
Description	<p>This command specifies the maximum number of hops that an LSP can traverse, including the ingress and egress routers. An LSP is not set up if the hop limit is exceeded. This value can be changed dynamically for an LSP that is already set up with the following implications:</p> <p style="padding-left: 40px;">If the new value is less than the current number of hops of the established LSP, the LSP is brought down. Software then tries to re-establish the LSP within the new hop-limit number. If the new value is equal to or greater than the current number hops of the established LSP, then the LSP is not affected.</p> <p>The no form of this command returns the parameter to the default value.</p>
Default	255
Parameters	<p><i>number</i> — The number of hops the LSP can traverse, expressed as an integer.</p> <p>Values 2 — 255</p> <p>Values 0 — 255</p>

ldp-over-rsvp

Syntax	ldp-over-rsvp [include exclude]
Context	config>router>mpls>lsp config>router>mpls>lsp-template
Description	This command specifies if this LSP will be included in LDP over RSVP.
Parameters	include — Specifies that this LSP will be included in LDP over RSVP. exclude — Specifies that this LSP will be excluded from LDP over RSVP.

igp-shortcut

Syntax	igp-shortcut [lfa-protect lfa-only] [relative-metric [offset]] [no] igp-shortcut
Context	config>router>mpls>lsp config>router>mpls>lsp-template
Description	<p>This command enables the use of a specific RSVP LSP by IS-IS and OSPF routing protocols as a shortcut or as a forwarding adjacency for resolving IGP routes.</p> <p>When the rsvp-shortcut or the advertise-tunnel-link option is enabled at the IGP instance level, all RSVP LSPs originating on this node are eligible by default as long as the destination address of the LSP, as configured in config>router>mpls>lsp>to, corresponds to a router-id of a remote node.</p> <p>The lfa-protect option allows an LSP to be included in both the main SPF and the Loop-Free Alternate (LFA) SPF. For a given prefix, the LSP can be used either as a primary next-hop or as an LFA next-hop, but not both. If the main SPF computation selected a tunneled primary next-hop for a prefix, the LFA SPF will not select an LFA next-hop for this prefix and the protection of this prefix will rely on the RSVP LSP FRR protection. If the main SPF computation selected a direct primary next-hop, then the LFA SPF will select an LFA next-hop for this prefix but will prefer a direct LFA next-hop over a tunneled LFA next-hop.</p> <p>The lfa-only option allows an LSP to be included in the LFA SPF only such that the introduction of IGP shortcuts does not impact the main SPF decision. For a given prefix, the main SPF always selects a direct primary next-hop. The LFA SPF will select an LFA next-hop for this prefix but will prefer a direct LFA next-hop over a tunneled LFA next-hop.</p> <p>When the relative-metric option is enabled, IGP will apply the shortest IGP cost between the endpoints of the LSP plus the value of the offset (instead of the LSP operational metric) when computing the cost of a prefix which is resolved to the LSP. The offset value is optional and it defaults to zero. The minimum net cost for a prefix is one (1) after applying the offset. Note that the TTM continues to show the LSP operational metric as provided by MPLS. In other words, applications such as LDP-over-RSVP (when IGP shortcut is disabled) and BGP and static route shortcuts will continue to use the LSP operational metric.</p> <p>The relative-metric option is mutually exclusive with the lfa-protect or the lfa-only options. In other words, an LSP with the relative-metric option enabled cannot be included in the LFA SPF and vice-versa when the rsvp-shortcut option is enabled in the IGP.</p>

Finally, the **relative-metric** option is ignored when forwarding adjacency is enabled in IS-IS or OSPF. In this case, IGP advertises the LSP as a point-to-point unnumbered link along with the LSP operational metric as returned by MPLS and capped to maximum link metric allowed in that IGP. Both the main SPF and the LFA SPFs will use the local IGP database to resolve the routes.

The **no** form of this command disables the use of a specific RSVP LSP by IS-IS and OSPF routing protocols as a shortcut or a forwarding adjacency for resolving IGP routes.

Default igp-shortcut. All RSVP LSPs originating on this node are eligible by default as long as the destination address of the LSP corresponds to a router-id of a remote node.

Parameters **lfa-protect** — An LSP is included in both the main SPF and the LFA SPF.

lfa-only — An LSP is included in the LFA SPF only.

relative-metric [*offset*] — The shortest IGP cost between the endpoints of the LSP plus the configured offset, instead of the LSP operational metric returned by MPLS, is used when calculating the cost of prefix resolved to this LSP. The offset parameter is an integer and is optional. An offset value of zero is used when the relative-metric option is enabled without specifying the offset parameter value.

Values [-10, +10]

least-fill

Syntax [no] least-fill

Context config>router>mpls>lsp
config>router>mpls>lsp-template

Description This command enables the use of the least-fill path selection method for the computation of the path of this LSP.

When MPLS requests the computation of a path for this LSP, CSPF will find all equal cost shortest paths which satisfy the constraints of this path. Then, CSPF identifies the single link in each of these paths which has the least available bandwidth as a percentage of its maximum reservable bandwidth. It then selects the path which has the largest value of this percentage least available bandwidth figure. CSPF identifies the least available bandwidth link in each equal cost path after it has accounted for the bandwidth of the new requested path of this LSP.

CSPF applies the least-fill path selection method to all requests for a path, primary and secondary, of an LSP for which this option is enabled. The bandwidth of the path can be any value, including zero.

CSPF applies the least-fill criterion separately to each pre-emption priority in the base TE. A higher setup priority path can pre-empt lower holding priority paths.

CSPF also applies the least-fill criterion separately to each Diff-Serv TE class if Diff-Serv TE is enabled on this node. A higher setup priority path can pre-empt lower holding priority paths within a Class Type.

MPLS will re-signal and move the LSP to the new path in the following cases:

- Initial LSP path signaling.
- Re-try of an LSP path after failure.
- Make-before-break (MBB) due to pending soft pre-emption of the LSP path.

- MBB due to LSP path configuration change, i.e., a user change to bandwidth parameter of primary or secondary path, or a user enabling of fast-reroute option for the LSP.
- MBB of secondary path due to an update to primary path SRLG.
- MBB due to FRR Global Revertive procedures on the primary path.
- Manual re-signaling of an LSP path or of all LSP paths by the user.

During a manual re-signaling of an LSP path, MPLS will always re-signal the path regardless of whether the new path is exactly the same or different than the current path and regardless of whether the metric of the new path is different or not from that of the current path.

During a timer-based re-signaling of an LSP path which has the least-fill option enabled, MPLS will only re-signal the path if the metric of the new path is different than the one of the current path.

The user deletes a specific node entry in this database by executing the no form of this command.

Default no least-fill. The path of an LSP is randomly chosen among a set of equal cost paths.

ldp-over-rsvp

Syntax `[no] ldp-over-rsvp [include | exclude]`

Context config>router>mpls>lsp
config>router>mpls>lsp-template

Description This command configures an LSP so that it can be used by the IGP to calculate its SPF tree. The IGP (OSPF/ISIS) will subsequently provide LDP with all ECMP IGP next-hops and tunnel endpoints that it considers to be the lowest cost path to its destination.

If an IGP calculation and an LDP-over-RSVP indicate the same cost then LDP will always prefer an LDP-over-RSVP tunnel over an IGP route and ECMP between the two types is not considered.

The type and number of tunnels considered by LDP depends on the IGP metrics (the lowest metric between the tunnel endpoint and the target is selected) assuming that each LSP has a TLDP session established between the endpoints.

Enter the command **ldporsvp include** to make the associated LSP available to be used by the LDP-over-RSVP feature.

The no form of the command reverts to default operation.

Default ldporsvp exclude

include

Syntax	[no] include <i>group-name</i> [<i>group-name...</i> (up to 5max)]
Context	config>router>mpls>lsp config>router>mpls>lsp>primary config>router>mpls>lsp>secondary config>router>mpls>lsp-template
Description	This command specifies the admin groups to be included when an LSP is set up. Up to 5 groups per operation can be specified, up to 32 maximum. The no form of the command deletes the specified groups in the specified context.
Default	no include
Parameters	<i>group-name</i> — Specifies admin groups to be included when an LSP is set up.

priority

Syntax	priority setup-priority <i>hold-priority</i> no priority
Context	config>router>mpls>lsp>primary config>router>mpls>lsp>secondary
Description	<p>This command enables the soft pre-emption procedures for this LSP path. The operator enables the soft pre-emption mechanism on a specific LSP name by explicitly configuring the setup and holding priorities for the primary path at the 7x50 head-end node. The operator can similarly configure priority values for a secondary path for this LSP name. Different values could be used for the primary and for any of the secondary paths. In the absence of explicit user configuration, the setup priority is internally set to the default value of 7 and the holding priority is set to the default value of 0. Note however that valid user-entered values for these two parameters require that the holding priority be numerically lower than or equal to the setup priority, otherwise pre-emption loops can occur.</p> <p>Pre-emption is effected when a 7x50 pre-empting node processes a new RSVP session reservation and there is not enough available bandwidth on the RSVP interface, or the Class Type (CT) when Diff-Serv is enabled, to satisfy the bandwidth in the Flowspec object while there exist other session reservations for LSP paths with a strictly lower holding priority (numerically higher holding priority value) than the setup priority of the new LSP reservation. If enough available bandwidth is freed on the link or CT to accommodate the new reservation by pre-empting one or more lower priority LSP paths, the pre-empting node allows temporary overbooking of the RSVP interface and honors the new reservation.</p> <p>The 7x50 pre-empting node will immediately set the 'Preemption pending' flag (0x10) in the IPv4 Sub-Object in the RRO object in the Resv refresh for each of the pre-empted LSP paths. The IPv4 Sub-Object corresponds to the outgoing interface being used by the pre-empting and pre-empted LSP paths. Note however that the bandwidth value in the Flowspec object is not changed. The Resv flag must also be set if the pre-empting node is a merge point for the primary LSP path and the backup bypass LSP or detour LSP and the backup LSP is activated.</p> <p>When evaluating if enough available bandwidth will be freed, the 7x50 pre-empting node considers the reservations in order from the lowest holding priority (numerically higher holding priority value)</p>

to the holding priority just below the setup priority of the new reservation. A new reservation cannot pre-empt a reservation which has a value of the holding priority equal to the new reservation setup priority.

When Diff-Serv is enabled on the pre-empting node and the MAM bandwidth allocation model is used, a new reservation can only pre-empt a reservation in the same Class Type (CT).

LSP paths which were not flagged at the head-end for soft pre-emption will be hard pre-empted. LSP paths with the default holding priority of 0 cannot be pre-empted. LSP paths with zero bandwidth do not pre-empt other LSP paths regardless of the values of the path setup priority and the path holding priority. They can also not be pre-empted.

When evaluating if enough available bandwidth will be freed, the 7x50 pre-empting node considers the reservations in order from the lowest holding priority (numerically higher holding priority) to the holding priority just below the setup priority of the new reservation. There is no specific order in which the reservations in the same holding priority are considered. Furthermore, LSP paths which were not flagged at the head-end for soft pre-emption cannot be pre-empted because their holding priority is set internally to 0.

The 7x50 pre-empting node starts a preemption timer for each of the pre-empted LSP paths. While this timer is on, the node should continue to refresh the Path and Resv for the pre-empted LSP paths. When the preemption timer expires, the node tears down the reservation if the head-end node has not already done so.

A 7x50 head-end node upon receipt of the Resv refresh message with the 'Preemption pending' flag must immediately perform a make-before-break on the affected adaptive CSPF LSP. Both IGP metric and TE metric based CSPF LSPs are included. If an alternative path that excludes the flagged interface is not found, then the LSP is put on a retry in a similar way to the Global Revertive procedure at a 7x50 head-end node. However, the number of retries and the retry timer are governed by the values of the retry-limit and retry-timer parameters: `config>router>mpls>lsp>retry-limit;` `config>router>mpls>lsp>retry-timer.`

Note that MPLS will keep the address list of flagged interfaces for a maximum of 60 seconds (not user-configurable) from the time the first Resv message with the 'Preemption pending' flag is received. This actually means that MPLS will request CSPF to find a path that excludes the flagged interfaces in the first few retries until success or until 60 seconds have elapsed. Subsequent retries after the 60 seconds will not exclude the flagged interfaces as it is assumed IGP has converged by then and the Unreserved Bandwidth sub-TLV for that priority, or TE Class, in the TE database will show the updated value taking into account the pre-empting LSP path reservation or a value of zero if overbooked.

If the LSP has a configured secondary standby which is operationally UP, the 7x50 will switch the path of the LSP to it and then start the MBB. If no standby path is available and a secondary non-standby is configured, the 7x50 will start the MBB and signal the path of the secondary. The LSP path will be switched to either the secondary or the new primary, whichever comes up first.

The no form of the command reverts the LSP path priority to the default values and results in setting the setup priority to 7, in setting the holding priority to 0, and in clearing the 'soft preemption desired' flag in the RRO in the Resv refresh message.

Default	no priority.
Parameters	<i>setup-priority</i> — The priority of the reservation for this session at setup time.
Values	0 — 7 (0 is the highest priority and 7 is the lowest priority.)
Default	7 — This session does not pre-empt any other session.

holding-priority — The priority of the reservation for this session at pre-emption action.

Values 0 — 7 (0 is the highest priority and 7 is the lowest priority.)

Default 0 — This session does not get pre-empted by any other session.

main-ct-retry-limit

Syntax	main-ct-retry-limit <i>number</i> no main-ct-retry-limit
Context	config>router>mpls>lsp
Description	<p>This command configures the maximum number of retries the LSP primary path should be retried with the LSP Diff-Serv main Class Type (CT).</p> <p>When an unmapped LSP primary path goes into retry, it uses the main CT until the number of retries reaches the value of the new main-ct-retry-limit parameter. If the path did not come up, it must start using the backup CT at that point in time. By default, this parameter is set to infinite value. The new main-ct-retry-limit parameter has no effect on an LSP primary path which retries due to a failure event.</p> <p>An unmapped LSP primary path is a path which has never received a Resv in response to the first Path message sent. This can occur when performing a “shut/no-shut” on the LSP or LSP primary path or when the node reboots. An unmapped LSP primary path goes into retry if the retry timer expired or the head-end node received a PathErr message before the retry timer expired.</p> <p>If the user entered a value of the main-ct-retry-limit parameter that is greater than the value of the LSP retry-limit, the number of retries will still stop when the LSP primary path reaches the value of the LSP retry-limit. In other words, the meaning of the LSP retry-limit parameter is not changed and always represents the upper bound on the number of retries. The unmapped LSP primary path behavior applies to both CSPF and non-CSPF LSPs.</p> <p>The no form of this command sets the parameter to the default value of zero (0) which means the LSP primary path will retry forever.</p>
Default	no main-ct-retry-limit
Parameters	<p><i>number</i> — The number of times MPLS will attempt to re-establish the LSP primary path using the Diff-Serv main CT. Allowed values are integers in the range of zero (0) to 10,000, where zero indicates to retry infinitely.</p> <p>Values 0-1000, integer</p>

metric

Syntax	[no] metric <i>metric</i>
Context	config>router>mpls>lsp config>router>mpls>lsp-template
Description	This command allows the user to override the LSP operational metric with a constant administrative value that will not change regardless of the actual path the LSP is using over its lifetime.

The LSP operational metric will match the metric the active path of this LSP is using at any given time. For a CSPF LSP, this metric represents the cumulative IGP metric of all the links the active path is using. If CSPF for this LSP is configured to use the TE metric, the LSP operational metric is set to the maximum value. For a non-CSPF LSP, the operational metric is the shortest IGP cost to the destination of the LSP.

The LSP operational metric is used by some applications to select an LSP among a set of LSPs that are destined to the same egress router. The LSP with the lowest operational metric will be selected. If more than one LSP with the same lowest LSP metric exists, the LSP with the lowest tunnel index will be selected. The configuration of a constant metric by the user will make sure the LSP always maintains its preference in this selection regardless of the path it is using at any given time. Applications that use the LSP operational metric include LDP-over-RSVP, VPRN auto-bind, and IGP, BGP and static route shortcuts.

The **no** form of this command disables the administrative LSP metric and reverts to the default setting in which the metric value will represent the LSP metric returned by MPLS. The same behavior is obtained if the user entered a metric of value zero (0).

Default no metric. The LSP operational metric defaults to the metric returned by MPLS.

Parameters *metric* — Specifies the integer value which specifies the value of the LSP administrative metric. A value of zero command reverts to the default setting and disables the administrative LSP metric.

Values 0— 16777215

to

Syntax to [*ip-address* | **node-id** [*a.b.c.d* | 1...4,294,967,295]]

Context config>router>mpls>lsp

Description This command specifies the system IP address or MPLS-TP node-id of the egress router for the LSP. This command is mandatory to create an LSP.

An IP address for which a route does not exist is allowed in the configuration. If the LSP signaling fails because the destination is not reachable, an error is logged and the LSP operational status is set to down.

For a non MPLS-TP LSP, the **to ip-address** *must* be the system IP address of the egress router. If the **to** address does not match the SDP address, the LSP is not included in the SDP definition.

For an MPLS-TP LSP, the **to node-id** may be either in 4-octet IPv4 address format, or a 32bit unsigned integer. This command is mandatory to create an MPLS-TP LSP. Note that a value of zero is invalid. This to address is used in the MPLS-TP LSP ID, and the MPLS-TP MEP ID for the LSP.

Default No default

Parameters *ip-address* — The system IP address of the egress router.

node-id *a.b.c.d. | 1...4,294,967,295* — 4-octet IPv4 formatted or unsigned 32-bit integer MPLS-TP node-id of the egress router.

propagate-admin-group

Syntax	[no] propagate-admin-group
Context	config>router>mpls>lsp config>router>mpls>lsp-template
Description	<p>This command enables propagation of session attribute object with resource affinity (C-type 1) in PATH message. If a session attribute with resource affinity is received at an LSR, then it will check the compatibility of admin-groups received in PATH message against configured admin-groups on the egress interface of LSP.</p> <p>To support admin-group for inter-area LSP, the ingress node must configure propagating admin-groups within the session attribute object. If a PATH message is received by an LSR node that has the cspf-on-loose option enabled and the message includes admin-groups, then the ERO expansion by CSPF to calculate the path to the next loose hop will include the admin-group constraints received from ingress node.</p> <p>If this option is disabled, then the session attribute object without resource affinity (C-Type 7) is propagated in PATH message and CSPF at the LSR node will not include admin-group constraints.</p> <p>This admin group propagation is supported with a P2P LSP, a P2MP LSP instance, and an LSP template.</p> <p>The user can change the value of the propagate-admin-group option on the fly. A RSVP P2P LSP will perform a Make-Before-Break (MBB) on changing the configuration. A S2L path of an RSVP P2MP LSP will perform a Break-Before-Make on changing the configuration.</p>
Default	no propagate-admin-group

vprn-auto-bind

Syntax	vprn-auto-bind [include exclude]
Context	config>router>mpls>lsp config>router>mpls>lsp-template
Description	<p>This command determines whether the associated names LSP can be used or no as part of the auto-bind feature for VPRN services. By default a names LSP is available for inclusion to used for the auto-bind feature.</p> <p>By configuring the command vprn-auto-bind exclude, the associated LSP will not be used by the auto-bind feature within VPRN services.</p> <p>The no form of the command resets the flag backto the default value.</p>
Default	include
Parameters	<p>include — Allows an associated LSPto be used by auto-bin for vprn services</p> <p>exclude — Disables the use of the associated LSP to be used with the auto-bind feature for VPRN services.</p>

retry-limit

Syntax	retry-limit <i>number</i> no retry-limit
Context	config>router>mpls>lsp config>router>mpls>lsp-template
Description	<p>This optional command specifies the number of attempts software should make to re-establish the LSP after it has failed LSP. After each successful attempt, the counter is reset to zero.</p> <p>When the specified number is reached, no more attempts are made and the LSP path is put into the shutdown state.</p> <p>Use the config router mpls lsp <i>lsp-name</i> no shutdown command to bring up the path after the retry-limit is exceeded.</p> <p>For P2MP LSP created based on LSP template, all S2Ls must attempt to retry-limit before client application is informed of failure.</p> <p>The no form of this command revert the parameter to the default value.</p>
Default	0 (no limit, retries forever)
Parameters	<i>number</i> — The number of times software will attempt to re-establish the LSP after it has failed. Allowed values are integers in the range of 0 to 10000 where 0 indicates to retry forever. Values 0 — 10000

retry-timer

Syntax	retry-timer <i>seconds</i> no retry-timer
Context	config>router>mpls>lsp config>router>mpls>lsp-template
Description	<p>This command configures the time, in seconds, for LSP re-establishment attempts after it has failed. The retry time is jittered to +/- 25% of its nominal value.</p> <p>For P2MP LSP created based on LSP template, all S2Ls must attempt to retry-limit before client application is informed of failure.</p> <p>The no form of this command reverts to the default value.</p>
Default	30
Parameters	<i>seconds</i> — The amount of time, in seconds, between attempts to re-establish the LSP after it has failed. Allowed values are integers in the range of 1 to 600. Values 1 — 600

rsvp-resv-style

Syntax	rsvp-resv-style [<i>se</i> <i>ff</i>]
Context	config>router>mpls>lsp config>router>mpls>lsp-template
Description	This command specifies the RSVP reservation style, shared explicit (<i>se</i>) or fixed filter (<i>ff</i>). A reservation style is a set of control options that specify a number of supported parameters. The style information is part of the LSP configuration.
Default	<i>se</i>
Parameters	<p><i>ff</i> — Fixed filter is single reservation with an explicit scope. This reservation style specifies an explicit list of senders and a distinct reservation for each of them. A specific reservation request is created for data packets from a particular sender. The reservation scope is determined by an explicit list of senders.</p> <p><i>se</i> — Shared explicit is shared reservation with a limited scope. This reservation style specifies a shared reservation environment with an explicit reservation scope. This reservation style creates a single reservation over a link that is shared by an explicit list of senders. Because each sender is explicitly listed in the RESV message, different labels can be assigned to different sender-receiver pairs, thereby creating separate LSPs.</p>

shutdown

Syntax	[no] shutdown
Context	config>router>mpls>lsp config>router>mpls>lsp-template
Description	<p>This command disables the existing LSP including the primary and any standby secondary paths. To shutdown only the primary enter the config router mpls lsp <i>lsp-name</i> primary <i>path-name</i> shutdown command.</p> <p>To shutdown a specific standby secondary enter the config router mpls lsp <i>lsp-name</i> secondary <i>path-name</i> shutdown command. The existing configuration of the LSP is preserved.</p> <p>Use the no form of this command to restart the LSP. LSPs are created in a shutdown state. Use this command to administratively bring up the LSP.</p>
Default	shutdown

lsp-template

Syntax	[no] lsp-template <i>lsp-template-name</i> p2mp-lsp
Context	config>router>mpls
Description	This command creates a template construct that can be referenced by client application where dynamic LSP creation is required. 'p2mp-lsp' keyword is mandatory.

The **no** form of command deletes LSP template. LSP template cannot be deleted if a client application is using it.

Default none

Parameters *lsp-template-name* — Name to identify LSP template. Any LSP template name and LSP name must not be same.

default-path

Syntax **[no] default-path** *path-name*

Context config>router>mpls>lsp-template

Description A default path binding must be provided before LSP template can be used for signaling LSP. LSP template must be shutdown to modify default-path binding.

The **no** form of command should delete path binding.

Default none

Parameters *path-name*

Primary and Secondary Path Commands

primary

Syntax	primary <i>path-name</i> no primary
Context	config>router>mpls>lsp
Description	<p>This command specifies a preferred path for the LSP. This command is optional only if the secondary <i>path-name</i> is included in the LSP definition. Only one primary path can be defined for an LSP.</p> <p>Some of the attributes of the LSP such as the bandwidth, and hop-limit can be optionally specified as the attributes of the primary path. The attributes specified in the primary path <i>path-name</i> command, override the LSP attributes.</p> <p>The no form of this command deletes the association of this <i>path-name</i> from the LSP <i>lsp-name</i>. All configurations specific to this primary path, such as record, bandwidth, and hop limit, are deleted. The primary path must be shutdown first in order to delete it. The no primary command will not result in any action except a warning message on the console indicating that the primary path is administratively up.</p>
Default	none
Parameters	<i>path-name</i> — The case-sensitive alphanumeric name label for the LSP path up to 32 characters in length.

secondary

Syntax	[no] secondary <i>path-name</i>
Context	config>router>mpls>lsp
Description	<p>This command specifies an alternative path that the LSP uses if the primary path is not available. This command is optional and is not required if the config router mpls lsp <i>lsp-name</i> primary <i>path-name</i> command is specified. After the switch over from the primary to the secondary, the software continuously tries to revert to the primary path. The switch back to the primary path is based on the retry-timer interval.</p> <p>Up to eight secondary paths can be specified. All the secondary paths are considered equal and the first available path is used. The software will not switch back among secondary paths.</p> <p>Software starts the signaling of all non-standby secondary paths at the same time. Retry counters are maintained for each unsuccessful attempt. Once the retry limit is reached on a path, software will not attempt to signal the path and administratively shuts down the path. The first successfully established path is made the active path for the LSP.</p> <p>The no form of this command removes the association between this <i>path-name</i> and <i>lsp-name</i>. All specific configurations for this association are deleted. The secondary path must be shutdown first in</p>

order to delete it. The **no secondary** *path-name* command will not result in any action except a warning message on the console indicating that the secondary path is administratively up.

Default none

Parameters *path-name* — The case-sensitive alphanumeric name label for the LSP path up to 32 characters in length.

adaptive

Syntax [no] **adaptive**

Context config>router>mpls>lsp>primary
config>router>mpls>lsp>secondary

Description This command enables the make-before-break functionality for an LSP or a primary or secondary LSP path. When enabled for the LSP, make-before-break will be performed for primary path and all the secondary paths of the LSP.

Default adaptive

backup-class-type

Syntax **backup-class-type** *ct-number*
no backup-class-type

Context config>router>mpls>lsp>primary

Description This command enables the use of the Diff-Serv backup Class-Type (CT), instead of the Diff-Serv main CT, to signal the LSP primary path when it fails and goes into retry. The Diff-Serv main CT is configured at the LSP level or at the primary path level using the following commands:

```
config>router>mpls>lsp>class-type ct-number
```

```
config>router>mpls>lsp>primary>class-type ct-number
```

When a LSP primary path retries due a failure, for example, it fails after being in the UP state, or undergoes any type of Make-Before-Break (MBB), MPLS will retry a new path for the LSP using the main CT. If the first attempt failed, the head-end node performs subsequent retries using the backup CT. This procedure must be followed regardless if the currently used CT by this path is the main or backup CT. This applies to both CSPF and non-CSPF LSPs.

The triggers for using the backup CT after the first retry attempt are:

1. A local interface failure or a control plane failure (hello timeout etc.).
2. Receipt of a PathErr message with a notification of a FRR protection becoming active downstream and/or Receipt of a Resv message with a 'Local-Protection-In-Use' flag set. This invokes the FRR Global Revertive MBB.
3. Receipt of a PathErr message with error code=25 ("Notify") and sub-code=7 ("Local link maintenance required") or a sub-code=8 ("Local node maintenance required"). This invokes the TE Graceful Shutdown MBB.

4. Receipt of a Resv refresh message with the 'Preemption pending' flag set or a PathErr message with error code=34 ("Reroute") and a value=1 ("Reroute request soft preemption"). This invokes the soft preemption MBB.
5. Receipt of a ResvTear message.
6. A configuration change MBB.
7. The user executing the `clear>router>mpls>lsp` command.

When an unmapped LSP primary path goes into retry, it uses the main CT until the number of retries reaches the value of the new **main-ct-retry-limit** parameter. If the path did not come up, it must start using the backup CT at that point in time. By default, this parameter is set to infinite value. The new **main-ct-retry-limit** parameter has no effect on an LSP primary path which retries due to a failure event.

An unmapped LSP primary path is a path which has never received a Resv in response to the first Path message sent. This can occur when performing a 'shut/no-shut' on the LSP or LSP primary path or when the node reboots. An unmapped LSP primary path goes into retry if the retry timer expired or the head-end node received a PathErr message before the retry timer expired.

When the re-signal timer expires, CSPF will try to find a path with the main CT. The head-end node must re-signal the LSP even if the new path found by CSPF is identical to the existing one since the idea is to restore the main CT for the primary path. A path with main CT is not found, the LSP remains on its current primary path using the backup CT.

When the user performs a manual re-signal of the primary path, CSPF will try to find a path with the main CT. The head-end node must re-signal the LSP as in current implementation.

The **no** form of this command disables the use of the Diff-Serv backup CT.

Default	no backup-class-type
Parameters	<i>ct-number</i> — The Diff-Serv Class Type number. One or more system forwarding classes can be mapped to a CT.
Values	0-7, integer

bandwidth

Syntax	bandwidth <i>rate-in-mbps</i> no bandwidth
Context	config>router>mpls>lsp>primary config>router>mpls>lsp>secondary config>router>mpls>lsp>template>fast-reroute
Description	This command specifies the amount of bandwidth to be reserved for the LSP path. The no form of this command resets bandwidth parameters (no bandwidth is reserved).
Default	no bandwidth (bandwidth setting in the global LSP configuration)
Parameters	<i>rate-in-mbps</i> — The amount of bandwidth reserved for the LSP path in Mbps. Allowed values are integers in the range of 1 to 100000.
Values	0 — 100000

exclude

Syntax	[no] exclude <i>group-name</i> [<i>group-name...</i> (up to 5 max)]
Context	config>router>mpls>lsp>primary config>router>mpls>lsp>secondary
Description	This command specifies the admin groups to be excluded when an LSP is set up. . Up to 5 groups per operation can be specified, up to 32 maximum. The admin groups are defined in the config>router>mpls>admin-group context. Use the no form of the command to remove the exclude command.
Default	no exclude
Parameters	<i>group-name</i> — Specifies the existing group-name to be excluded when an LSP is set up.

hop-limit

Syntax	hop-limit <i>number</i> no hop-limit
Context	config>router>mpls>lsp>primary config>router>mpls>lsp>secondary
Description	This optional command overrides the config router mpls lsp lsp-name hop-limit command. This command specifies the total number of hops that an LSP traverses, including the ingress and egress routers. This value can be changed dynamically for an LSP that is already set up with the following implications: If the new value is less than the current hops of the established LSP, the LSP is brought down. MPLS then tries to re-establish the LSP within the new hop-limit number. If the new value is equal or more than the current hops of the established LSP then the LSP will be unaffected. The no form of this command reverts the values defined under the LSP definition using the config router mpls lsp lsp-name hop-limit command.
Default	no hop-limit
Parameters	<i>number</i> — The number of hops the LSP can traverse, expressed as an integer. Values 2 — 255

record

Syntax	[no] record
Context	config>router>mpls>lsp>primary config>router>mpls>lsp>secondary config>router>mpls>lsp>template
Description	<p>This command enables recording of all the hops that an LSP path traverses. Enabling record increases the size of the PATH and RESV refresh messages for the LSP since this information is carried end-to-end along the path of the LSP. The increase in control traffic per LSP may impact scalability.</p> <p>The no form of this command disables the recording of all the hops for the given LSP. There are no restrictions as to when the no command can be used. The no form of this command also disables the record-label command.</p>
Default	record

record-label

Syntax	[no] record-label
Context	config>router>mpls>lsp>primary config>router>mpls>lsp>secondary config>router>mpls>lsp>template
Description	<p>This command enables recording of all the labels at each node that an LSP path traverses. Enabling the record-label command will also enable the record command if it is not already enabled.</p> <p>The no form of this command disables the recording of the hops that an LSP path traverses.</p>
Default	record-label

srlg

Syntax	[no] srlg
Context	config>router>mpls>lsp>secondary
Description	<p>This command enables the use of the SRLG constraint in the computation of a secondary path for an LSP at the head-end LER.</p> <p>When this feature is enabled, CSPF includes the SRLG constraint in the computation of the secondary LSP path. This requires that the primary LSP already be established and is up since the head-end LER needs the most current ERO computed by CSPF for the primary path. CSPF would return the list of SRLG groups along with the ERO during primary path CSPF computation. At a subsequent establishment of a secondary path with the SRLG constraint, the MPLS/RSVP task will query again CSPF providing the list of SLRG group numbers to be avoided. CSPF prunes all links with interfaces which belong to the same SRLGs as the interfaces included in the ERO of the primary</p>

path. If CSPF finds a path, the secondary is setup. If not, MPLS/RSVP will keep retrying the requests to CSPF.

If CSPF is not enabled on the LSP name, then a secondary path of that LSP which has the SRLG constraint included will be shut down and a specific failure code will indicate the exact reason for the failure in **show>router>mpls>lsp>path>detail** output.

At initial primary LSP path establishment, if primary does not come up or primary is not configured, SRLG secondary will not be signaled and will put to down state. A specific failure code will indicate the exact reason for the failure in **show>router>mpls>lsp>path>detail** output. However, if a non-SRLG secondary path was configured, such as a secondary path with the SRLG option disabled, MPLS/RSVP task will signal it and the LSP use it.

As soon as the primary path is configured and successfully established, MPLS/RSVP moves the LSP to the primary and signals all SRLG secondary paths.

Any time the primary path is re-optimized, has undergone MBB, or has come back up after being down, MPLS/RSVP task checks with CSPF if the SRLG secondary should be re-signaled. If MPLS/RSVP finds that current secondary path is no longer SRLG disjoint, for example, it became ineligible, it puts it on a delayed MBB immediately after the expiry of the retry timer. If MBB fails at the first try, the secondary path is torn down and the path is put on retry.

At the next opportunity the primary goes down, the LSP will use the path of an eligible SRLG secondary if it is UP. If all secondary eligible SLRG paths are Down, MPLS/RSVP will use a non SRLG secondary if configured and UP. If while the LSP is using a non SRLG secondary, an eligible SRLG secondary came back up, MPLS/RSVP will not switch the path of the LSP to it. As soon as primary is re-signaled and comes up with a new SLRG list, MPLS/RSVP will re-signal the secondary using the new SRLG list.

A secondary path which becomes ineligible as a result of an update to the SRLG membership list of the primary path will have the ineligibility status removed on any of the following events:

8. A successful MBB of the standby SRLG path which makes it eligible again.
9. The standby path goes down. MPLS/RSVP puts the standby on retry at the expiry of the retry timer. If successful, it becomes eligible. If not successful after the retry-timer expired or the number of retries reached the number configured under the retry-limit parameter, it is left down.
10. The primary path goes down. In this case, the ineligible secondary path is immediately torn down and will only be re-signaled when the primary comes back up with a new SRLG list.

Once primary path of the LSP is setup and is operationally up, any subsequent changes to the SRLG group membership of an interface the primary path is using would not be considered until the next opportunity the primary path is re-signaled. The primary path may be re-signaled due to a failure or to a make-before-break operation. Make-before-break occurs as a result of a global revertive operation, a timer based or manual re-optimization of the LSP path, or an operator change to any of the path constraints.

One an SRLG secondary path is setup and is operationally UP, any subsequent changes to the SRLG group membership of an interface the secondary path is using would not be considered until the next opportunity secondary path is re-signaled. The secondary path is re-signaled due to a failure, to a re-signaling of the primary path, or to a make before break operation. Make-before break occurs as a result of a timer based or manual re-optimization of the secondary path, or an operator change to any of the path constraints of the secondary path, including enabling or disabling the SRLG constraint itself.

Also, the user-configured include/exclude admin group statements for this secondary path are also checked together with the SRLG constraints by CSPF. Finally, note that enabling SRPG on a secondary standby path that is in the up state will cause the path to be torn down and re-signaled using the SRLG constraint.

The **no** form of the command reverts to the default value.

Default no srlg

standby

Syntax [no] standby

Context config>router>mpls>lsp>secondary

Description The secondary path LSP is normally signaled once the primary path LSP fails. The **standby** keyword ensures that the secondary path LSP is signaled and maintained indefinitely in a hot-standby state. When the primary path is re-established then the traffic is switched back to the primary path LSP.

The **no** form of this command specifies that the secondary LSP is signaled when the primary path LSP fails.

Default none

path-preference

Syntax [no] path-preference *value*

Context config>router>mpls>lsp>secondary

Description This command enables use of path preference among configured standby secondary paths per LSP. If all standby secondary paths have a default path-preference value then a non-standby secondary path will remain the active path while a standby secondary is available. A standby secondary path configured with highest priority (lowest path-preference value) must be made the active path when the primary is not in use. Path preference can be configured on standby secondary path.

The **no** form of this command resets the path-preference to the default value.

Default 255

Parameters *value* — Specifies an alternate path for the LSP if the primary path is not available,
1–255

LSP Path Commands

hop

Syntax	hop <i>hop-index ip-address</i> { strict loose } no hop <i>hop-index</i>
Context	config>router>mpls>path
Description	<p>This command specifies the IP address of the hops that the LSP should traverse on its way to the egress router. The IP address can be the interface IP address or the system IP address. If the system IP address is specified then the LSP can choose the best available interface.</p> <p>Optionally, the LSP ingress and egress IP address can be included as the first and the last hop. A hop list can include the ingress interface IP address, the system IP address, and the egress IP address of any of the hops being specified.</p> <p>The no form of this command deletes hop list entries for the path. All the LSPs currently using this path are affected. Additionally, all services actively using these LSPs are affected. The path must be shutdown first in order to delete the hop from the hop list. The no hop hop-index command will not result in any action except a warning message on the console indicating that the path is administratively up.</p>
Default	none
Parameters	<p><i>hop-index</i> — The hop index is used to order the hops specified. The LSP always traverses from the lowest hop index to the highest. The hop index does not need to be sequential.</p> <p>Values 1 — 1024</p> <p><i>ip-address</i> — The system or network interface IP address of the transit router. The IP address can be the interface IP address or the system IP address. If the system IP address is specified then the LSP can choose the best available interface. A hop list can also include the ingress interface IP address, the system IP address, and the egress IP address of any of the specified hops.</p> <p>loose — This keyword specifies that the route taken by the LSP from the previous hop to this hop can traverse through other routers. Multiple hop entries with the same IP address are flagged as errors. Either the loose or strict keyword must be specified.</p> <p>strict — This keyword specifies that the LSP must take a direct path from the previous hop router to this router. No transit routers between the previous router and this router are allowed. If the IP address specified is the interface address, then that is the interface the LSP must use. If there are direct parallel links between the previous router and this router and if system IP address is specified, then any one of the available interfaces can be used by the LSP. The user must ensure that the previous router and this router have a direct link. Multiple hop entries with the same IP address are flagged as errors. Either the loose or strict keyword must be specified.</p>

path

Syntax	[no] path <i>path-name</i>
Context	config>router>mpls
Description	<p>This command creates the path to be used for an LSP. A path can be used by multiple LSPs. A path can specify some or all hops from ingress to egress and they can be either strict or loose. A path can also be empty (no <i>path-name</i> specified) in which case the LSP is set up based on IGP (best effort) calculated shortest path to the egress router. Paths are created in a shutdown state. A path must be shutdown before making any changes (adding or deleting hops) to the path. When a path is shutdown, any LSP using the path becomes operationally down.</p> <p>To create a strict path from the ingress to the egress router, the ingress and the egress routers must be included in the path statement.</p> <p>The no form of this command deletes the path and all its associated configuration information. All the LSPs that are currently using this path will be affected. Additionally all the services that are actively using these LSPs will be affected. A path must be shutdown and unbound from all LSPs using the path before it can be deleted. The no path <i>path-name</i> command will not result in any action except a warning message on the console indicating that the path may be in use.</p>
Parameters	<i>path-name</i> — Specify a unique case-sensitive alphanumeric name label for the LSP path up to 32 characters in length.

shutdown

Syntax	[no] shutdown
Context	config>router>mpls>path
Description	<p>This command disables the existing LSPs using this path. All services using these LSPs are affected. Binding information, however, is retained in those LSPs. Paths are created in the shutdown state.</p> <p>The no form of this command administratively enables the path. All LSPs, where this path is defined as primary or defined as standby secondary, are (re)established.</p>
Default	shutdown

Static LSP Commands

static-lsp

Syntax	[no] static-lsp <i>lsp-name</i>
Context	config>router>mpls
Description	<p>This command is used to configure a static LSP on the ingress router. The static LSP is a manually set up LSP where the nexthop IP address and the outgoing label (push) must be specified.</p> <p>The no form of this command deletes this static LSP and associated information.</p> <p>The LSP must be shutdown first in order to delete it. If the LSP is not shut down, the no static-lsp <i>lsp-name</i> command does nothing except generate a warning message on the console indicating that the LSP is administratively up.</p>
Parameters	<i>lsp-name</i> — Name that identifies the LSP.
Values	Up to 32 alphanumeric characters.

static-lsp-fast-retry

Syntax	static-lsp-fast-retry <i>seconds</i> [no] static-lsp-fast-retry
Context	config>router>mpls
Description	<p>This command specifies the value used as the fast retry timer for a static LSP.</p> <p>When a static LSP is trying to come up, the MPLS request for the ARP entry of the LSP next-hop may fail when it is made while the next-hop is still down or unavailable. In that case, MPLS starts a retry timer before making the next request. This enhancement allows the user to configure the retry timer, so that the LSP comes up as soon as the next-hop is up.</p> <p>The no form of the command reverts to the default.</p>
Default	no static-fast-retry-timer
Parameters	<i>seconds</i> — specifies the value, in seconds, used as the fast retry timer for a static LSP.
Values	1-30

push

Syntax	push { <i>label</i> implicit-null-label } nexthop <i>ip-address</i> no push { <i>out-label</i> implicit-null-label }
Context	config>router>mpls>static-lsp
Description	This command specifies the label to be pushed on the label stack and the next hop IP address for the static LSP. The no form of this command removes the association of the label to push for the static LSP.
Parameters	implicit-null-label — Specifies the use of the implicit label value for the push operation. <i>label</i> — The label to push on the label stack. Label values 16 through 1,048,575 are defined as follows: Label values 16 through 31 are reserved. Label values 32 through 1,023 are available for static assignment. Label values 1,024 through 2,047 are reserved for future use. Label values 2,048 through 18,431 are statically assigned for services. Label values 28,672 through 131,071 are dynamically assigned for both MPLS and services. Label values 131,072 through 1,048,575 are reserved for future use. Values 16 — 1048575 nexthop <i>ip-address</i> — This command specifies the IP address of the next hop towards the LSP egress router. If an ARP entry for the next hop exists, then the static LSP is marked operational. If ARP entry does not exist, software sets the operational status of the static LSP to down and continues to ARP for the configured nexthop. Software continuously tries to ARP for the configured nexthop at a fixed interval.

shutdown

Syntax	[no] shutdown
Context	config>router>mpls>static-lsp
Description	This command is used to administratively disable the static LSP. The no form of this command administratively enables the static LSP.
Default	shutdown

to

Syntax	to <i>ip-address</i>
Context	config>router>mpls>static-lsp
Description	This command specifies the system IP address of the egress router for the static LSP. When creating an LSP this command is required. For LSPs that are used as transport tunnels for services, the to IP address <i>must</i> be the system IP address. If the to address does not match the SDP address, the LSP is not included in the SDP definition.
Parameters	<i>ip-address</i> — The system IP address of the egress router.
Default	none

Point-to-Multipoint MPLS (P2MP) Commands

p2mp-id

Syntax	p2mp-id <i>id</i>
Context	config>router>mpls>lsp
Description	<p>This command configures the identifier of an RSVP P2MP LSP. An RSVP P2MP LSP is fully identified by the combination of: <P2MP ID, tunnel ID, extended tunnel ID> part of the P2MP session object, and <tunnel sender address, LSP ID> fields in the p2mp sender_template object.</p> <p>The p2mp-id is a 32-bit identifier used in the session object that remains constant over the life of the P2MP tunnel. It is unique within the scope of the ingress LER.</p> <p>The no form restores the default value of this parameter.</p>
Default	0
Parameters	<i>id</i> — Specifies a P2MP identifier.
Values	0 — 65535

primary-p2mp-instance

Syntax	[no] primary-p2mp-instance <i>instance-name</i>
Context	config>router>mpls>lsp
Description	<p>This command creates the primary instance of a P2MP LSP. The primary instance of a P2MP LSP is modeled as a set of root-to-leaf (S2L) sub-LSP's. The root, for example a head-end node triggers signaling using one path message per S2L path. The leaf sub-LSP paths are merged at branching points.</p>
Default	none
Parameters	<i>instance-name</i> — Specifies a name that identifies the P2MP LSP instance. The instance name can be up to 32 characters long and must be unique.

s2l-path

Syntax	[no] s2l-path <i>path-name to ip-address</i>
Context	config>router>mpls>lsp>primary-inst
Description	<p>This command creates a root-to-leaf (S2L) sub-LSP path for the primary instance of a P2MP LSP. The primary instance of a P2MP LSP is modeled as a set of root-to-leaf (S2L) sub-LSPs. The root, for example, head-end node, triggers signaling using one path message per S2L path. The leaf sub-LSP paths are merged at branching points.</p>

Each S2L sub-LSP is signaled in a separate path message. Each leaf node will respond with its own RESV message. A branch LSR node will forward the path message of each S2L sub-LSP to the downstream LSR without replicating it. It will also forward the RESV message of each S2L sub-LSP to the upstream LSR without merging it with the RESV messages of other S2L sub-LSPs of the same P2MP LSP. The same is done for subsequent refreshes of the path and RESV states.

The S2L paths can be empty paths or can specify a list of explicit hops. The path name must exist and must have been defined using the **config>router>mpls>path** command. The same path name can be re-used by more than one S2L of the primary P2MP instance. However, the **to** keyword must have a unique argument per S2L as it corresponds to the address of the egress LER node.

Default none

Parameters *path-name* — Specifies the name of the path which consists of up to 32 alphanumeric characters.
to *ip-address* — Specifies the system IP address of the egress router.

p2mp-resignal-timer

Syntax **p2mp-resignal-timer** *minutes*
no p2mp-resignal-timer

Context config>router>mpls

Description This command configures the re-signal timer for a P2MP LSP instance. MPLS will request CSPF to re-compute the whole set of S2L paths of a given active P2MP instance each time the P2MP re-signal timer expires. The P2MP re-signal timer is configured separately from the P2P LSP parameter. MPLS performs a global MBB and moves each S2L sub-LSP in the instance into its new path using a new P2MP LSP ID if the global MBB is successful, regardless of the cost of the new S2L path. The **no** form of this command disables the timer-based re-signaling of P2MP LSPs on this system.

Parameters *minutes* — Specifies the time MPLS waits before attempting to re-signal the P2MP LSP instance.

Values 60 — 10080

RSVP Configuration Commands

Generic Commands

shutdown

Syntax	[no] shutdown
Context	config>router>rsvp config>router>rsvp>interface
Description	<p>This command disables the RSVP protocol instance or the RSVP-related functions for the interface. The RSVP configuration information associated with this interface is retained. When RSVP is administratively disabled, all the RSVP sessions are torn down. The existing configuration is retained.</p> <p>The no form of this command administratively enables RSVP on the interface.</p>
Default	shutdown

RSVP Commands

rsvp

Syntax	[no] rsvp
Context	config>router
Description	<p>This command enables the context to configure RSVP protocol parameters. RSVP is not enabled by default and must be explicitly enabled (no shutdown).</p> <p>RSVP is used to set up LSPs. RSVP should be enabled on all router interfaces that participate in signaled LSPs.</p> <p>The no form of this command deletes this RSVP protocol instance and removes all configuration parameters for this RSVP instance. To suspend the execution and maintain the existing configuration, use the shutdown command. RSVP must be shutdown before the RSVP instance can be deleted. If RSVP is not shutdown, the no rsvp command does nothing except issue a warning message on the console indicating that RSVP is still administratively enabled.</p>
Default	no shutdown

diffserv-te

Syntax	diffserv-te [mam rdm] no diffserv-te
Context	config>router>rsvp
Description	<p>This command enabled Diff-Serv Traffic Engineering on the node.</p> <p>When this command is enabled, IS-IS and OSPF will start advertising available bandwidth for each TE class configured under the diffserv-te node. This command will only have effect if the operator has already enabled traffic engineering at the IS-IS and/or OSPF routing protocol levels:</p> <pre>config>router>isis>traffic-engineering</pre> <p>and/or:</p> <pre>config>router>ospf>traffic-engineering</pre> <p>IGP will advertize for each RSVP interface in the system the available bandwidth in each TE class in the unreserved bandwidth TE parameter for that class. In addition, IGP will continue to advertize the existing Maximum Reservable Link Bandwidth TE parameter to mean the maximum bandwidth that can be booked on a given interface by all classes. The value advertized is adjusted with the link subscription percentage factor configured in the config>router>rsvp>interface context.</p> <p>The user configures the following parameters for the operation of Diff-Serv:</p> <ul style="list-style-type: none">• Definition of TE classes, TE Class = {Class Type (CT), LSP priority}.• Mapping of the system forwarding classes to the Diff-Serv Class Type (CT).

- Configuration of the percentage of RSVP interface bandwidth each CT shares, i.e., the Bandwidth Constraint (BC).

When Diff-Serv TE is enabled, the system will automatically enable the Max Allocation Model (MAM) Admission Control Policy. MAM represents the bandwidth constraint model for the admission control of an LSP reservation to a link. This is the only Admission Control Policy supported in this release.

Each CT shares a percentage of the Maximum Reservable Link Bandwidth via the user configured Bandwidth Constraint (BC) for this CT. The Maximum Reservable Link Bandwidth is the link bandwidth multiplied by the RSVP interface subscription factor.

The sum of all BC values across all CTs will not exceed the Maximum Reservable Link Bandwidth. In other words, the following rule is enforced:

$$\text{SUM}(\text{BC}_c) \leq \text{Max-Reservable-Bandwidth}, 0 \leq c \leq 7$$

An LSP of class-type CT_c , setup priority p , holding priority h ($h \leq p$), and bandwidth B is admitted into a link if the following condition is satisfied:

$$B \leq \text{Unreserved Bandwidth for TE-Class}[i]$$

where TE-Class $[i]$ maps to $\langle \text{CT}_c, p \rangle$ in the definition of the TE classes on the node. The bandwidth reservation is effected at the holding priority, i.e., in TE-class $[j] = \langle \text{CT}_c, h \rangle$. Thus, the reserved bandwidth for CT_c and the unreserved bandwidth for the TE classes using CT_c are updated as follows:

$$\text{Reserved}(\text{CT}_c) = \text{Reserved}(\text{CT}_c) + B$$

$$\text{Unreserved TE-Class } [j] = \text{BC}_c - \text{SUM}(\text{Reserved}(\text{CT}_c, q)) \text{ for } 0 \leq q \leq h$$

$$\text{Unreserved TE-Class } [i] = \text{BC}_c - \text{SUM}(\text{Reserved}(\text{CT}_c, q)) \text{ for } 0 \leq q \leq p$$

The same is done to update the unreserved bandwidth for any other TE class making use of the same CT_c . These new values are advertised to the rest of the network at the next IGP-TE flooding.

The Russian Doll Model (RDM) LSP admission control policy allows bandwidth sharing across Class Types. It provides a hierarchical model by which the reserved bandwidth of a CT is the sum of the reserved bandwidths of the numerically equal and higher CTs.

The RDM model is defined using the following equations:

$$\text{SUM}(\text{Reserved}(\text{CT}_c)) \leq \text{BC}_b,$$

where the SUM is across all values of c in the range $b \leq c \leq (\text{MaxCT} - 1)$, and BC_b is the bandwidth constraint of CT_b .

$$\text{BC}_0 = \text{Max-Reservable-Bandwidth}, \text{ so that}$$

$$\text{SUM}(\text{Reserved}(\text{CT}_c)) \leq \text{Max-Reservable-Bandwidth},$$

where the SUM is across all values of c in the range $0 \leq c \leq (\text{MaxCT} - 1)$.

When Diff-Serv is disabled on the node, this model degenerates into a single default CT internally with eight pre-emption priorities and a non-configurable BC equal to the Maximum Reservable Link Bandwidth. This would behave exactly like CT_0 with eight pre-emption priorities and $\text{BC} = \text{Maximum Reservable Link Bandwidth}$ if Diff-Serv was enabled.

The enabling or disabling of Diff-Serv TE on the system requires the RSVP and MPLS protocol be shutdown.

The **no** form of this command reverts to the default value.

Default no diffserv-te

Parameters **mam** — Defines the default admission control policy for Diff-Serv LSPs.
rdm — Defines Russian doll model for the admission control policy of Diff-Serv LSPs.

class-type-bw

Syntax **class-type-bw ct0** *%-link-bandwidth* **ct1** *%-link-bandwidth* **ct2** *%-link-bandwidth* **ct3** *%-link-bandwidth* **ct4** *%-link-bandwidth* **ct5** *%-link-bandwidth* **ct6** *%-link-bandwidth* **ct7** *%-link-bandwidth*
no class-type-bw

Context config>router>rsvp>diffserv-te
config>router>rsvp>interface

Description This command configures the percentage of RSVP interface bandwidth each CT shares, for example, the Bandwidth Constraint (BC).

The absolute value of the CT share of the interface bandwidth is derived as the percentage of the bandwidth advertised by IGP in the Maximum Reservable Link Bandwidth TE parameter, for example, the link bandwidth multiplied by the RSVP interface **subscription percentage** parameter.

Note this configuration also exists at RSVP interface level and the interface specific configured value overrides the global configured value. The BC value can be changed at any time.

The RSVP interface **subscription percentage** parameter is configured in the **config>router>rsvp>interface** context.

The operator can specify the Bandwidth Constraint (BC) for a CT which is not used in any of the TE class definition but that does not get used by any LSP originating or transiting this node.

When Diff-Serv is disabled on the node, this model degenerates into a single default CT internally with eight pre-emption priorities and a non configurable BC equal to the Maximum Reservable Link Bandwidth. This would behave exactly like CT0 with eight pre-emption priorities and BC= Maximum Reservable Link Bandwidth if Diff-Serv was enabled.

The **no** form of this command reverts to the default value.

Parameters **ct0 (ct1/ct2/ — ct7)** *% link-bandwidth* — The Diff-Serv Class Type number. One or more system forwarding classes can be mapped to a CT.

Values 0 — 100 %

Default 0

fc

- Syntax** **fc** *fc-name* **class-type** *ct-number*
no fc *fc-name*
- Context** config>router>rsvp>diffserv-te
- Description** This command maps one or more system forwarding classes to a Diff-Serv Class Type (CT). The default mapping is shown in the following table.

FC ID	FC Name	FC Designation	Class Type (CT)
7	Network Control	NC	7
6	High-1	H1	6
5	Expedited	EF	5
4	High-2	H2	4
3	Low-1	L1	3
2	Assured	AF	2
1	Low-2	L2	1
0	Best Effort	BE	0

The **no** form of this command reverts to the default mapping for the forwarding class name.

- Parameters** **class-type** *ct-number* — The Diff-Serv Class Type number. One or more system forwarding classes can be mapped to a CT.

Values 0 — 7

te-class

- Syntax** **te-class** *te-class-number* **class-type** *ct-number* **priority** *priority*
no te-class *te-class-number*
- Context** config>router>rsvp>diffserv-te
- Description** This command configures a traffic engineering class. A TE class is defined as:
TE Class = {Class Type (CT), LSP priority}
- Eight TE classes are supported. There is no default TE class once Diff-Serv is enabled. The user has to explicitly define each TE class.
- When when Diff-Serv is disabled there will be an internal use of the default CT (CT0) and eight pre-emption priorities as shown in the following table.

Class Type (CT internal)	LSP Priority
0	7
0	6
0	5
0	4
0	3
0	2
0	1
0	0

The **no** form of this command deletes the TE class.

Parameters **te-class** *te-class-number* — The traffic engineering class number.

Values 0 — 7

class-type *ct-number* — The Diff-Serv Class Type number. One or more system forwarding classes can be mapped to a CT.

Values 0 — 7

priority *priority* — The LSP priority.

Values 0 — 7

gr-helper

Syntax **gr-helper** [enable | disable]

Context config>router>rsvp>if

Description This command enables the RSVP Graceful Restart Helper feature.

The RSVP-TE Graceful Restart helper mode allows the SR OS based system (the helper node) to provide another router that has requested it (the restarting node) a grace period, during which the system will continue to use RSVP sessions to neighbors requesting the grace period. This is typically used when another router is rebooting its control plane but its forwarding plane is expected to continue to forward traffic based on the previously available Path and Resv states.

The user can enable Graceful Restart helper on each RSVP interface separately. When the GR helper feature is enabled on an RSVP interface, the node starts inserting a new Restart_Cap Object in the Hello packets to its neighbor. The restarting node does the same and indicates to the helper node the desired Restart Time and Recovery Time.

The GR Restart helper consists of a couple of phases. Once it loses Hello communication with its neighbor, the helper node enters the Restart phase. During this phase, it preserves the state of all RSVP sessions to its neighbor and waits for a new Hello message.

Once the Hello message is received indicating the restarting node preserved state, the helper node enters the recovery phase in which it starts refreshing all the sessions that were preserved. The restarting node will activate all the stale sessions that are refreshed by the helper node. Any Path state which did not get a Resv message from the restarting node once the Recovery Phase time is over is considered to have expired and is deleted by the helper node causing the proper Path Tear generation downstream.

The duration of the restart phase (recovery phase) is equal to the minimum of the neighbor's advertised Restart Time (Recovery Time) in its last Hello message and the locally configured value of the max-restart (max-recovery) parameter.

When GR helper is enabled on an RSVP interface, its procedures apply to the state of both P2P and P2MP RSVP LSP to a neighbor over this interface.

Default disable

graceful-shutdown

Syntax	<code>[no] graceful-shutdown</code>
Context	config>router>rsvp config>router>rsvp>interface
Description	<p>This command initiates a graceful shutdown of the specified RSVP interface or all RSVP interfaces on the node if applied at the RSVP level. These are referred to as maintenance interface and maintenance node, respectively.</p> <p>To initiate a graceful shutdown the maintenance node generates a PathErr message with a specific error sub-code of Local Maintenance on TE Link required for each LSP that is exiting the maintenance interface.</p> <p>The node performs a single make-before-break attempt for all adaptive CSPF LSPs it originates and LSP paths using the maintenance interfaces. If an alternative path for an affected LSP is not found, then the LSP is maintained on its current path. The maintenance node also tears down and re-signals any detour LSP path using listed maintenance interfaces as soon as they are not active.</p> <p>The maintenance node floods an IGP TE LSA/LSP containing Link TLV for the links under graceful shutdown with Traffic Engineering metric set to 0xffffffff and Unreserved Bandwidth parameter set to zero (0).</p> <p>A head-end LER node, upon receipt of the PathErr message performs a single make-before-break attempt on the affected adaptive CSPF LSP. If an alternative path is not found, then the LSP is maintained on its current path.</p> <p>A node does not take any action on the paths of the following originating LSPs after receiving the PathErr message:</p> <ol style="list-style-type: none"> a. An adaptive CSPF LSP for which the PathErr indicates a node address in the address list and the node corresponds to the destination of the LSP. In this case, there are no alternative paths which can be found.

- b. An adaptive CSPF LSP whose path has explicit hops defined using the listed maintenance interface(s)/node(s).
- c. A CSPF LSP with the adaptive option disabled and which current path is over the listed maintenance interfaces in the PathErr message. These are not subject to make-before-break.
- d. A non CSPF LSP which current path is over the listed maintenance interfaces in the PathErr message.

The head-end LER node upon receipt of the updates IPG TE LSA/LSP for the maintenance interfaces updates the TE database. This information will be used at the next scheduled CSPF computation for any LSP which path may traverse any of the maintenance interfaces.

The **no** form of the command disables the graceful shutdown operation at the RSVP interface level or at the RSVP level. The configured TE parameters of the maintenance links are restored and the maintenance node floods the links.

Default none

gr-helper-time

Syntax **gr-helper-time max-recovery** *recovery-interval* [1..1800] seconds **max-restart** *restart-interval* [1..300] seconds
no gr-helper-time

Context config>router>rsvp

Description This command configures the local values for the max-recovery and the max-restart intervals used in the RSVP Graceful Restart Helper feature.

The values are configured globally in RSVP but separate instances of the timers are applied to each RSVP interface that has the RSVP Graceful Restart Helper enabled.

The **no** version of this command re-instates the default value for the delay timer.

Parameters *recovery-interval* — Specifies the max recovery interval value in seconds.

Values 1—1800

Default 300

restart-interval — Specifies the max restart interval value in seconds.

Values 1—300

Default 120

implicit-null-label

Syntax [**no**] **implicit-null-label**
implicit-null-label

Context config>router>rsvp

Description This command enables the use of the implicit null label.

Signalling the IMPLICIT NULL label value for all RSVP LSPs can be enabled for which this node is the egress LER. RSVP must be shutdown before being able to change this configuration option.

The egress LER does not signal the implicit null label value on P2MP RSVP LSPs. However, the Penultimate Hop Popping (PHP) node can honor a resv message with the label value set to the implicit null.

The **no** form of this command disables the signaling of the implicit null label.

Default no implicit-null-label

keep-multiplier

Syntax **[no] keep-multiplier** *number*
no keep-multiplier

Context config>router>rsvp

Description The **keep-multiplier** *number* is an integer used by RSVP to declare that a reservation is down or the neighbor is down.

The **no** form of this command reverts to the default value.

Default 3

Parameters *number* — The **keep-multiplier** value.

Values 1 — 255

refresh-reduction-over-bypass

Syntax **refresh-reduction-over-bypass** [**enable** | **disable**]

Context config>router>rsvp

Description This command enables the refresh reduction capabilities over all bypass tunnels originating on this PLR node or terminating on this Merge Point (MP) node.

By default, this is disabled. Since a bypass tunnel may merge with the primary LSP path in a node downstream of the next-hop, there is no direct interface between the PLR and the MP node and it is possible the latter will not accept summary refresh messages received over the bypass.

When disabled, the node as a PLR or MP will not set the “Refresh-Reduction-Capable” bit on RSVP messages pertaining to LSP paths tunneled over the bypass. It will also not send Message-ID in RSVP messages. This effectively disables summary refresh.

Default disable

rapid-retransmit-time

Syntax	rapid-retransmit-time <i>hundred-milliseconds</i> no rapid-retransmit-time
Context	config>router>rsvp
Description	<p>This command defines the value of the Rapid Retransmission Interval. It is used in the re-transmission mechanism to handle unacknowledged message_id objects and is based on an exponential back-off timer.</p> <p>Re-transmission interval of a RSVP message with the same message_id = 2 * rapid-retransmit-time interval of time.</p> <p>The node stops re-transmission of unacknowledged RSVP messages:</p> <ul style="list-style-type: none">• If the updated back-off interval exceeds the value of the regular refresh interval.• If the number of re-transmissions reaches the value of the rapid-retry-limit parameter, whichever comes first. <p>The Rapid Retransmission Interval must be smaller than the regular refresh interval configured in config>router>rsvp>refresh-time.</p> <p>The no form of this command reverts to the default value.</p>
Default	5
Parameters	<i>hundred-milliseconds</i> — Specifies the rapid retransmission interval. Values 1 – 100, in units of 100 msec.

rapid-retry-limit

Syntax	rapid-retry-limit <i>number</i> no rapid-retry-limit
Context	config>router>rsvp
Description	<p>This command is used to define the value of the Rapid Retry Limit. This is used in the retransmission mechanism based on an exponential backoff timer in order to handle unacknowledged message_id objects. The RSVP message with the same message_id is retransmitted every 2 * rapid-retransmit-time interval of time. The node will stop retransmission of unacknowledged RSVP messages whenever the updated backoff interval exceeds the value of the regular refresh interval or the number of retransmissions reaches the value of the rapid-retry-limit parameter, whichever comes first.</p> <p>The no form of this command reverts to the default value.</p>
Default	3
Parameters	<i>number</i> — Specifies the value of the Rapid Retry Limit. Values 1 – 6, integer values

refresh-time

Syntax	refresh-time <i>seconds</i> no refresh-time
Context	config>router>rsvp
Description	The refresh-time controls the interval, in seconds, between the successive Path and Resv refresh messages. RSVP declares the session down after it misses keep-multiplier <i>number</i> consecutive refresh messages. The no form of this command reverts to the default value.
Default	30 seconds
Parameters	<i>seconds</i> — The refresh time in seconds. Values 1 — 65535

te-threshold-update

Syntax	[no] te-threshold-update
Context	config>router>rsvp
Description	This command is used to control threshold-based IGP TE updates. The te-threshold-update command must enable IGP TE update based only on bandwidth reservation thresholds per interface and must block IGP TE update on bandwidth changes for each reservation. Threshold levels can be defined using the te-up-threshold and te-down-threshold commands at the global RSVP or per-interface level. The no form of this command should reset te-threshold-update to the default value and disable threshold based update.
Default	no te-threshold-update

on-cac-failure

Syntax	[no] on-cac-failure
Context	config>router>rsvp>te-threshold-update
Description	This command is used to enable a CAC failure-triggered IGP update. The no form of this command should reset on-cac-failure to the default value and disable the CAC failure-triggered IGP update.
Default	no on-cac-failure

update-timer

Syntax	update-timer <i>seconds</i> no update-timer
Context	config>router>rsvp>te-threshold-update
Description	This command is to control timer-based IGP TE updates. Timer-based IGP updates can be enabled by specifying a non-zero time value. Default value of update-timer is 0. The no form of this command should reset update-timer to the default value and disable timer-based IGP update.
Default	no update-timer (time - 0 seconds)
Parameters	<i>seconds</i> — The time in seconds. Values 0-300

te-up-threshold

Syntax	te-up-threshold <i>threshold-level</i> [<i>threshold-level</i> ...(up to 16 max)] no te-up-threshold
Context	config>router>rsvp config>router>rsvp>interface
Description	This command configures the specific threshold levels per node and per interface. Threshold levels are for reserved bandwidth per interface. The te-threshold-update command is used to enable or disable threshold-based IGP TE updates. Any reserved bandwidth change per interface is compared with all the threshold levels and trigger an IGP TE update if a defined threshold level is crossed in either direction (LSP setup or teardown). Threshold-based updates must be supported with both ISIS and OSPF. A minimum of one and a maximum of 16 threshold levels must be supported. Threshold levels configured per node is inherited by all configured RSVP interfaces. Threshold levels defined under the RSVP interface is used to trigger IGP updates if non-default threshold levels are configured. The no form of this command resets te-up-threshold to its default value.
Default	0 15 30 45 60 75 80 85 90 95 96 97 98 99 100
Parameters	<i>threshold-level</i> — Integer value Values 0 — 100

te-down-threshold

Syntax	te-down-threshold <i>threshold-level</i> [<i>threshold-level...</i> (up to 16 max)] no te-down-threshold
Context	config>router>rsvp config>router>rsvp>interface
Description	<p>This command configures the specific threshold levels per node and per interface. Threshold levels are for reserved bandwidth per interface. The te-threshold-update command is used to enable or disable threshold-based IGP TE updates. Any reserved bandwidth change per interface is compared with all the threshold levels and trigger an IGP TE update if a defined threshold level is crossed in either direction (LSP setup or teardown). Threshold-based updates must be supported with both ISIS and OSPF. A minimum of one and a maximum of 16 threshold levels is supported.</p> <p>Threshold levels configured per node is inherited by all configured RSVP interfaces. Threshold levels defined under the RSVP interface is used to trigger IGP updates if non-default threshold levels are configured.</p> <p>The no form of this command resets te-down-threshold to its default value.</p>
Default	100 99 98 97 96 95 90 85 80 75 60 45 30 15 0
Parameters	<i>threshold-level</i> — Integer value
	Values 0 — 100

Interface Commands

interface

Syntax	[no] interface <i>ip-int-name</i>
Context	config>router>rsvp
Description	<p>This command enables RSVP protocol support on an IP interface. No RSVP commands are executed on an IP interface where RSVP is not enabled.</p> <p>The no form of this command deletes all RSVP commands such as hello-interval and subscription, which are defined for the interface. The RSVP interface must be shutdown it can be deleted. If the interface is not shut down, the no interface <i>ip-int-name</i> command does nothing except issue a warning message on the console indicating that the interface is administratively up.</p>
Default	shutdown
Parameters	<p><i>ip-int-name</i> — The name of the network IP interface. An interface name cannot be in the form of an IP address. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.</p> <p>Values 1 — 32 alphanumeric characters.</p>

authentication-key

Syntax	authentication-key [<i>authentication-key</i> <i>hash-key</i>] [hash hash2] no authentication-key
Context	config>router>rsvp>interface
Description	<p>This command specifies the authentication key to be used between RSVP neighbors to authenticate RSVP messages. Authentication uses the MD-5 message-based digest.</p> <p>When enabled on an RSVP interface, authentication of RSVP messages operates in both directions of the interface.</p> <p>A node maintains a security association using one authentication key for each interface to a neighbor. The following items are stored in the context of this security association:</p> <ul style="list-style-type: none">• The HMAC-MD5 authentication algorithm.• Key used with the authentication algorithm.• Lifetime of the key. The user-entered key is valid until the user deletes it from the interface.• Source Address of the sending system.• Latest sending sequence number used with this key identifier. <p>A router RSVP sender transmits an authenticating digest of the RSVP message, computed using the shared authentication key and a keyed-hash algorithm. The message digest is included in an integrity object which also contains a flags field, a key identifier field, and a sequence number field. The</p>

RSVP sender complies to the procedures for RSVP message generation in RFC 2747, *RSVP Cryptographic Authentication*.

A RSVP receiver uses the key together with the authentication algorithm to process received RSVP messages.

When a PLR node switches the path of the LSP to a bypass LSP, it does not send the Integrity object in the RSVP messages sent over the bypass tunnel. If the PLR receives an RSVP message with an Integrity object, it will perform the digest verification for the key of the interface over which the packet was received. If this fails, the packet is dropped. If the received RSVP message is a RESV message and does not have an Integrity object, then the PLR node will accept it only if it originated from the MP node.

An MP node will accept RSVP messages received over the bypass tunnel with and without the Integrity object. If an Integrity object is present, the proper digest verification for the key of the interface over which the packet was received is performed. If this fails, the packet is dropped.

The MD5 implementation does not support the authentication challenge procedures in RFC 2747.

The **no** form of this command disables authentication.

Default	no authentication-key - The authentication key value is the null string.
Parameters	<p><i>authentication-key</i> — The authentication key. The key can be any combination of ASCII characters up to 16 characters in length (unencrypted). If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.</p> <p><i>hash-key</i> — The hash key. The key can be any combination of up to 33 alphanumeric characters. If spaces are used in the string, enclose the entire string in quotation marks (“ ”)</p> <p>This is useful when a user must configure the parameter, but for security purposes, the actual unencrypted key value is not provided.</p> <p>hash — Specifies the key is entered in an encrypted form. If the hash parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the hash parameter specified.</p> <p>hash2 — Specifies the key is entered in a more complex encrypted form. If the hash2 parameter is not used, the less encrypted hash form is assumed.</p>

bfd-enable

Syntax	[no] bfd-enable
Context	config>router>rsvp>interface
Description	<p>This command enables the use of bi-directional forwarding (BFD) to control the state of the associated RSVP interface. This causes RSVP to register the interface with the BFD session on that interface.</p> <p>The user configures the BFD session parameters, such as, transmit-interval, receive-interval, and multiplier, under the IP interface in the config>router> interface>bfd context.</p> <p>Note that it is possible that the BFD session on the interface was started because of a prior registration with another protocol, for example, OSPF or IS-IS.</p>

The registration of an RSVP interface with BFD is performed at the time of neighbor gets its first session. This means when this node sends or receives a new Path message over the interface. If however the session did not come up, due to not receiving a Resv for a new path message sent after the maximum number of re-tries, the LSP is shutdown and the node de-registers with BFD. In general, the registration of RSVP with BFD is removed as soon as the last RSVP session is cleared.

The registration of an RSVP interface with BFD is performed independent of whether RSVP hello is enabled on the interface or not. However, hello timeout will clear all sessions towards the neighbor and RSVP de-registers with BFD at clearing of the last session.

Note that an RSVP session is associated with a neighbor based on the interface address the path message is sent to. If multiple interfaces exist to the same node, then each interface is treated as a separate RSVP neighbor. The user will have to enable BFD on each interface and RSVP will register with the BFD session running with each of those neighbors independently

Similarly the disabling of BFD on the interface results in removing registration of the interface with BFD.

When a BFD session transitions to DOWN state, the following actions are triggered. For RSVP signaled LSPs, this triggers activation of FRR bypass/detour backup (PLR role), global revertive (head-end role), and switchover to secondary if any (head-end role) for affected LSPs with FRR enabled. It triggers switchover to secondary if any and scheduling of re-tries for signaling the primary path of the non-FRR affected LSPs (head-end role).

The **no** form of this command removes BFD from the associated RSVP protocol adjacency.

Default no bfd-enable

hello-interval

Syntax **hello-interval** *milli-seconds*
no hello-interval

Context config>router>rsvp>interface

Description This command configures the time interval between RSVP hello messages.

RSVP hello packets are used to detect loss of RSVP connectivity with the neighboring node. Hello packets detect the loss of neighbor far quicker than it would take for the RSVP session to time out based on the refresh interval. After the loss of the of number keep-multiplier consecutive hello packets, the neighbor is declared to be in a down state.

The **no** form of this command reverts to the default value of the hello-interval. To disable sending hello messages, set the value to zero.

Default 3000 milliseconds

Parameters *milli-seconds* — Specifies the RSVP hello interval in milliseconds, in multiples of 1000. A 0 (zero) value disables the sending of RSVP hello messages.

Values 0 — 60000 milliseconds (in multiples of 1000)

implicit-null-label

Syntax	implicit-null-label [enable disable] no implicit-null-label
Context	config>router>rsvp>interface
Description	<p>This command enables the use of the implicit null label over a specific RSVP interface.</p> <p>All LSPs for which this node is the egress LER and for which the path message is received from the previous hop node over this RSVP interface will signal the implicit null label. This means that if the egress LER is also the merge-point (MP) node, then the incoming interface for the path refresh message over the bypass dictates if the packet will use the implicit null label or not. The same for a 1-to-1 detour LSP.</p> <p>The user must shutdown the RSVP interface before being able to change the implicit null configuration option.</p> <p>The no form of this command returns the RSVP interface to use the RSVP level configuration value.</p>
Default	disable
Parameters	<p>enable — This parameter enables the implicit null label.</p> <p>disable — This parameter disables the implicit null label.</p>

refresh-reduction

Syntax	[no] refresh-reduction
Context	config>router>rsvp>interface
Description	<p>This command enables the use of the RSVP overhead refresh reduction capabilities on this RSVP interface.</p> <p>When this option is enabled, a node will enable support for three capabilities. It will accept bundles RSVP messages from its peer over this interface, it will attempt to perform reliable RSVP message delivery to its peer, and will use summary refresh messages to refresh path and resv states. The reliable message delivery must be explicitly enabled by the user after refresh reduction is enabled. The other two capabilities are enabled immediately.</p> <p>A bundle message is intended to reduce overall message handling load. A bundle message consists of a bundle header followed by one or more bundle sub-messages. A sub-message can be any regular RSVP message except another bundle message. A node will only process received bundled RSVP messages but will not generate them.</p> <p>When reliable message delivery is supported by both the node and its peer over the RSVP interface, an RSVP message is sent with a message_id object. A message_id object can be added to any RSVP message when sent individually or as a sub-message of a bundled message.</p> <p>if the sender sets the ack_desired flag in the message_id object, the receiver acknowledges the receipt of the RSVP message by piggy-backing a message_ack object to the next RSVP message it sends to its peer. Alternatively, an ACK message can also be used to send the message_ack object. In both cases, one or many message_ack objects could be included in the same message.</p>

The router supports the sending of separate ACK messages only but is capable of processing received message_ack objects piggy-backed to hop-by-hop RSVP messages, such as path and resv.

The router sets the ack_desired flag only in non refresh RSVP messages and in refresh messages which contain new state information.

A retransmission mechanism based on an exponential backoff timer is supported in order to handle unacknowledged message_id objects. The RSVP message with the same message_id is retransmitted every 2 * rapid-retransmit-time interval of time. The rapid-retransmit-time is referred to as the rapid retransmission interval as it must be smaller than the regular refresh interval configured in the **config>router>rsvp>refresh-time** context. There is also a maximum number of retransmissions of an unacknowledged RSVP message rapid-retry-limit. The node will stop retransmission of unacknowledged RSVP messages whenever the updated backoff interval exceeds the value of the regular refresh interval or the number of retransmissions reaches the value of the rapid-retry-limit parameter, whichever comes first. These two parameters are configurable globally on a system in the **config>router>rsvp** context.

Refresh summary consists of sending a summary refresh message containing a message_id list object. The fields of this object are populated each with the value of the message_identifier field in the message_id object of a previously sent individual path or resv message. The summary refresh message is sent every refresh regular interval as configured by the user using the refresh-time command in the **config>router>rsvp** context. The receiver checks each message_id object against the saved path and resv states. If a match is found, the state is updated as if a regular path or resv refresh message was received from the peer. If a specific message_identifier field does not match, then the node sends a message_id_nack object to the originator of the message.

The above capabilities are referred to collectively as “refresh overhead reduction extensions”. When the refresh-reduction is enabled on an RSVP interface, the node indicates this to its peer by setting a “refresh-reduction-capable” bit in the flags field of the common RSVP header. If both peers of an RSVP interface set this bit, all the above three capabilities can be used. Furthermore, the node monitors the settings of this bit in received RSVP messages from the peer on the interface. As soon as this bit is cleared, the router stops sending summary refresh messages. If a peer did not set the “refresh-reduction-capable” bit, a node does not attempt to send summary refresh messages.

However, if the peer did not set the “refresh-reduction-capable” bit, a node, with refresh reduction enabled and reliable message delivery enabled, will still attempt to perform reliable message delivery with this peer. If the peer does not support the message_id object, it returns an error message “unknown object class”. In this case, the node retransmits the RSVP message without the message_id object and reverts to using this method for future messages destined to this peer. The RSVP Overhead Refresh Reduction is supported with both RSVP P2P LSP path and the S2L path of an RSVP P2MP LSP instance over the same RSVP instance.

The **no** form of the command reverts to the default value.

Default no refresh-reduction

reliable-delivery

Syntax	[no] reliable-delivery
Context	config>router>rsvp>interface>refresh-reduction
Description	<p>This command enables reliable delivery of RSVP messages over the RSVP interface. When refresh-reduction is enabled on an interface and reliable-delivery is disabled, then the router will send a message_id and not set ACK desired in the RSVP messages over the interface. Thus 7750 does not expect an ACK and but will accept it if received. The node will also accept message ID and reply with an ACK when requested. In this case, if the neighbor set the “refresh-reduction-capable” bit in the flags field of the common RSVP header, the node will enter summary refresh for a specific message_id it sent regardless if it received an ACK or not to this message from the neighbor.</p> <p>Finally, when ‘reliable-delivery’ option is enabled on any interface, RSVP message pacing is disabled on all RSVP interfaces of the system, for example, the user cannot enable the msg-pacing option in the config>router>rsvp context, and error message is returned in CLI. Conversely, when the msg-pacing option is enabled, the user cannot enable the reliable delivery option on any interface on this system. An error message will also generated in CLI after such an attempt.</p> <p>The no form of the command reverts to the default value.</p>
Default	no reliable-delivery

subscription

Syntax	subscription <i>percentage</i> no subscription
Context	config>router>rsvp>interface
Description	<p>This command configures the percentage of the link bandwidth that RSVP can use for reservation and sets a limit for the amount of over-subscription or under-subscription allowed on the interface.</p> <p>When the subscription is set to zero, no new sessions are permitted on this interface. If the <i>percentage</i> is exceeded, the reservation is rejected and a log message is generated.</p> <p>The no form of this command reverts the <i>percentage</i> to the default value.</p>
Default	100
Parameters	<i>percentage</i> — The percentage of the interface's bandwidth that RSVP allows to be used for reservations.
	Values 0 — 1000

te-up-threshold

Syntax	te-up-threshold <i>threshold-level</i> [<i>threshold-level...</i> (up to 16 max)] no te-up-threshold
Context	config>router>rsvp config>router>rsvp>interface
Description	<p>This command configures the specific threshold levels per node and per interface. Threshold levels are for reserved bandwidth per interface. The te-threshold-update command is used to enable or disable threshold-based IGP TE updates. Any reserved bandwidth change per interface is compared with all the threshold levels and trigger an IGP TE update if a defined threshold level is crossed in either direction (LSP setup or teardown). Threshold-based updates must be supported with both ISIS and OSPF. A minimum of one and a maximum of 16 threshold levels must be supported.</p> <p>Threshold levels configured per node is inherited by all configured RSVP interfaces. Threshold levels defined under the RSVP interface is used to trigger IGP updates if non-default threshold levels are configured.</p> <p>The no form of this command resets the default value.</p>
Default	0 15 30 45 60 75 80 85 90 95 96 97 98 99 100
Parameters	<i>threshold-level</i> — Integer value Values 0 — 100

te-down-threshold

Syntax	te-down-threshold <i>threshold-level</i> [<i>threshold-level...</i> (up to 16 max)] no te-down-threshold
Context	config>router>rsvp config>router>rsvp>interface
Description	<p>This command configures the specific threshold levels per node and per interface. Threshold levels are for reserved bandwidth per interface. The te-threshold-update command is used to enable or disable threshold-based IGP TE updates. Any reserved bandwidth change per interface is compared with all the threshold levels and trigger an IGP TE update if a defined threshold level is crossed in either direction (LSP setup or teardown). Threshold-based updates is supported with both ISIS and OSPF. A minimum of one and a maximum of 16 threshold levels is supported.</p> <p>Threshold levels configured per node is inherited by all configured RSVP interfaces. Threshold levels defined under the RSVP interface must be used to trigger IGP updates if non-default threshold levels are configured.</p> <p>The no form of this command resets the default value.</p>
Default	100 99 98 97 96 95 90 85 80 75 60 45 30 15 0
Parameters	<i>threshold-level</i> — Integer value Values 0 — 100

Message Pacing Commands

msg-pacing

Syntax	[no] msg-pacing
Context	config>router>rsvp
Description	This command enables RSVP message pacing in which the specified number of RSVP messages, specified in the max-burst command, are sent in a configured interval, specified in the period command. A count is kept of the messages that were dropped because the output queue for the interface used for message pacing was full.
Default	no msg-pacing

max-burst

Syntax	max-burst <i>number</i> no max-burst
Context	config>router>rsvp>msg-pacing
Description	This command specifies the maximum number of RSVP messages that are sent in the specified period under normal operating conditions.
Default	650
Parameters	<i>number</i> —
	Values 100 — 1000 in increments of 10

period

Syntax	period <i>milli-seconds</i> no period
Context	config>router>rsvp>msg-pacing
Description	This command specifies the time interval, in milliseconds, when the router can send the specified number of RSVP messages which is specified in the max-burst command.
Default	100
Parameters	<i>milli-seconds</i> —
	Values 100 — 1000 milliseconds in increments of 10 milliseconds

Show Commands

admin-group

- Syntax** `admin-group group-name`
- Context** `show>router>mpls`
- Description** This command displays MPLS administrative group information.
- Parameters** *group-name* — Specify a group name up to 32 characters.
- Output** **MPLS Administrative Group Output Fields** — The following table describes MPLS administrative group output fields.

Label	Description
Group Name	The name of the group. The name identifies the administrative group within a virtual router instance.
Group Value	The unique group value associated with the administrative group. If the value displays -1, then the group value for this entry has not been set.
No. of Groups	The total number of configured admin groups within the virtual router instance.

Sample Output

```
A:ALA-1# show router mpls admin-group
=====
MPLS Administrative Groups
=====
Group Name                               Group Value
-----
green                                     15
red                                       25
yellow                                    20
-----
No. of Groups: 3
=====
A:ALA-1#
```

auto-lsp

- Syntax** **auto-lsp** [*lsp-name*] **auto-bandwidth**
auto-lsp [*lsp-name*] [**status** {up|down}] [**detail**] [**to** *ip-address*]
auto-lsp [*lsp-name*] [**status** {up|down}] {**mesh-p2p** | **one-hop-p2p**} [**detail**] [**to** *ip-address*]
- Context** show>router>mpls
- Parameters** *lsp-name* — Specifies the LSP name.
- Values** 80 characters max
- up|down** — Specifies the state.
- mesh-p2p|one-hop-p2p** — Specifies the auto LSP type.

Sample Output

```
*A:Dut-C# show router mpls auto-lsp
```

```
=====
MPLS Auto-LSP Template
=====
LSP Name                               Type                Fastfail   Admin   Oper
                               Config              State     State
-----
meshP2pLsp3-10.20.1.6-61441           MeshP2P             Yes        Up      Up
meshP2pLsp2-10.20.1.1-61442           MeshP2P             Yes        Up      Up
meshP2pLsp2-10.20.1.2-61443           MeshP2P             Yes        Up      Up
meshP2pLsp2-10.20.1.4-61444           MeshP2P             Yes        Up      Up
meshP2pLsp2-10.20.1.5-61445           MeshP2P             Yes        Up      Up
meshP2pLsp2-10.20.1.6-61446           MeshP2P             Yes        Up      Up
meshP2pLsp10-10.20.1.1-61447          MeshP2P             Yes        Up      Up
meshP2pLsp10-10.20.1.2-61448          MeshP2P             Yes        Up      Up
```

bypass-tunnel

- Syntax** **bypass-tunnel** [**to** *ip-address*] [**protected-lsp** [*lsp-name*]] [**dynamic** | **manual** | **p2mp**] [**detail**]
- Context** show>router>mpls
- Description** If fast reroute is enabled on an LSP and the facility method is selected, instead of creating a separate LSP for every LSP that is to be backed up, a single LSP is created which serves as a backup for a set of LSPs. Such an LSP tunnel is called a bypass tunnel.
- Parameters** *ip-address* — Specify the IP address of the egress router.
- lsp-name* — Specify the name of the LSP protected by the bypass tunnel.
- dynamic** — Displays dynamically assigned labels for bypass protection.
- manual** — Displays manually assigned labels for bypass protection.
- detail** — Displays detailed information.

p2mp — Displays P2MP bypass tunnel information.

Output MPLS Bypass Tunnel Output Fields — The following table describes MPLS bypass tunnel output fields.

Label	Description
To	The system IP address of the egress router.
State	The LSP's administrative state.
Out I/F	Specifies the name of the network IP interface.
Out Label	Specifies the incoming MPLS label on which to match.
Reserved BW (Kbps)	Specifies the amount of bandwidth in megabits per second (Mbps) reserved for the LSP.

Sample Output

```
*A:Dut-B# show router mpls bypass-tunnel detail

=====
MPLS Bypass Tunnels (Detail)
=====
-----
bypass-node10.20.1.4
-----
To           : 10.20.1.7           State          : Up
Out I/F      : 1/1/4             Out Label     : 131071
Up Time     : 0d 01:17:22       Active Time   : n/a
Reserved BW  : 0 Kbps          Protected LSP Count : 1
Type        : Dynamic
Setup Priority : 7              Hold Priority  : 0
Class Type   : 0
Exclude Node : 10.20.1.4       Inter-Area    : True
Computed Hops :
  10.10.8.2(S)                Egress Admin Groups : None
-> 10.10.8.6(SA)              Egress Admin Groups : None
-> 10.20.1.7(L)                Egress Admin Groups : None
Actual Hops  :
  10.10.8.2(10.20.1.2)        Record Label   : N/A
-> 10.10.8.6(10.20.1.6)      Record Label   : 131071
-> 10.20.1.7(10.20.1.7)     Record Label   : 131068
-> 10.10.22.7                 Record Label   : 131068

=====

*A:Dut-A>config>router>mpls>lsp$ /show router mpls bypass-tunnel detail

=====
MPLS Bypass Tunnels (Detail)
=====
-----
bypass-node10.20.1.2
-----
To           : 10.20.1.4           State          : Up
Out I/F      : 1/1/2             Out Label     : 131070
```

```

Up Time          : 0d 00:00:18      Active Time       : n/a
Reserved BW     : 0 Kbps           Protected LSP Count : 1
Type            : Dynamic
Setup Priority   : 7                Hold Priority      : 0
Class Type      : 0
Exclude Node    : None             Inter-Area        : False
Computed Hops   :
    10.20.1.1, If Index : 3(S)      Egress Admin Groups : None
-> 10.20.1.3, If Index : 2(S)      Egress Admin Groups : None
-> 10.20.1.4, If Index : 5(S)      Egress Admin Groups : None
Actual Hops     :
    10.20.1.1, If Index : 3         Record Label      : N/A
-> 10.20.1.3, If Index : 2         Record Label      : 131070
-> 10.20.1.4, If Index : 5         Record Label      : 131070

```

=====

```
B:Dut-B>config>router>mpls>lsp# show router mpls bypass-tunnel detail
```

```
=====
MPLS Bypass Tunnels (Detail)
=====
```

```
bypass-node10.20.1.4
```

```
-----
To          : 10.10.10.6      State          : Up
Out I/F     : lag-1          Out Label      : 131071
Up Time     : 0d 00:00:06    Active Time    : n/a
Reserved BW : 0 Kbps         Protected LSP Count : 1
Type        : Dynamic
Setup Priority : 7           Hold Priority   : 0
Class Type   : 0
Exclude Node : None
Actual Hops  :
    10.10.12.2(S)          Egress Admin Groups:
                            lime
                            olive
                            blue
                            black
                            acqua
-> 10.10.12.3(S)          Egress Admin Groups:
                            olive
                            Unknown Group 9
                            Unknown Group 11
                            black
                            Unknown Group 16
                            Unknown Group 18
-> 10.10.5.5(S)          Egress Admin Groups:
                            purple
                            Unknown Group 7
                            Unknown Group 11
                            orange
                            acqua
                            Unknown Group 16
                            Unknown Group 19
                            Unknown Group 21
                            Unknown Group 22
                            Unknown Group 26
                            khaki
-> 10.10.10.6(S)          Egress Admin Groups: None

```

```

=====
*A:SRU4>show>router>mpls# bypass-tunnel
=====
MPLS Bypass Tunnels
=====
Legend :  m - Manual      d - Dynamic      p - P2mp
=====
To          State  Out I/F      Out Label    Reserved   Protected   Type
          BW (Kbps)  LSP Count
-----
No Matching Entries Found
=====
*A:SRU4>show>router>mpls#

*A:Dut-B# show router mpls bypass-tunnel detail
=====
MPLS Bypass Tunnels (Detail)
=====
bypass-link10.10.104.4
-----
To          : 10.10.101.4      State          : Up
Out I/F     : 1/1/2:1             Out Label     : 129994
Up Time    : 0d 00:02:33   Active Time    : n/a
Reserved BW : 0 Kbps        Protected LSP Count : 1
Type       : Dynamic
SetupPriority : 7          Hold Priority   : 0
Class Type  : 0
Actual Hops :
          10.10.101.2    -> 10.10.101.4
=====
*A:Dut-B#

*A:Dut-B# show router mpls bypass-tunnel detail
=====
MPLS Bypass Tunnels (Detail)
=====
bypass-link10.10.104.4
-----
To          : 10.10.101.4      State          : Up
Out I/F     : 1/1/2:1             Out Label     : 129994
Up Time    : 0d 00:02:33   Active Time    : n/a
Reserved BW : 0 Kbps        Protected LSP Count : 1
Type       : Dynamic
SetupPriority : 7          Hold Priority   : 0
Class Type  : 0
Actual Hops :
          10.10.101.2    -> 10.10.101.4
=====
*A:Dut-B#

```

interface

Syntax **interface** [*ip-int-name* | *ip-address*] [**label-map** *label*]
interface [*ip-int-name* | *ip-address*] **statistics**

Context show>router>mpls

Description This command displays MPLS interface information.

Parameters *ip-int-name* — The name of the network IP interface. An interface name cannot be in the form of an IP address. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

ip-address — The system or network interface IP address.

label-map *label* — The MPLS label on which to match.

Values 32 — 1048575

statistics — Displays MPLS interface name and the number of packets and octets sent and received on an MPLS interface.

Output **MPLS Interface Output Fields** — The following table describes MPLS interface output fields.

Label	Description
Interface	The interface name.
Port-id	The port ID displayed in the <i>slot/mda/port</i> format.
Adm	Specifies the administrative state of the interface.
Oper	Specifies the operational state of the interface.
Te-metric	Specifies the traffic engineering metric used on the interface.
Srlg Groups	Specifies the shared risk loss group (SRLG) name(s).
Interfaces	The total number of interfaces.
Transmitted	Displays the number of packets and octets transmitted from the interface.
Received	Displays the number of packets and octets received.
In Label	Specifies the ingress label.
In I/F	Specifies the ingress interface.
Out Label	Specifies the egress label.
Out I/F	Specifies the egress interface.
Next Hop	Specifies the next hop IP address for the static LSP.
Type	Specifies whether the label value is statically or dynamically assigned.

Sample Output

```

*A:SRU4>config>router>mpls# show router mpls interface
=====
MPLS Interfaces
=====
Interface                               Port-id           Adm   Opr   TE-metric
-----
system                                  system            Up    Up    None
  Admin Groups                           None
  Srlg Groups                             None
aps-1                                    aps-1             Up    Up    None
  Admin Groups                           None
  Srlg Groups                             3410
aps-2                                    aps-2             Up    Up    None
  Admin Groups                           None
  Srlg Groups                             3420
aps-3                                    aps-3             Up    Up    None
  Admin Groups                           None
  Srlg Groups                             3430
sr4-1                                    1/1/4            Up    Up    None
  Admin Groups                           None
  Srlg Groups                             3440
ess-7-1                                  3/2/4            Up    Up    None
  Admin Groups                           None
  Srlg Groups                             45100
ess-7-2                                  3/2/5            Up    Up    None
  Admin Groups                           None
  Srlg Groups                             45110
...
g7600                                    3/1/2            Up    Up    None
  Admin Groups                           None
  Srlg Groups                             41.80
ml60                                      3/2/1            Up    Up    None
  Admin Groups                           None
  Srlg Groups                             420.40
-----
Interfaces : 35
=====
*A:SRU4>config>router>mpls#

*A:SRU4>config>router>mpls# show router mpls interface "hubA"
=====
MPLS Interface : hubA
=====
Interface                               Port-id           Adm   Opr   TE-metric
-----
hubA                                    3/2/8            Up    Up    None
  Admin Groups                           None
  Srlg Groups                             44.200
-----
Interfaces : 1
=====
*A:SRU4>config>router>mpls#

*A:SRU4>config>router>mpls# show router mpls interface "hubA" label-map 203
=====
MPLS Interface : hubA (Label-Map 203)
=====

```

```

In Label  In I/F      Out Label Out I/F      Next Hop      Type      Adm  Opr
-----
203      3/2/8      403         1/1/9      11.22.10.3   Static    Up   Up
-----
Interfaces : 1
=====
*A:SRU4>config>router>mpls#

*A:SRU4>config>router>mpls# show router mpls interface statistics
=====
MPLS Interface (statistics)
=====
Interface      : aps-1
  Transmitted  : Pkts - 76554          Octets - 7930285
  Received     : Pkts - 17068       Octets - 3626842

Interface      : aps-2
  Transmitted  : Pkts - 0          Octets - 0
  Received     : Pkts - 1311       Octets - 219888

Interface      : aps-3
  Transmitted  : Pkts - 0          Octets - 0
  Received     : Pkts - 3          Octets - 234

Interface      : sr4-1
  Transmitted  : Pkts - 0          Octets - 0
  Received     : Pkts - 0          Octets - 0

Interface      : ess-7-1
  Transmitted  : Pkts - 113537     Octets - 15058332
  Received     : Pkts - 13193     Octets - 1091492

Interface      : ess-7-2
  Transmitted  : Pkts - 166133     Octets - 22762482
  Received     : Pkts - 16672     Octets - 1368464

Interface      : ess-7-3
  Transmitted  : Pkts - 122934     Octets - 11033246
  Received     : Pkts - 12256     Octets - 1026826
...

Interface      : m160
  Transmitted  : Pkts - 17188024    Octets - 2183076528
  Received     : Pkts - 677745     Octets - 59367236
=====
*A:SRU4>config>router>mpls#

```

label

Syntax `label start-label [end-label | in-use | owner]`

Context `show>router>mpls`

Description Displays MPLS labels exchanged.

- Parameters**
- start-label* — The label value assigned at the ingress router.
 - end-label* — The label value assigned for the egress router.
 - in-use* — The number of in-use labels displayed.

Output **MPLS Label Output Fields** — The following table describes MPLS label output fields.

Label	Description
Label	Displays the value of the label being displayed.
Label Type	Specifies whether the label value is statically or dynamically assigned.
Label Owner	The label owner.
In-use labels in entire range	The total number of labels being used by RSVP.

Sample Output

```
*A:mlstp-dutA# show router mpls label-range
=====
Label Ranges
=====
Label Type      Start Label      End Label      Aging      Total Available
-----
Static-lsp      32                16415         -           16364
Static-svc      16416             32799         -           16376
Dynamic         32800             131071        0           98268
=====

*A:SRU4>config>router>mpls# show router mpls label 202
=====
MPLS Label 202
=====
Label           Label Type      Label Owner
-----
202             static-lsp      STATIC
-----
In-use labels in entire range      : 5057
=====
*A:SRU4>config>router>mpls#
```

label-range

- Syntax** **label-range**
- Context** show>router>mpls
- Description** This command displays the MPLS label range.

Output MPLS Label Range Output — The following table describes the MPLS label range output fields.

Label	Description
Label Type	Displays the information about static-lsp , static-svc , and dynamic label types.
Start Label	The label value assigned at the ingress router.
End Label	The label value assigned for the egress router.
Aging	The number of labels released from a service which are transitioning back to the label pool. Labels are aged 15 seconds.
Total Available	The number of label values available.

Sample Output

```
*A:SRU4>config>router>mpls# show router mpls label-range
=====
Label Ranges
=====
Label Type      Start Label    End Label      Aging          Total Available
-----
Static-lsp      32             1023           -              736
Static-svc      2048           18431          -              16384
Dynamic         32768          131071         258            93232
=====
*A:SRU4>config>router>mpls#
```

lsp

Syntax **lsp** *lsp-name* [**status** {**up**|**down**}] [**from** *ip-address* | **to** *ip-address*] [**detail**]
lsp {**transit** | **terminate**} [**status** {**up** | **down**}] [**from** *ip-address* | **to** *ip-address* | *lsp-name* *name*] [**detail**]
lsp count
lsp *lsp-name* **activepath**
lsp *lsp-name* **path** [*path-name*] [**status** {**up** | **down**}] [**detail**]
lsp [*lsp-name*] **path** [*path-name*] **mbb**

Context show>router>mpls

Description This command displays LSP details.

Parameters **lsp** *lsp-name* — The name of the LSP used in the path.
status up — Displays an LSP that is operationally up.
status down — Displays an LSP that is operationally down.
from ip-address — Displays the IP address of the ingress router for the LSP.
to ip-address — Displays the IP address of the egress router for the LSP.
transit — Displays the number of static LSPs that transit through the router.
terminate — Displays the number of static LSPs that terminate at the router.

lsp count — Displays the total number of LSPs.

activepath — Displays the present path being used to forward traffic.

mbb — Displays make-before-break (MBB) information.

detail — Displays detailed information.

Output MPLS LSP Output — The following table describes MPLS LSP output fields.

Label	Description
LSP Name	The name of the LSP used in the path.
To	The system IP address of the egress router for the LSP.
Adm State	Down — The path is administratively disabled. Up — The path is administratively enabled.
Oper State	Down — The path is operationally down. Up — The path is operationally up.
Oper State	Down — The path is operationally down. Up — The path is operationally up.
LSPs	The total number of LSPs configured.
From	The IP address of the ingress router for the LSP.
LSP Up Time	The length of time the LSP has been operational.
Transitions	The number of transitions that have occurred for the LSP.
Retry Limit	The number of attempts that the software should make to re-establish the LSP after it has failed.
Signaling	Specifies the signaling style.
Hop Limit	The maximum number of hops that an LSP can traverse, including the ingress and egress routers.
Fast Reroute/ FastFail Config	enabled — Fast reroute is enabled. In the event of a failure, traffic is immediately rerouted on the pre-computed detour LSP, thus minimizing packet loss. disabled — There is no detour LSP from each node on the primary path.
ADSPEC	enabled — The LSP will include advertising data (ADSPEC) objects in RSVP messages. disabled — The LSP will not include advertising data (ADSPEC) objects in RSVP messages.
Primary	The preferred path for the LSP.

Label	Description (Continued)
Secondary	The alternate path that the LSP will use if the primary path is not available.
Bandwidth	The amount of bandwidth in megabits per second (Mbps) reserved for the LSP path.
LSP Up Time	The total time in increments that the LSP path has been operational.
LSP Tunnel ID	The value which identifies the label switched path that is signaled for this entry.
To	The IP address of the egress router for the LSP.
LSP Down Time	The total time in increments that the LSP path has not been operational.
Path Changes	The number of path changes this LSP has had. For every path change (path down, path up, path change), a corresponding syslog/trap (if enabled) is generated.
Retry Timer	The time, in seconds, for LSP re-establishment attempts after an LSP failure.
Resv Style	<p><i>se</i> – Specifies a shared reservation environment with a limited reservation scope. This reservation style creates a single reservation over a link that is shared by an explicit list of senders.</p> <p><i>ff</i> – Specifies a shared reservation environment with an explicit reservation scope. Specifies an explicit list of senders and a distinct reservation for each of them.</p>
Negotiated MTU	The size of the maximum transmission unit (MTU) that is negotiated during establishment of the LSP.
FR Hop Limit	The total number of hops a detour LSP can take before merging back onto the main LSP path.
LastResignalAttempt	Displays the system up time when the last attempt to resignal this LSP was made.
MBB Type	Displays an enumerated integer that specifies the type of make-before-break (MBB). If none displays then there is no MBB in progress or no last MBB.
MBB State	Displays the state of the most recent invocation of the make-before-break functionality.
End at	Displays the system up time when the last MBB ended.
Old Metric	Displays the cost of the traffic engineered path for the LSP path prior to MBB.
NextRetryIn	Displays the amount of time remaining, in seconds, before the next attempt is made to retry the in-progress MBB.

Label	Description (Continued)
RetryAttempt	Displays the number attempts for the MBB is in progress.
Failure Code	Displays the reason code for in-progress MBB failure. A value of none indicates that no failure has occurred.
Failure Node	Displays the IP address of the node in the LSP path at which the in-progress MBB failed. When no failure has occurred, this value is none .

Sample Output

```
*A:Dut-A>config>router>mpls>lsp$ /show router mpls lsp "1" path detail

=====
MPLS LSP 1 Path (Detail)
=====
Legend :
  @ - Detour Available           # - Detour In Use
  b - Bandwidth Protected       n - Node Protected
  s - Soft Preemption
  S - Strict                     L - Loose
  A - ABR

=====
-----
LSP 1 Path 1
-----
LSP Name       : 1                      Path LSP ID : 30208
From           : 10.20.1.1              To           : 10.20.1.6
Adm State      : Up                     Oper State   : Up
Path Name      : 1                      Path Type    : Primary
Path Admin     : Up                     Path Oper    : Up
OutInterface   : 1/1/1                  Out Label    : 131071
Path Up Time   : 0d 00:00:05            Path Dn Time : 0d 00:00:00
Retry Limit    : 0                      Retry Timer  : 30 sec
RetryAttempt   : 0                      NextRetryIn  : 0 sec

Adspec        : Disabled                Oper Adspec  : Disabled
CSPF          : Enabled                 Oper CSPF    : Enabled
Least Fill    : Disabled                Oper LeastF* : Disabled
FRR           : Enabled                 Oper FRR     : Enabled
FRR NodePro*  : Enabled                Oper FRR NP  : Enabled
FR Hop Limit  : 16                     Oper FRHopL* : 16
FR Prop Adm*  : Disabled                Oper FRProp* : Disabled
Prop Adm Grp  : Disabled                Oper PropAG  : Disabled
Inter-area    : False

Neg MTU       : 1496                    Oper MTU     : 1496
Bandwidth     : No Reservation           Oper Bw      : 0 Mbps
Hop Limit     : 255                     Oper HopLim* : 255
Record Route  : Record                  Oper RecRou* : Record
Record Label  : Record                  Oper RecLab* : Record
SetupPriori*  : 7                      Oper SetupP* : 7
Hold Priori*  : 0                      Oper HoldPr* : 0
Class Type    : 0                      Oper CT      : 0
Backup CT     : None
MainCT Retry  : n/a
Rem           :
```

```

MainCT Retry: 0
  Limit      :
Include Grps:
None
Exclude Grps:
None
Oper InclGr*:
None
Oper ExclGr*:
None

Adaptive    : Enabled
Preference  : n/a
Path Trans  : 1
Failure Code: noError
ExplicitHops:
  No Hops Specified
Actual Hops :
  10.20.1.1, If Index : 2 @ n
-> 10.20.1.2, If Index : 2 @ n
-> 10.20.1.4, If Index : 2
-> 10.20.1.6, If Index : 2
Record Label      : N/A
Record Label      : 131071
Record Label      : 131071
Record Label      : 131071
ComputedHops:
  10.20.1.1, If Index : 2(S)
-> 10.20.1.2, If Index : 2(S)
-> 10.20.1.4, If Index : 2(S)
-> 10.20.1.6, If Index : 2(S)
ResigEligib*: False
LastResignal: n/a
CSPF Metric : 3000

```

=====
* indicates that the corresponding row element may have been truncated.

*A:Dut-A# show router mpls lsp "AtoL1" path detail

```

=====
MPLS LSP AtoL1 Path (Detail)
=====
Legend :
  @ - Detour Available          # - Detour In Use
  b - Bandwidth Protected      n - Node Protected
  s - Soft Preemption
  S - Strict                    L - Loose
  A - ABR
=====
-----
LSP AtoL1 Path empty
-----
LSP Name      : AtoL1
From          : 10.20.1.1
Adm State     : Up
Path Name     : empty
Path Admin    : Up
OutInterface  : 1/1/1
Path Up Time  : 0d 01:19:46
Retry Limit   : 0
RetryAttempt  : 0
Path LSP ID   : 13316
To            : 10.20.1.12
Oper State    : Up
Path Type     : Primary
Path Oper     : Up
Out Label     : 131069
Path Dn Time  : 0d 00:00:00
Retry Timer   : 20 sec
NextRetryIn   : 0 sec

Adspec       : Disabled
CSPF         : Enabled
Least Fill   : Disabled
FRR          : Enabled
FRR NodePro* : Enabled
FR Hop Limit : 6
FR Prop Adm* : Disabled
Oper Adspec  : Disabled
Oper CSPF    : Enabled
Oper LeastF* : Disabled
Oper FRR     : Enabled
Oper FRR NP  : Enabled
Oper FRHopL* : 6
Oper FRProp* : Disabled

```



```

Prop Adm Grp: Enabled
Inter-area : True

Neg MTU      : 1496
Bandwidth   : 1 Mbps
Hop Limit    : 255
Record Route: Record
Record Label: Record
SetupPriori*: 7
Hold Priori*: 0
Class Type   : 0
Backup CT    : None
MainCT Retry: n/a
  Rem        :
MainCT Retry: 0
  Limit      :
Include Grps:
None
Exclude Grps:
None

Adaptive     : Enabled
Preference   : n/a
Path Trans   : 1
Failure Code: noError
ExplicitHops:
  No Hops Specified
Actual Hops :
  10.10.1.1(10.20.1.1) @ n
-> 10.10.1.2(10.20.1.2) @ n
-> 10.10.5.4(10.20.1.4) @ n
-> 10.20.1.7(10.20.1.7) @ n
-> 10.10.17.7 @ n
-> 10.20.1.9(10.20.1.9) @
-> 10.10.25.9 @
-> 10.20.1.12(10.20.1.12)
-> 10.10.33.12
ComputedHops:
  10.10.1.1(S)
-> 10.10.1.2(S)
-> 10.10.5.4(SA)
-> 10.20.1.12(L)
ResigEligib*: False
LastResignal: n/a

Oper PropAG : Enabled
Oper MTU     : 1496
Oper Bw      : 1 Mbps
Oper HopLim* : 255
Oper RecRou* : Record
Oper RecLab* : Record
Oper SetupP* : 7
Oper HoldPr* : 0
Oper CT      : 0

Oper InclGr*:
None
Oper ExclGr*:
None

Oper Metric : 1500
CSPF Queries: 3
Failure Node: n/a

Record Label : N/A
Record Label : 131069
Record Label : 131069
Record Label : 131069
Record Label : 131069
Record Label : 131069
Record Label : 131069
Record Label : 131068
Record Label : 131068

CSPF Metric : 1500
=====
* indicates that the corresponding row element may have been truncated.

```

*A:Dut-C# show router mpls lsp detail

```

=====
MPLS LSPs (Originating) (Detail)
=====
-----
Type : Originating
-----
LSP Name      : to_D_10.20.1.4_viaBD
LSP Type      : RegularLsp
From          : 10.20.1.3
Adm State     : Up
LSP Tunnel ID : 1
To           : 10.20.1.4
Oper State    : Up

```

```

LSP Up Time : 0d 00:05:38
Transitions : 1
Retry Limit : 0
Signaling : RSVP
Hop Limit : 255
Adaptive : Enabled
FastReroute : Disabled
CSPF : Enabled
Metric : 0
Include Grps:
None
Least Fill : Disabled

LSP Down Time : 0d 00:00:00
Path Changes : 1
Retry Timer : 30 sec
Resv. Style : SE
Negotiated MTU : 1500
ClassType : 0
Oper FR : Disabled
ADSPEC : Disabled
Use TE metric : Disabled
Exclude Grps :
None

```

```

Auto BW : Disabled
LdpOverRsvp : Enabled
IGP Shortcut: Enabled
IGP LFA : Disabled
BGPTransTun : Enabled
Oper Metric : 20
Prop Adm Grp: Disabled

VprnAutoBind : Enabled
BGP Shortcut : Enabled
IGP Rel Metric : -1

```

```

Primary(a) : to_D_10.20.1.4_viaBD
Bandwidth : 0 Mbps
Up Time : 0d 00:05:38

```

```

=====
*A:Dut-C#

```

```

*A:Dut-A# show router mpls lsp "AtoL1" detail

```

```

=====
MPLS LSPs (Originating) (Detail)
=====
-----

```

```

Type : Originating
-----

```

```

LSP Name : AtoL1
LSP Type : RegularLsp
From : 10.20.1.1
Adm State : Up
LSP Up Time : 0d 01:19:30
Transitions : 1
Retry Limit : 0
Signaling : RSVP
Hop Limit : 255
Adaptive : Enabled
FastReroute : Enabled
FR Method : Facility
FR Bandwidth: 0 Mbps
FR Object : Enabled
CSPF : Enabled
Metric : 0
Include Grps:
None
Least Fill : Disabled

LSP Tunnel ID : 1
To : 10.20.1.12
Oper State : Up
LSP Down Time : 0d 00:00:00
Path Changes : 1
Retry Timer : 20 sec
Resv. Style : SE
Negotiated MTU : 1496
ClassType : 0
Oper FR : Enabled
FR Hop Limit : 6
FR Node Protect: Enabled
FR Prop Adm Grp: Disabled
ADSPEC : Disabled
Use TE metric : Disabled
Exclude Grps :
None

```

```

Auto BW : Disabled
LdpOverRsvp : Enabled
IGP Shortcut: Enabled
IGP LFA : Disabled
BGPTransTun : Enabled
Oper Metric : 1500

VprnAutoBind : Enabled
BGP Shortcut : Enabled
IGP Rel Metric : Disabled

```

Prop Adm Grp: Enabled

Primary(a) : empty Up Time : 0d 01:19:30
 Bandwidth : 1 Mbps

=====
 *A:Dut-A# show router mpls lsp "AtoL1" path detail

=====
 MPLS LSP AtoL1 Path (Detail)

Legend :

@ - Detour Available	# - Detour In Use
b - Bandwidth Protected	n - Node Protected
s - Soft Preemption	
S - Strict	L - Loose
A - ABR	

 LSP AtoL1 Path empty

LSP Name : AtoL1	Path LSP ID : 13316
From : 10.20.1.1	To : 10.20.1.12
Adm State : Up	Oper State : Up
Path Name : empty	Path Type : Primary
Path Admin : Up	Path Oper : Up
OutInterface: 1/1/1	Out Label : 131069
Path Up Time: 0d 01:19:46	Path Dn Time: 0d 00:00:00
Retry Limit : 0	Retry Timer : 20 sec
RetryAttempt: 0	NextRetryIn : 0 sec

Adspec : Disabled	Oper Adspec : Disabled
CSPF : Enabled	Oper CSPF : Enabled
Least Fill : Disabled	Oper LeastF* : Disabled
FRR : Enabled	Oper FRR : Enabled
FRR NodePro*: Enabled	Oper FRR NP : Enabled
FR Hop Limit: 6	Oper FRHopL* : 6
FR Prop Adm*: Disabled	Oper FRProp* : Disabled
Prop Adm Grp: Enabled	Oper PropAG : Enabled
Inter-area : True	

Neg MTU : 1496	Oper MTU : 1496
Bandwidth : 1 Mbps	Oper Bw : 1 Mbps
Hop Limit : 255	Oper HopLim* : 255
Record Route: Record	Oper RecRou* : Record
Record Label: Record	Oper RecLab* : Record
SetupPriori* : 7	Oper SetupP* : 7
Hold Priori* : 0	Oper HoldPr* : 0
Class Type : 0	Oper CT : 0

MainCT Retry: n/a
 Rem :

MainCT Retry: 0
 Limit :

Include Grps:
 None
 Exclude Grps:
 None

Oper InclGr*:
 None
 Oper ExclGr*:
 None

Adaptive : Enabled

Oper Metric : 1500

```

Preference : n/a
Path Trans : 1
Failure Code: noError
ExplicitHops:
    No Hops Specified
Actual Hops :
    10.10.1.1(10.20.1.1) @ n
-> 10.10.1.2(10.20.1.2) @ n
-> 10.10.5.4(10.20.1.4) @ n
-> 10.20.1.7(10.20.1.7) @ n
-> 10.10.17.7 @ n
-> 10.20.1.9(10.20.1.9) @
-> 10.10.25.9 @
-> 10.20.1.12(10.20.1.12)
-> 10.10.33.12
ComputedHops:
    10.10.1.1(S)
-> 10.10.1.2(S)
-> 10.10.5.4(SA)
-> 10.20.1.12(L)
ResigEligib*: False
LastResignal: n/a
CSPF Queries: 3
Failure Node: n/a
CSPF Metric : 1500
=====
* indicates that the corresponding row element may have been truncated.

```

```
A:sim1>config>router>mpls>lsp$ show router mpls lsp path detail
```

```

=====MPLS LSP
Path (Detail)
=====
Legend :
    @ - Detour Available          # - Detour In Use
    b - Bandwidth Protected       n - Node Protected
    s - Soft Preemption
    S - Strict                    L - Loose
=====
-----
LSP 11 Path 1
-----
LSP Name      : 11                Path LSP ID : 30208
From          : 10.20.1.1         To           : 10.20.1.3
Adm State     : Up                Oper State   : Down
Path Name     : 1                 Path Type    :
Primary
Path Admin    : Up                Path Oper    : Down
OutInterface  : n/a               Out Label    : n/a
Path Up Time  : 0d 00:00:00       Path Dn Time : 0d 00:00:02
Retry Limit   : 0                 Retry Timer  : 30 sec
RetryAttempt  : 0                 NextRetryIn : 7 sec (Fast)
SetupPriorit*: 7                 Hold Priorit*: 0
Preference    : n/a
Bandwidth     : No Reservation     Oper Bw      : 0 Mbps
Hop Limit     : 255                Class Type   : 0
Backup CT     : None
MainCT Retry  : n/a                MainCT Retry: 0
Rem           :                    Limit        :
Oper CT       : None
Record Route  : Record             Record Label: Record
Oper MTU      : 0                  Neg MTU      : 0

```

```

Adaptive      : Enabled                               Oper Metric : 65535
Include Grps:                                     Exclude Grps:
None                                                  None
Path Trans   : 2                                    CSPF Queries: 0
Failure Code: noError                               Failure Node: n/a
ExplicitHops:
  10.20.1.2(S)
Actual Hops  :
  No Hops Specified
ResigEligib*: False
LastResignal: n/a                                  CSPF Metric : 0
=====

*A:# show router mpls lsp path detail
=====
MPLS LSP Path (Detail)
=====
Legend :
  @ - Detour Available          # - Detour In Use
  b - Bandwidth Protected      n - Node Protected
  s - Soft Preemption
  S - Strict                    L - Loose
=====

LSP to_C Path 1000_S
-----
LSP Name      : to_C                               Path LSP ID : 17926
From          : 10.20.1.2                           To          : 10.20.1.3
Adm State     : Up                                  Oper State  : Up
Path Name     : 1000_S                              Path Type   : Standby
Path Admin    : Up                                  Path Oper   : Up
OutInterface  : 1/1/2                               Out Label   : 131068
Path Up Time  : 0d 00:06:46                         Path Dn Time: 0d 00:00:00
Retry Limit   : 0                                    Retry Timer : 20 sec
RetryAttempt  : 0                                    NextRetryIn: 0 sec

Adspec       : Disabled                             Oper Adspec : Disabled
CSPF         : Enabled                               Oper CSPF   : Enabled
CSPF-FL     : Enabled                               Oper CSPF-FL: Enabled
Least Fill   : Disabled                             Oper LeastF*: Enabled
FRR NodePro*: Disabled                             Oper FRR NP : Enabled
Prop Adm Grp: Disabled                             Oper PropAG : Disabled
Neg MTU      : 1496                                 Oper MTU    : 1496
Bandwidth    : No Reservation                       Oper Bw     : 0 Mbps
Hop Limit    : 255                                 Oper HopLim*: 255
Record Route : Record                             Oper RecRou*: Record
Record Label : Record                             Oper RecLab*: Record
SetupPriori*: 7                                  Oper SetupP*: 7
Hold Priori* : 0                                  Oper HoldPr*: 0
Class Type   : 0
Backup CT    : None                                Oper CT     : 0
MainCT Retry: n/a
  Rem       :
MainCT Retry: n/a
  Limit    :
Include Grps:                                     Oper InclGr*:
  silver                                       silver
Exclude Grps:                                     Oper ExclGr*:
  None                                           None

Adaptive     : Enabled                               Oper Metric : 2999

```

```

Preference : 255
Path Trans : 0
Failure Code: noError
ExplicitHops:
    No Hops Specified
Actual Hops :
    10.10.4.2(10.20.1.2)
-> 10.10.4.4(10.20.1.4)
-> 10.10.6.5(10.20.1.5)
-> 10.10.5.3(10.20.1.3)
ComputedHops:
    10.10.4.2(S)      -> 10.10.4.4(S)      -> 10.10.6.5(S)
-> 10.10.5.3(S)
Srlg : Disabled
SrlgDisjoint: False
ResigEligib*: False
LastResignal: n/a
CSPF Queries: 0
Failure Node: n/a
Record Label : N/A
Record Label : 131068
Record Label : 131068
Record Label : 131065
CSPF Metric : 2999
=====

```

```
*A:Dut-C>config>router>mpls>lsp$ /show router mpls lsp path detail
```

```
=====
MPLS LSP Path (Detail)
=====
```

```
Legend :
```

```

@ - Detour Available          # - Detour In Use
b - Bandwidth Protected      n - Node Protected
s - Soft Preemption
S - Strict                    L - Loose

```

```
-----
LSP 2 Path 1
-----
```

```

LSP Name      : 2
From          : 10.20.1.3
Adm State     : Up
Path Name     : 1
Path Admin    : Up
OutInterface  : n/a
Path Up Time  : 0d 00:00:00
Retry Limit   : 0
RetryAttempt  : 1
Timeout In   : 19 sec

Path LSP ID   : 54272
To            : 10.20.1.1
Oper State    : Down
Path Type     : Primary
Path Oper     : In Progress
Out Label     : n/a
Path Dn Time  : 0d 00:00:13
Retry Timer   : 30 sec
NextRetryIn  : 0 sec

```

```

Adspec        : Disabled
CSPF          : Disabled
CSPF-FL       : Disabled
Least Fill    : Disabled
FRR           : Disabled
FR Hop Limit  : 16
Prop Adm Grp  : Disabled

Oper Adspec   : N/A
Oper CSPF     : N/A
Oper CSPF-FL  : N/A
Oper LeastF*  : N/A
Oper FRR      : N/A
Oper FRHopL*  : N/A
Oper PropAG   : N/A

```

```

Neg MTU       : 0
Bandwidth     : No Reservation
Hop Limit     : 255
Record Route  : Record
Record Label  : Record
SetupPriori* : 7
Hold Priori*  : 0
Class Type    : 0

Oper MTU      : N/A
Oper Bw       : N/A
Oper HopLim*  : N/A
Oper RecRou*  : N/A
Oper RecLab*  : N/A
Oper SetupP*  : N/A
Oper HoldPr*  : N/A
Oper CT       : N/A

```

```

Backup CT      : None
MainCT Retry: Infinite
      Rem      :
MainCT Retry: 0
      Limit    :
Include Grps:  Oper InclGr*:
None           N/A
Exclude Grps:  Oper ExclGr*:
None           N/A

Adaptive      : Enabled           Oper Metric : 65535
Preference    : n/a
Path Trans    : 0                CSPF Queries: 0
Failure Code: noError           Failure Node: n/a
ExplicitHops:
      10.10.2.1(S)
Actual Hops   :
      No Hops Specified
ResigEligib*: False
LastResignal: n/a                CSPF Metric : 0
=====
* indicates that the corresponding row element may have been truncated.

```

A:sim1>config>router>mpls>lsp\$ show router mpls lsp path detail

```

=====
MPLS LSP Path (Detail)
=====
Legend :
  @ - Detour Available           # - Detour In Use
  b - Bandwidth Protected        n - Node Protected
  s - Soft Preemption
  S - Strict                     L - Loose
=====
-----
LSP 11 Path 1
-----
LSP Name      : 11                Path LSP ID : 30208
From          : 10.20.1.1         To           : 10.20.1.3
Adm State     : Up                Oper State   : Down
Path Name     : 1                 Path Type    : Primary
Path Admin    : Up                Path Oper    : Down
OutInterface: n/a                 Out Label    : n/a
Path Up Time: 0d 00:00:00         Path Dn Time: 0d 00:00:02
Retry Limit   : 0                 Retry Timer  : 30 sec
RetryAttempt: 0                   NextRetryIn : 7 sec (Fast)
SetupPriori*: 7                   Hold Priori*: 0
Preference    : n/a
Bandwidth     : No Reservation     Oper Bw      : 0 Mbps
Hop Limit     : 255                Class Type   : 0
Backup CT     : None
MainCT Retry: n/a                 MainCT Retry: 0
      Rem      :                   Limit        :
Oper CT       : None
Record Route: Record              Record Label: Record
Oper MTU      : 0                  Neg MTU      : 0
Adaptive      : Enabled            Oper Metric  : 65535
Include Grps:  Exclude Grps:
None           None
Path Trans    : 2                  CSPF Queries: 0

```

```

Failure Code: noError                               Failure Node: n/a
ExplicitHops:
  10.20.1.2(S)
Actual Hops :
  No Hops Specified
ResigEligib*: False
LastResignal: n/a                                  CSPF Metric : 0
=====

```

```

*A:SRU4>config>router>mpls# show router mpls lsp path
=====
MPLS LSP Path (Detail)
=====
Legend :
  @ - Detour Available          # - Detour In Use
  b - Bandwidth Protected      n - Node Protected
  s - Soft Preemption
  S - Strict                    L - loose
=====
ExplicitHops:
  10.20.1.3(L)      -> 10.20.1.4(S)
Actual Hops :
  10.10.1.1(10.20.1.1)      Record Label : N/A
-> 10.10.1.2(10.20.1.2)    Record Label : 131071
-> 10.10.5.3(10.20.1.3)    Record Label : 131071
-> 10.10.7.4(10.20.1.4)    Record Label : 131071
-> 10.10.8.5(10.20.1.5)    Record Label : 131071
ComputedHops:
  10.10.1.1(S)      -> 10.10.1.2(S)      -> 10.10.5.3(S)
-> 10.20.1.4(S)      -> 10.20.1.5(L)
=====

```

```

*A:SRU4>config>router>mpls# show router mpls lsp
=====
MPLS LSPs (Originating)
=====
LSP Name                               To                               Fastfail   Adm   Opr
                                   Config
-----
to_110_20_1_1_cspf                     110.20.1.1                       No         Up    Up
to_110_20_1_2_cspf                     110.20.1.2                       No         Up    Dwn
to_110_20_1_3_cspf                     110.20.1.3                       No         Up    Up
to_110_20_1_4_cspf                     110.20.1.4                       No         Up    Dwn
to_110_20_1_5_cspf                     110.20.1.5                       No         Up    Up
to_110_20_1_6_cspf                     110.20.1.6                       No         Up    Dwn
to_110_20_1_110_cspf                   110.20.1.110                     No         Up    Up
to_10_8_100_15_cspf                    10.8.100.15                      No         Up    Dwn
to_10_20_1_20_cspf                     10.20.1.20                       No         Up    Up
to_10_20_1_22_cspf                     10.20.1.22                       No         Up    Up
to_10_100_1_1_cspf                     10.100.1.1                       No         Up    Dwn
to_110_20_1_1_cspf_2                   110.20.1.1                       No         Up    Up
to_110_20_1_1_cspf_3                   110.20.1.1                       No         Up    Up
to_110_20_1_1_cspf_4                   110.20.1.1                       No         Up    Up
to_110_20_1_1_cspf_5                   110.20.1.1                       No         Up    Up
to_110_20_1_1_cspf_6                   110.20.1.1                       No         Up    Up
to_110_20_1_1_cspf_7                   110.20.1.1                       No         Up    Up
to_110_20_1_1_cspf_8                   110.20.1.1                       No         Up    Up
to_110_20_1_1_cspf_9                   110.20.1.1                       No         Up    Up
to_110_20_1_1_cspf_10                  110.20.1.1                       No         Up    Up
to_110_20_1_1_cspf_11                  110.20.1.1                       No         Up    Up

```



```

to_110_20_1_1_cspf_12          110.20.1.1          No          Up          Up
to_110_20_1_1_cspf_13          110.20.1.1          No          Up          Up
to_110_20_1_1_cspf_14          110.20.1.1          No          Up          Up
to_110_20_1_1_cspf_15          110.20.1.1          No          Up          Up
...

```

```
-----
LSPs : 201
=====
```

```
*A:SRU4>config>router>mpls#
```

Label	Description
Auto BW	Enabled – Auto-bandwidth adjustment is configured on this LSP.
AB OpState	Up – Auto-bandwidth is operationally enabled on this LSP Down – Auto-bandwidth is operationally disabled on this LSP
Auto BW Min	The minimum bandwidth of the LSP that auto-bandwidth can request (in Mbps).
Auto BW Max	The maximum bandwidth of the LSP that auto-bandwidth can request (in Mbps).
AB Up Thresh	The percent threshold for increasing LSP bandwidth.
AB Down Thresh	The percent threshold for decreasing LSP bandwidth.
AB Up BW	The absolute bandwidth threshold for increasing LSP bandwidth (in Mbps).
AB Down BW	The absolute bandwidth threshold for decreasing LSP bandwidth (in Mbps).
AB Coll Intv	The auto-bandwidth collection interval.
AB Adj Mul	The adjust-multiplier for this LSP (may be configured or inherited).
AB Samp Mul	The sample-multiplier for this LSP (may be configured or inherited).
AB Adj Time	The adjust-multiplier times the collection-interval (in Mins).
AB Sample Time	The sample-multiplier times the collection-interval (in Mins).
AB Adj Cnt	The adjust count (number of whole collection intervals since the start of the current adjust interval).
AB Samp Cnt	The sample count (number of whole collection intervals since the start of the current sample interval).
AB Last Adj	The system time of the last auto-bandwidth adjustment.
AB Next Adj	The approximate remaining time in the current adjust interval (adjust-multiplier – adjust count) times the collection interval (in Mins). This overstates the actual remaining time because the elapsed time in the current collection interval is not accounted for.

Label	Description (Continued)
AB Adj Cause	The cause of the last auto-bandwidth adjustment: <ul style="list-style-type: none"> • none – no adjustment has occurred • manual • adj-count • overflow
AB Max AvgR*	The maximum average data rate in any sample interval of the current adjust interval.
AB Lst AvgR*	The average data rate measured in the sample interval that ended most recently.
AB Ovfl Lmt	The configured value of the auto-bandwidth overflow-limit.
AB Ovfl Cnt	The number of overflow samples since the last reset.
ABOvflThres	The percent threshold for declaring an overflow sample.
AB Ovfl BW	The absolute bandwidth threshold for declaring an overflow sample (in Mbps).
AB Monitor BW	True – monitor bandwidth is enabled on the LSP. False – monitor bandwidth is not enabled on the LSP.

```
*A:SRU4>config>router>mpls# show router mpls lsp detail
=====
MPLS LSPs (Originating) (Detail)
=====
Type : Originating
-----
LSP Name      : to_110_20_1_1_cspf
LSP Type      : RegularLsp
From          : 110.20.1.4
Adm State     : Up
LSP Up Time   : 0d 01:47:49
Transitions   : 11
Retry Limit   : 0
Signaling     : RSVP
Hop Limit     : 255
Adaptive      : Enabled
FastReroute   : Disabled
CSPF          : Enabled
Metric        : 0
Include Grps  :
None
Least Fill    : Disabled
LdpOverRsvp   : Enabled
IGP Shortcut  : Enabled
Oper Metric   : 1001

LSP Tunnel ID : 1
Oper State     : Up
LSP Down Time  : 0d 00:00:00
Path Changes   : 11
Retry Timer    : 30 sec
Resv. Style    : SE
Negotiated MTU : 1500
ClassType      : 0
Oper FR        : Disabled
ADSPEC        : Disabled
Use TE metric  : Disabled
Exclude Grps   :
None
VprnAutoBind  : Enabled

Primary(a)    : to_110_20_1_1
Bandwidth     : 0 Mbps

Up Time       : 0d 01:47:49
```

```

-----
...
-----
Type : Originating
-----
LSP Name      : to_10_100_1_1_cspf_20
LSP Type      : RegularLsp
From          : 110.20.1.4
Adm State     : Up
LSP Up Time   : 0d 00:00:00
Transitions   : 0
Retry Limit   : 0
Signaling     : RSVP
Hop Limit     : 255
Adaptive      : Enabled
FastReroute   : Disabled
CSPF          : Enabled
Metric        : 0
Include Grps  :
None
Least Fill    : Disabled
LdpOverRsvp   : Enabled
IGP Shortcut  : Enabled
Oper Metric   : 65535

LSP Tunnel ID : 201
Oper State     : Down
LSP Down Time : 0d 13:30:49
Path Changes   : 0
Retry Timer    : 30 sec
Resv. Style    : SE
Negotiated MTU : 0
ClassType      : 0
Oper FR        : Disabled
ADSPEC         : Disabled
Use TE metric  : Disabled
Exclude Grps   :
None

VprnAutoBind  : Enabled

Primary       : to_10_100_1_1
Bandwidth     : 0 Mbps
Down Time     : 0d 13:30:49
=====
*A:SRU4>config>router>mpls#

*A:SRU4>config>router>mpls# show router mpls lsp path detail
=====
MPLS LSP Path (Detail)
=====
Legend :
  @ - Detour Available      # - Detour In Use
  b - Bandwidth Protected   n - Node Protected
  s - Soft Preemption
=====
-----
LSP to_110_20_1_1_cspf Path to_110_20_1_1
-----
LSP Name      : to_110_20_1_1_cspf
From          : 110.20.1.4
Adm State     : Up
Path Name     : to_110_20_1_1
Path Admin    : Up
OutInterface  : 3/2/1
Path Up Time  : 0d 01:43:19
Retry Limit   : 0
RetryAttempt  : 0
SetupPriorit : 7
Preference    : n/a
Bandwidth     : No Reservation
Hop Limit     : 255
Backup CT     : None
MainCT Retry  : n/a
Rem          :
Oper CT       : 0
Record Route  : Record

Path LSP ID   : 12856
To            : 110.20.1.1
Oper State    : Up
Path Type     : Primary
Path Oper     : Up
Out Label     : 336302
Path Dn Time  : 0d 00:00:00
Retry Timer   : 30 sec
NextRetryIn  : 0 sec
Hold Priorit : 0

Oper Bw       : 0 Mbps
Class Type    : 0

MainCT Retry  : 0
Limit        :

Record Label  : Record

```

```

Oper MTU      : 1500                      Neg MTU      : 1500
Adaptive     : Enabled                    Oper Metric  : 1001
Include Grps :                            Exclude Grps :
None                                                None
Path Trans   : 13                         CSPF Queries: 56
Failure Code : noError                    Failure Node : n/a
ExplicitHops :
    No Hops Specified
Actual Hops  :
    10.100.30.4(110.20.1.4)                Record Label : N/A
-> 10.100.30.20(10.20.1.20)                Record Label : 336302
-> 10.100.14.1(110.20.1.1)                Record Label : 126325
ComputedHops :
    10.100.30.4 -> 10.100.30.20 -> 10.100.14.1
ResigEligib* : False
LastResignal : n/a                        CSPF Metric  : 1001
Last MBB     :
    MBB Type  : TimerBasedResignal        MBB State   : Fail
    Ended At  : 03/04/2010 08:53:40      Old Metric  : 0

```

```

...
LSP to_10_100_1_1_cspf_20 Path to_10_100_1_1

```

```

-----
LSP Name      : to_10_100_1_1_cspf_20      Path LSP ID  : 40960
From          : 110.20.1.4                 To           : 10.100.1.1
Adm State     : Up                         Oper State   : Down
Path Name     : to_10_100_1_1              Path Type    : Primary
Path Admin    : Up                         Path Oper    : Down
OutInterface  : n/a                        Out Label    : n/a
Path Up Time  : 0d 00:00:00                 Path Dn Time : 0d 13:26:06
Retry Limit   : 0                           Retry Timer  : 30 sec
RetryAttempt  : 1612                        NextRetryIn  : 19 sec
SetupPriori* : 7                            Hold Priori* : 0
Preference    : n/a
Bandwidth     : No Reservation              Oper Bw      : 0 Mbps
Hop Limit     : 255                         Class Type   : 0
Backup CT     : None
MainCT Retry  : Infinite                    MainCT Retry : 0
    Rem       :                             Limit       :
Oper CT       : None
Record Route  : Record                      Record Label : Record
Oper MTU      : 0                           Neg MTU      : 0
Adaptive     : Enabled                    Oper Metric  : 65535
Include Grps :                            Exclude Grps :
None                                                None
Path Trans   : 0                         CSPF Queries: 0
Failure Code  : noCspfRouteOwner          Failure Node  : 110.20.1.4
ExplicitHops :
    No Hops Specified
Actual Hops  :
    No Hops Specified
ComputedHops :
    No Hops Specified
ResigEligib* : False
LastResignal : n/a                        CSPF Metric  : 0

```

```

=====
* indicates that the corresponding row element may have been truncated.
*A:SRU4>config>router>mpls#

```

```

*A:SRU4>config>router>mpls# show router mpls lsp "to_110_20_1_1_cspf"

```

```

=====
MPLS LSPs (Originating)
=====
LSP Name                               To                Fastfail    Adm   Opr
                               Fastfail          Config
-----
to_110_20_1_1_cspf                    110.20.1.1       No          Up    Up
-----
LSPs : 1
=====
*A:SRU4>config>router>mpls#

```

```

*A:SRU4>config>router>mpls# show router mpls lsp "to_110_20_1_1_cspf" detail
=====
MPLS LSPs (Originating) (Detail)
=====
Type : Originating
-----
LSP Name      : to_110_20_1_1_cspf
LSP Type      : RegularLsp
From          : 110.20.1.4
Adm State     : Up
LSP Up Time   : 0d 01:47:02
Transitions   : 11
Retry Limit    : 0
Signaling     : RSVP
Hop Limit     : 255
Adaptive      : Enabled
FastReroute   : Disabled
CSPF          : Enabled
Metric        : 0
Include Grps  :
None
Least Fill    : Disabled
LdpOverRsvp   : Enabled
IGP Shortcut  : Enabled
Oper Metric   : 1001

LSP Tunnel ID : 1
Oper State     : Up
LSP Down Time : 0d 00:00:00
Path Changes   : 11
Retry Timer    : 30 sec
Resv. Style    : SE
Negotiated MTU : 1500
ClassType      : 0
Oper FR        : Disabled
ADSPEC         : Disabled
Use TE metric  : Disabled
Exclude Grps   :
None
VprnAutoBind  : Enabled

Primary(a)    : to_110_20_1_1
Bandwidth     : 0 Mbps
Up Time       : 0d 01:47:02
=====
*A:SRU4>config>router>mpls#

```

```

*A:SRU4>config>router>mpls# show router mpls lsp detail to 110.20.1.2
=====
MPLS LSPs (Originating) (Detail)
=====
Type : Originating
-----
LSP Name      : 1
LSP Type      : RegularLsp
From          : 0.0.0.0
Adm State     : Down
LSP Up Time   : 0d 00:00:00
Transitions   : 0
Retry Limit    : 0
Signaling     : RSVP

LSP Tunnel ID : 1
Oper State     : Down
LSP Down Time : 0d 00:00:07
Path Changes   : 0
Retry Timer    : 30 sec
Resv. Style    : SE

```

```

Hop Limit      : 255
Adaptive       : Enabled
FastReroute    : Disabled
CSPF           : Disabled
Metric         : 0
Include Grps   :
None
Least Fill     : Disabled

Auto BW        : Enabled
Auto BW Min    : 0 Mbps
AB Up Thresh   : 5 percent
AB Up BW       : 0 Mbps
AB Curr BW     : 0 Mbps
AB Adj Mul     : 288+
AB Adj Time    : 0 Mins
AB Adj Cnt     : 0
AB Last Adj    : n/a
ABMaxAvgRt     : 0 Mbps
AB Ovfl Lmt    : 0
ABOvflThres   : 0 percent
AB Adj Cause   : none
LdpOverRsvp    : Enabled
IGP Shortcut   : Enabled
Oper Metric    : 65535

Negotiated MTU : 0
ClassType      : 0
Oper FR        : Disabled
ADSPEC         : Disabled

Exclude Grps   :
None

AB OpState     : Down
Auto BW Max    : 100000 Mbps
AB Down Thresh : 5 percent
AB Down BW     : 0 Mbps
AB Samp Intv   : 0
AB Samp Mul    : 1+
AB Samp Time   : 0 Mins
AB Samp Cnt    : 0
AB Next Adj    : 0 Mins
AB Lst AvgRt   : 0 Mbps
AB Ovfl Cnt    : 0
AB Ovfl BW     : 0
AB Monitor BW  : False
VprnAutoBind   : Enabled

```

+ indicates inherited values

```

=====
*A:SRU4>config>router>mpls#

```

```

*A:SRU4>config>router>mpls# show router mpls lsp count

```

```

=====
MPLS LSP Count

```

```

=====
Originate      Transit      Terminate
-----
Static LSPs    0            136         0
Dynamic LSPs   140          421        1620
Detour LSPs    0            0           0
P2MP S2Ls     0            0           0
=====

```

```

*A:SRU4>config>router>mpls#

```

```

*A:SRU4>config>router>mpls# show router mpls lsp path mbb

```

```

=====
MPLS LSP Paths

```

```

-----
LSP to_110_20_1_1_cspf Path to_110_20_1_1

```

```

-----
LastResignal: n/a                                CSPF Metric : 1001
Last MBB      :
  MBB Type    : TimerBasedResignal              MBB State   : Fail
  Ended At    : 03/04/2010 09:23:58           Old Metric   : 0
-----

```

```

LSP to_110_20_1_2_cspf Path to_110_20_1_2

```

```

-----
LastResignal: 03/04/2010 09:23:58              CSPF Metric : 65535
-----

```

```

LSP to_110_20_1_3_cspf Path to_110_20_1_3
-----
LastResignal: n/a                                CSPF Metric : 1001
Last MBB      :
  MBB Type    : TimerBasedResignal                MBB State   : Fail
  Ended At    : 03/04/2010 09:23:58              Old Metric  : 0
-----

LSP to_110_20_1_4_cspf Path to_110_20_1_4
-----
LastResignal: n/a                                CSPF Metric : 0
-----

LSP to_110_20_1_5_cspf Path to_110_20_1_5
-----
...
-----
LastResignal: n/a                                CSPF Metric : 0
-----

LSP to_10_100_1_1_cspf_19 Path to_10_100_1_1
-----
LastResignal: n/a                                CSPF Metric : 0
-----

LSP to_10_100_1_1_cspf_20 Path to_10_100_1_1
-----
LastResignal: n/a                                CSPF Metric : 0
=====
*A:SRU4>config>router>mpls#

In Prog MBB :
  MBB Type    : SoftPreemption                    NextRetryIn : 19 sec
  Started At  : 12/08/2008 22:21:11              RetryAttempt: 0
  FailureCode: noError                            Failure Node: n/a
=====
*A:Dut-B#

*A:SRU4>config>router>mpls# show router mpls lsp transit
=====
MPLS LSPs (Transit)
=====
Legend : @ - Active Detour
=====
From          To          In I/F    Out I/F    State LSP Name
-----
110.20.1.5    10.20.1.22  3/2/1    3/2/7     Up    to_10_20_1_22_cspf::to*
110.20.1.5    10.20.1.20  3/2/7    3/2/1     Up    to_10_20_1_20_cspf_3::*
110.20.1.5    10.20.1.20  3/2/7    3/2/1     Up    to_10_20_1_20_cspf_4::*
110.20.1.5    10.20.1.20  3/2/7    3/2/1     Up    to_10_20_1_20_cspf_2::*
110.20.1.5    10.20.1.20  3/2/7    3/2/1     Up    to_10_20_1_20_cspf_20::*
110.20.1.5    10.20.1.20  3/2/7    3/2/1     Up    to_10_20_1_20_cspf_18::*
110.20.1.5    10.20.1.20  3/2/7    3/2/1     Up    to_10_20_1_20_cspf_19::*
110.20.1.5    10.20.1.20  3/2/7    3/2/1     Up    to_10_20_1_20_cspf_17::*
110.20.1.5    10.20.1.20  3/2/7    3/2/1     Up    to_10_20_1_20_cspf_16::*
110.20.1.5    10.20.1.20  3/2/7    3/2/1     Up    to_10_20_1_20_cspf_15::*
110.20.1.5    10.20.1.20  3/2/7    3/2/1     Up    to_10_20_1_20_cspf_13::*
110.20.1.5    10.20.1.20  3/2/7    3/2/1     Up    to_10_20_1_20_cspf_14::*
110.20.1.5    10.20.1.20  3/2/7    3/2/1     Up    to_10_20_1_20_cspf_12::*
110.20.1.5    10.20.1.20  3/2/7    3/2/1     Up    to_10_20_1_20_cspf_10::*
...
110.20.1.3    10.20.1.22  aps-1    3/2/7     Up    to_10_20_1_22_cspf_6::*
110.20.1.3    10.20.1.22  aps-1    3/2/7     Up    to_10_20_1_22_cspf::to*
110.20.1.3    10.20.1.22  aps-1    3/2/7     Up    to_10_20_1_22_cspf_9::*

```

```

-----
LSPs : 520
=====
* indicates that the corresponding row element may have been truncated.
*A:SRU4>config>router>mpls#
*A:SRU4>config>router>mpls# show router mpls lsp terminate
=====
MPLS LSPs (Terminate)
=====
Legend : @ - Active Detour
=====
From          To          In I/F    Out I/F    State LSP Name
-----
110.20.1.5    110.20.1.4  3/2/1     n/a        Up    b4-1::b4-1
110.20.1.5    110.20.1.4  3/2/7     n/a        Up    gsr::gsr
10.20.1.22    110.20.1.4  3/2/7     n/a        Up    gsr2_t10
110.20.1.6    110.20.1.4  3/2/3:10  n/a        Up    1::2
110.20.1.6    110.20.1.4  3/2/3:3   n/a        Up    1::stby
110.20.1.6    110.20.1.4  3/2/3:10  n/a        Up    2::2
110.20.1.6    110.20.1.4  3/2/3:6   n/a        Up    2::stby
110.20.1.6    110.20.1.4  3/2/3:10  n/a        Up    3::2
110.20.1.6    110.20.1.4  3/2/3:6   n/a        Up    3::stby
...
110.20.1.3    110.20.1.4  aps-1     n/a        Up    to_110_20_1_4_cspf_20:*
110.20.1.3    110.20.1.4  aps-1     n/a        Up    to_110_20_1_4_cspf_4:*
-----
LSPs : 1603
=====
* indicates that the corresponding row element may have been truncated.
*A:SRU4>config>router>mpls#

*A:SRU4>config>router>mpls# show router mpls lsp terminate detail
=====
MPLS LSPs (Terminate) (Detail)
=====
-----
LSP b4-1::b4-1
-----
From          : 110.20.1.5          To          : 110.20.1.4
State         : Up
SetupPriority : 7                   Hold Priority : 0
Class Type    : 0
In Interface  : 3/2/1              In Label    : 131071
Previous Hop  : 10.100.30.20
-----
LSP gsr::gsr
-----
From          : 110.20.1.5          To          : 110.20.1.4
State         : Up
SetupPriority : 7                   Hold Priority : 0
Class Type    : 0
In Interface  : 3/2/7              In Label    : 128547
Previous Hop  : 160.60.60.2
-----
...
-----
From          : 110.20.1.3          To          : 110.20.1.4
State         : Up
SetupPriority : 7                   Hold Priority : 0
Class Type    : 0
In Interface  : aps-1              In Label    : 130409

```



```

Previous Hop          : 104.104.0.3
=====
*A:SRU4>config>router>mpls#

```

lsp-egress-stats

Syntax **lsp-egress-stats**
lsp-egress-stats *lsp-name*

Context show>router>mpls

Description This command displays MPLS LSP egress statistics information.

Output **Sample Output**

```
*A:Dut-C>config>router>mpls>lsp$ show router mpls lsp-egress-stats lsp "1"
```

```

=====
MPLS LSP Egress Statistics
=====
-----
LSP Name          : 1
-----
Collect Stats   : Enabled                Accting Plcy. : Default
Adm State       : Up                     PSB Match     : True
FC BE
InProf Pkts    : 0                      OutProf Pkts  : 0
InProf Octets  : 0                      OutProf Octets: 0
FC L2
InProf Pkts    : 0                      OutProf Pkts  : 0
InProf Octets  : 0                      OutProf Octets: 0
FC AF
InProf Pkts    : 0                      OutProf Pkts  : 0
InProf Octets  : 0                      OutProf Octets: 0
FC L1
InProf Pkts    : 0                      OutProf Pkts  : 0
InProf Octets  : 0                      OutProf Octets: 0
FC H2
InProf Pkts    : 0                      OutProf Pkts  : 0
InProf Octets  : 0                      OutProf Octets: 0
FC EF
InProf Pkts    : 0                      OutProf Pkts  : 0
InProf Octets  : 0                      OutProf Octets: 0
FC H1
InProf Pkts    : 0                      OutProf Pkts  : 0
InProf Octets  : 0                      OutProf Octets: 0
FC NC
InProf Pkts    : 0                      OutProf Pkts  : 0
InProf Octets  : 0                      OutProf Octets: 0
=====

```

```
*A:Dut-C# show router mpls lsp-egress-stats lsp "ipmsi-1-73728"
```

```

=====
MPLS LSP Egress Statistics
=====

```

```

-----
LSP Name      : ipmsi-1-73728
-----
Collect Stats : Enabled          Accting Plcy. : Default
Adm State     : Up              PSB Match    : True
FC BE
InProf Pkts   : 0              OutProf Pkts : 0
InProf Octets : 0              OutProf Octets: 0
FC L2
InProf Pkts   : 0              OutProf Pkts : 0
InProf Octets : 0              OutProf Octets: 0
FC AF
InProf Pkts   : 0              OutProf Pkts : 0
InProf Octets : 0              OutProf Octets: 0
FC L1
InProf Pkts   : 0              OutProf Pkts : 0
InProf Octets : 0              OutProf Octets: 0
FC H2
InProf Pkts   : 0              OutProf Pkts : 0
InProf Octets : 0              OutProf Octets: 0
FC EF
InProf Pkts   : 0              OutProf Pkts : 0
InProf Octets : 0              OutProf Octets: 0
FC H1
InProf Pkts   : 0              OutProf Pkts : 0
InProf Octets : 0              OutProf Octets: 0
FC NC
InProf Pkts   : 0              OutProf Pkts : 0
InProf Octets : 0              OutProf Octets: 0
=====

```

lsp-ingress-stats

- Syntax** **lsp-ingress-stats**
lsp-ingress-stats *ip-address* **lsp** *lsp-name*
- Context** show>router>mpls
- Description** This command displays MPLS LSP ingress statistics information.

Sample Output

```

*A:Dut-A# show router mpls lsp-ingress-stats lsp "1" sender 10.20.1.3
=====
MPLS LSP Ingress Statistics
=====
-----
LSP Name      : 1
Sender       : 10.20.1.3
-----
Collect Stats : Disabled          Accting Plcy. : None
Adm State     : Up              PSB Match    : True

```

```

FC BE
InProf Pkts : 0
InProf Octets : 0
OutProf Pkts : 0
OutProf Octets: 0
FC L2
InProf Pkts : 0
InProf Octets : 0
OutProf Pkts : 0
OutProf Octets: 0
FC AF
InProf Pkts : 0
InProf Octets : 0
OutProf Pkts : 0
OutProf Octets: 0
FC L1
InProf Pkts : 0
InProf Octets : 0
OutProf Pkts : 0
OutProf Octets: 0
FC H2
InProf Pkts : 0
InProf Octets : 0
OutProf Pkts : 0
OutProf Octets: 0
FC EF
InProf Pkts : 0
InProf Octets : 0
OutProf Pkts : 0
OutProf Octets: 0
FC H1
InProf Pkts : 0
InProf Octets : 0
OutProf Pkts : 0
OutProf Octets: 0
FC NC
InProf Pkts : 0
InProf Octets : 0
OutProf Pkts : 0
OutProf Octets: 0
=====

*A:Dut-A# show router mpls lsp-ingress-stats lsp "ipmsi-1-73728" sender 10.20.1.3
=====
MPLS LSP Ingress Statistics
=====
-----
LSP Name      : ipmsi-1-73728
Sender        : 10.20.1.3
-----
Collect Stats : Disabled
Adm State     : Up
Accting Plcy. : None
PSB Match     : True
FC BE
InProf Pkts : 0
InProf Octets : 0
OutProf Pkts : 0
OutProf Octets: 0
FC L2
InProf Pkts : 0
InProf Octets : 0
OutProf Pkts : 0
OutProf Octets: 0
FC AF
InProf Pkts : 0
InProf Octets : 0
OutProf Pkts : 0
OutProf Octets: 0
FC L1
InProf Pkts : 0
InProf Octets : 0
OutProf Pkts : 0
OutProf Octets: 0
FC H2
InProf Pkts : 0
InProf Octets : 0
OutProf Pkts : 0
OutProf Octets: 0
FC EF
InProf Pkts : 0
InProf Octets : 0
OutProf Pkts : 0
OutProf Octets: 0
FC H1
InProf Pkts : 0
InProf Octets : 0
OutProf Pkts : 0
OutProf Octets: 0
FC NC
InProf Pkts : 0
InProf Octets : 0
OutProf Pkts : 0
OutProf Octets: 0
=====

```

```
*A:Dut-A>config>router>mpls>ingr-stats# show router mpls lsp-ingress-stats
type p2mp active template-match
```

```
=====
MPLS LSP Ingress Statistics
=====
-----
LSP Name       : ipmsi-1-73728
Sender        : 10.20.1.3
-----
Collect Stats : Disabled          Accting Plcy. : None
Adm State     : Up                PSB Match    : True
FC BE
InProf Pkts  : 0                  OutProf Pkts : 0
InProf Octets : 0                  OutProf Octets : 0
FC L2
InProf Pkts  : 0                  OutProf Pkts : 0
InProf Octets : 0                  OutProf Octets : 0
FC AF
InProf Pkts  : 0                  OutProf Pkts : 0
InProf Octets : 0                  OutProf Octets : 0
FC L1
InProf Pkts  : 0                  OutProf Pkts : 0
InProf Octets : 0                  OutProf Octets : 0
FC H2
InProf Pkts  : 0                  OutProf Pkts : 0
InProf Octets : 0                  OutProf Octets : 0
FC EF
InProf Pkts  : 0                  OutProf Pkts : 0
InProf Octets : 0                  OutProf Octets : 0
FC H1
InProf Pkts  : 0                  OutProf Pkts : 0
InProf Octets : 0                  OutProf Octets : 0
FC NC
InProf Pkts  : 0                  OutProf Pkts : 0
InProf Octets : 0                  OutProf Octets : 0
-----
LSP Statistics : 1
```

Isp-template

Syntax `isp-template [isp-template-name] [detail]`

Context `show>router>mpls`

Description This command displays MPLS LSP template information.

Sample Output

```
*A:Dut-C# show router mpls lsp-template detail
=====
MPLS LSP Templates (Detail)
=====
-----
```

```

LSP Template : ipmsi
-----
Type           : P2MP           Admin State      : Up
Default Path   : path_ipmsi     Adaptive         : Enabled
Bandwidth      : 0 Mbps         Hop Limit        : 255
CSPF           : Enabled        Use TE metric    : Disabled
Include Groups :                 Exclude Groups   :
None
FastReroute    : Enabled
FR Method      : Facility       FR Hop Limit     : 16
Record Route   : Record        Record Label     : Record
Retry Limit    : 0              Retry Timer      : 30 sec
LSP Count      : 3              Ref Count        : 3
=====

```

oam-template

- Syntax** **oam-template**
- Context** show>router>mpls>mpls-tp
- Description** This command displays MPLS-TP OAM template information.

Sample Output

```

*A:mlstp-dutA# show router mpls mpls-tp oam-template
=====
MPLS-TP OAM Templates
=====
Template Name : privatebed-oam-template Router ID      : 1
BFD Template  : privatebed-bfd-template Hold-Down Time: 0 centiseconds
                                           Hold-Up Time  : 20 deciseconds
=====

```

protection-template

- Syntax** **protection-template**
- Context** show>router>mpls>mpls-tp
- Description** This command displays MPLS-TP protection template information.

Sample Output

```

*A:mlstp-dutA# show router mpls mpls-tp protection-template
=====
MPLS-TP Protection Templates
=====
Template Name : privatebed-protection-template Router ID      : 1
Protection Mode: one2one                        Direction        : bidirectional
Revertive     : revertive                       Wait-to-Restore: 300sec

```

```
Rapid-PSC-Timer: 10ms                               Slow-PSC-Timer : 5sec
=====
```

status

- Syntax** **status**
- Context** show>router>mpls>mpls-tp
- Description** This command displays MPLS-TP system configuration information.

Sample Output

```
*A:mlstp-dutA# show router mpls mpls-tp status

=====
MPLS-TP Status
=====
Admin Status   : Up
Global ID      : 42
Node ID        : 0.0.3.233
Tunnel Id Min  : 1
Tunnel Id Max  : 4096
=====
```

transit-path

- Syntax** **transit-path** [*path-name*] [**detail**]
- Context** show>router>mpls>mpls-tp
- Description** This command displays MPLS-TP tunnel information.
- Parameters** *path-name* — Specifies the path name, up to 32 characters max.

Sample Output

```
A:mlstp-dutC# show router mpls mpls-tp transit-path
<path-name>
"tp-32" "tp-33" "tp-34" "tp-35" "tp-36" "tp-37" "tp-38" "tp-39"
"tp-40" "tp-41"
detail

A:mlstp-dutC# show router mpls mpls-tp transit-path "tp-32"

=====
MPLS-TP Transit tp-32 Path Information
=====
Path Name      : tp-32
Admin State    : Up
Oper State     : Up

-----
Path           NextHop           InLabel  OutLabel  Out I/F
-----
FP             2080             2081     CtoB_1
```

```

RP                2081      2080      CtoA_1
=====
A:mplstp-dutC# show router mpls mpls-tp transit-path "tp-32" detail

=====
MPLS-TP Transit tp-32 Path Information (Detail)
=====
Path Name       : tp-32
Admin State    : Up
Oper State     : Up
-----
Path ID configuration
Src Global ID  : 42
Src Node ID   : 0.0.3.234
LSP Number    : 2
Dst Global ID  : 42
Dst Node ID   : 0.0.3.233
Dst Tunnel Num: 32

Forward Path configuration
In Label      : 2080
Out Interface : CtoB_1
Out Label     : 2081
Next Hop Addr: n/a

Reverse Path configuration
In Label      : 2081
Out Interface : CtoA_1
Out Label     : 2080
Next Hop Addr: n/a
=====
A:mplstp-dutC#

```

p2mp-info

- Syntax** `p2mp-info [type {originate | transit | terminate}] [s2l-endpoint ip-address]`
- Context** `show>router>mpls`
- Description** This command displays P2MP cross-connect information.
- Parameters** `type` — Specifies the P2MP type.
- Values**
- originate** — Specifies to display the static LSPs that originate at this virtual router.
 - transit** — Specifies to display the static LSPs that transit through this virtual router.
 - terminate** — Specifies to display the static LSPs that terminate at this virtual router.

Sample Output

```

*A:Dut-C# show router mpls p2mp-info

=====
MPLS P2MP Cross Connect Information
=====
-----
S2L ipmsi-4000-73729::path_ipmsi

```

```
-----  
Source IP Address      : 10.20.1.1          Tunnel ID      : 61441  
P2MP ID               : 4000                    Lsp ID        : 29696  
S2L Name              : ipmsi-4000-73729::pa* To      : 10.20.1.3  
In Interface         : 1/1/1                    In Label      : 262129  
Num. of S2ls        : 1  
-----
```

S2L ipmsi-65535-73730::path_ipmsi

```
-----  
Source IP Address      : 10.20.1.1          Tunnel ID      : 61442  
P2MP ID               : 65535                    Lsp ID        : 30208  
S2L Name              : ipmsi-65535-73730::p* To      : 10.20.1.3  
In Interface         : 1/1/1                    In Label      : 262128  
Num. of S2ls        : 1  
-----
```

S2L ipmsi-1001-73728::path_ipmsi

```
-----  
Source IP Address      : 10.20.1.1          Tunnel ID      : 61440  
P2MP ID               : 1001                    Lsp ID        : 35840  
S2L Name              : ipmsi-1001-73728::pa* To      : 10.20.1.3  
In Interface         : 1/1/1                    In Label      : 262127  
Num. of S2ls        : 1  
-----
```

S2L ipmsi-1001-73732::path_ipmsi

```
-----  
Source IP Address      : 10.20.1.2          Tunnel ID      : 64944  
P2MP ID               : 1001                    Lsp ID        : 34816  
S2L Name              : ipmsi-1001-73732::pa* To      : 10.20.1.3  
In Interface         : 1/1/2                    In Label      : 262114  
Num. of S2ls        : 1  
-----
```


S2L ipmsi-4000-73729::path_ipmsi

```

-----
Source IP Address      : 10.20.1.3          Tunnel ID      : 61441
P2MP ID               : 4000                Lsp ID        : 16384
S2L Name              : ipmsi-4000-73729::pa* To      : 10.20.1.1
Out Interface         : 1/1/1                Out Label     : 262131
Num. of S2ls         : 1
-----

```

S2L ipmsi-4000-73729::path_ipmsi

```

-----
Source IP Address      : 10.20.1.3          Tunnel ID      : 61441
P2MP ID               : 4000                Lsp ID        : 16384
S2L Name              : ipmsi-4000-73729::pa* To      : 10.20.1.4
Out Interface         : 2/1/1                Out Label     : 262121
Num. of S2ls         : 1
-----

```

S2L ipmsi-1001-73728::path_ipmsi

```

-----
Source IP Address      : 10.20.1.3          Tunnel ID      : 61440
P2MP ID               : 1001                Lsp ID        : 22016
S2L Name              : ipmsi-1001-73728::pa* To      : 10.20.1.1
Out Interface         : 1/1/1                Out Label     : 262129
Num. of S2ls         : 1
-----

```

S2L ipmsi-1001-73728::path_ipmsi

```

-----
Source IP Address      : 10.20.1.3          Tunnel ID      : 61440
P2MP ID               : 1001                Lsp ID        : 22016
S2L Name              : ipmsi-1001-73728::pa* To      : 10.20.1.2
Out Interface         : 1/1/2                Out Label     : 262115
-----

```

Num. of S2ls : 1

S2L ipmsi-1001-73728::path_ipmsi

Source IP Address : 10.20.1.3 Tunnel ID : 61440
P2MP ID : 1001 Lsp ID : 22016
S2L Name : ipmsi-1001-73728::pa* To : 10.20.1.4
Out Interface : 2/1/1 Out Label : 262108
Num. of S2ls : 2

S2L ipmsi-1001-73728::path_ipmsi

Source IP Address : 10.20.1.3 Tunnel ID : 61440
P2MP ID : 1001 Lsp ID : 22016
S2L Name : ipmsi-1001-73728::pa* To : 10.20.1.5
Out Interface : 2/1/1 Out Label : 262108
Num. of S2ls : 2

S2L ipmsi-65535-73730::path_ipmsi

Source IP Address : 10.20.1.3 Tunnel ID : 61442
P2MP ID : 65535 Lsp ID : 46592
S2L Name : ipmsi-65535-73730::p* To : 10.20.1.1
Out Interface : 1/1/1 Out Label : 262130
Num. of S2ls : 1

S2L ipmsi-65535-73730::path_ipmsi

Source IP Address : 10.20.1.3 Tunnel ID : 61442
P2MP ID : 65535 Lsp ID : 46592
S2L Name : ipmsi-65535-73730::p* To : 10.20.1.4

```

Out Interface      : 2/1/1                Out Label      : 262109
Num. of S2ls      : 1

```

```
-----
P2MP Cross-connect instances : 12

```

p2mp-lsp

- Syntax** **p2mp-lsp** [*lsp-name*] [**detail**]
p2mp-lsp [*lsp-name*] **p2mp-instance** [*p2mp-instance-name*] [**mbb**]
p2mp-lsp [*lsp-name*] **p2mp-instance** [*p2mp-instance-name*] **s2l** [*s2l-name* [**to** *s2l-to-address*]] [**status** {**up** | **down**}] [**detail**]
p2mp-lsp [*lsp-name*] **p2mp-instance** [*p2mp-instance-name*] **s2l** [*s2l-name* [**to** *s2l-to-address*]] **mbb**
- Context** show>router>mpls
- Description** This command displays MPLS P2MP LSP information.
- Parameters** *lsp-name* — Specifies the name of the LSP used in the path.
- p2mp-instance**[*p2mp-instance-name* — Specifies the administrative name for the P2MP instance which must be unique within a virtual router instance.
- mbb** — Specifies to display make-before-break (MBB) information.
- s2l** — Specifies the source-to-leaf (S2L) name.
- to** *s2l-to-address* —
- status** — Displays the status of the p2mp LSP.
- Values** up — Displays the total time that this S2L has been operational.
down — Displays the total time that this S2L has not been operational.

Sample Output

```

*A:Dut-C# show router mpls p2mp-lsp
- p2mp-lsp [<lsp-name>] [detail]
- p2mp-lsp [<lsp-name>] p2mp-instance [<p2mp-instance-name>] [mbb]
- p2mp-lsp [<lsp-name>] p2mp-instance [<p2mp-instance-name>] s2l [<s2l-name>
  [to <s2l-to-address>]][status {up|down}] [detail]
- p2mp-lsp [<lsp-name>] p2mp-instance [<p2mp-instance-name>] s2l [<s2l-name>
  [to <s2l-to-address>]] <mbb>
- p2mp-lsp using-template [lsp-template <template-name>] [detail]

```

```

<lsp-name>          : [64 chars max] - accepts * as wildcard char
<p2mp-instance>    : keyword
<p2mp-instance-name> : [max 32 chars]
<s2l>              : keyword
<s2l-name>         : [max 32 chars]
<up|down>         : keywords
<detail>          : keyword
<mbb>            : keyword
<s2l-to-address>  : [a.b.c.d]
<using-template>  : keyword
<lsp-template>   : [32 chars max]

```

```
*A:Dut-C# show router mpls p2mp-lsp
```

```

=====
MPLS P2MP LSPs (Originating)
=====

```

LSP Name	Adm	Opr
ipmsi-1001-73728	Up	Up
ipmsi-4000-73729	Up	Up
ipmsi-65535-73730	Up	Up

```
-----
LSPs : 3
=====
```

```
*A:Dut-C# show router mpls p2mp-lsp detail
```

```

=====
MPLS P2MP LSPs (Originating) (Detail)
=====

```

 Type : Originating

LSP Name : ipmsi-1001-73728
 LSP Type : P2mpAutoLsp LSP Tunnel ID : 61440
 From : 10.20.1.3
 Adm State : Up Oper State : Up
 LSP Up Time : 6d 21:08:37 LSP Down Time : 0d 00:00:00
 Transitions : 1 Path Changes : 1
 Retry Limit : 0 Retry Timer : 30 sec
 Signaling : RSVP Resv. Style : SE
 Hop Limit : 255 Negotiated MTU : n/a
 Adaptive : Enabled ClassType : 0
 FastReroute : Enabled Oper FR : Enabled
 FR Method : Facility FR Hop Limit : 16
 FR Bandwidth: 0 Mbps FR Node Protect: Disabled
 FR Object : Enabled
 CSPF : Enabled ADSPEC : Disabled
 Metric : Disabled Use TE metric : Disabled
 Include Grps: Exclude Grps :
 None None
 Least Fill : Disabled
 Auto BW : Disabled
 LdpOverRsvp : Disabled VprnAutoBind : Disabled
 IGP Shortcut: Disabled BGP Shortcut : Disabled
 BGPTransTun : Disabled
 Oper Metric : Disabled
 Prop Adm Grp: Disabled CSPFFirstLoose : Disabled
 P2MPInstance: 1001 P2MP-Inst-type : Primary

```

S2L Cfg Cou*: 4
S2L Oper Count*: 4
S2l-Name : path_ipmsi To : 10.20.1.1
S2l-Name : path_ipmsi To : 10.20.1.2
S2l-Name : path_ipmsi To : 10.20.1.4
S2l-Name : path_ipmsi To : 10.20.1.5

```

```

-----
Type : Originating
-----

```

```

LSP Name : ipmsi-4000-73729
LSP Type : P2mpAutoLsp LSP Tunnel ID : 61441
From : 10.20.1.3
Adm State : Up Oper State : Up
LSP Up Time : 6d 21:08:38 LSP Down Time : 0d 00:00:00
Transitions : 1 Path Changes : 1
Retry Limit : 0 Retry Timer : 30 sec
Signaling : RSVP Resv. Style : SE
Hop Limit : 255 Negotiated MTU : n/a
Adaptive : Enabled ClassType : 0
FastReroute : Enabled Oper FR : Enabled
FR Method : Facility FR Hop Limit : 16
FR Bandwidth: 0 Mbps FR Node Protect: Disabled
FR Object : Enabled
CSPF : Enabled ADSPEC : Disabled
Metric : Disabled Use TE metric : Disabled
Include Grps: Exclude Grps :
None None
Least Fill : Disabled
Auto BW : Disabled
LdpOverRsvp : Disabled VprnAutoBind : Disabled
IGP Shortcut: Disabled BGP Shortcut : Disabled

```

```

BGPTransTun : Disabled
Oper Metric : Disabled
Prop Adm Grp: Disabled
CSPFFirstLoose : Disabled

P2MPInstance: 4000
P2MP-Inst-type : Primary
S2L Cfg Cou*: 2
S2L Oper Count*: 2
S2l-Name : path_ipmsi To : 10.20.1.1
S2l-Name : path_ipmsi To : 10.20.1.4

```

```

-----
Type : Originating
-----

```

```

LSP Name : ipmsi-65535-73730
LSP Type : P2mpAutoLsp LSP Tunnel ID : 61442
From : 10.20.1.3
Adm State : Up Oper State : Up
LSP Up Time : 6d 21:08:39 LSP Down Time : 0d 00:00:00
Transitions : 1 Path Changes : 1
Retry Limit : 0 Retry Timer : 30 sec
Signaling : RSVP Resv. Style : SE
Hop Limit : 255 Negotiated MTU : n/a
Adaptive : Enabled ClassType : 0
FastReroute : Enabled Oper FR : Enabled
FR Method : Facility FR Hop Limit : 16
FR Bandwidth: 0 Mbps FR Node Protect: Disabled
FR Object : Enabled
CSPF : Enabled ADSPEC : Disabled
Metric : Disabled Use TE metric : Disabled
Include Grps: Exclude Grps :
None None
Least Fill : Disabled

```

```

Auto BW      : Disabled
LdpOverRsvp  : Disabled
IGP Shortcut: Disabled
BGP TransTun: Disabled
Oper Metric  : Disabled
Prop Adm Grp: Disabled

VprnAutoBind : Disabled
BGP Shortcut  : Disabled

CSPFFirstLoose : Disabled

P2MPInstance: 65535
P2MP-Inst-type : Primary
S2L Cfg Cou*: 2
S2L Oper Count*: 2
S2l-Name      : path_ipmsi
To             : 10.20.1.1
S2l-Name      : path_ipmsi
To             : 10.20.1.4

```

=====

* indicates that the corresponding row element may have been truncated.

*A:Dut-C#

*A:sim1>config>router>mpls>lsp\$ show router mpls p2mp-lsp p2mp-instance s2l detail

=====

MPLS LSP S2L (Detail)

Legend :

```

@ - Detour Available          # - Detour In Use
b - Bandwidth Protected      n - Node Protected
S - Strict                    L - Loose
s - Soft Preemption

```

LSP 1 S2L 1

```

-----
LSP Name      : 1
P2MP ID       : 0
Adm State     : Up
S2l State:    : Inactive
S2L Name      : 1
S2l Admin     : Up
OutInterface  : n/a
S2L Up Time   : 0d 00:00:00
RetryAttempt  : 0
S2L Trans     : 8
Failure Code  : noError
ExplicitHops  :
  10.20.1.2(S)
Actual Hops   :
  No Hops Specified
ComputedHops  :
  No Hops Specified

S2l LSP ID    : 26624
S2l Grp Id    : 0
Oper State    : Down
To            : 10.20.1.3
S2l Oper      : Down
Out Label     : n/a
S2L Dn Time   : 0d 00:00:01
NextRetryIn   : 9 sec (Fast)
CSPF Queries  : 4
Failure Node  : n/a

```


LastResignal: n/a

show router mpls p2mp-lsp p2mp-instance s2l detail

LSP 2 S2L 2

```

-----
LSP Name      : 2                      S2l LSP ID   : 52230
P2MP ID       : 0                      S2l Grp Id   : 2
Adm State     : Up                     Oper State   : Up
S2l State     : Active                  :
S2L Name      : 2                      To           : 10.20.1.3
S2l Admin     : Up                     S2l Oper    : Up
OutInterface  : 1/1/1                  Out Label    : 131071
S2L Up Time   : 0d 00:04:43           S2L Dn Time  : 0d 00:00:00
RetryAttempt  : 0                      NextRetryIn  : 0 sec
S2L Trans     : 5                      CSPF Queries : 21
Failure Code  : tunnelLocallyRepaired  Failure Node : 10.20.1.2
ExplicitHops  :
    10.20.1.2(S)
Actual Hops   :
    10.10.1.1(10.20.1.1)              Record Label  : N/A
-> 10.10.1.2(10.20.1.2) @ #          Record Label  : 131071
-> 10.10.6.3(10.20.1.3)              Record Label  : 131068
ComputedHops  :
    10.10.1.1(S)      -> 10.10.1.2(S)  -> 10.10.5.3(S)
LastResignal  : n/a
In Prog MBB   :
  MBB Type    : GlobalRevert          NextRetryIn   : n/a
  Timeout In  : 23 sec
  Started At  : 06/29/2011 11:06:09   RetryAttempt  : 7
  FailureCode : noError                Failure Node  : n/a
-----

```

*A:Dut-C>config>router>mpls>lsp\$ /show router mpls lsp path detail

MPLS LSP Path (Detail)

Legend :

```

@ - Detour Available          # - Detour In Use
b - Bandwidth Protected      n - Node Protected
s - Soft Preemption
S - Strict                    L - Loose

```

LSP 1 Path 1

```

-----
LSP Name      : 1                      Path LSP ID  : 56320
From          : 10.20.1.3              To           : 10.10.1.1
Adm State     : Up                     Oper State   : Up
Path Name     : 1                      Path Type    : Primary
Path Admin    : Up                     Path Oper    : Up
OutInterface  : 1/1/1                  Out Label    : 131071
Path Up Time  : 0d 00:03:09           Path Dn Time : 0d 00:00:00
Retry Limit   : 0                      Retry Timer  : 30 sec
RetryAttempt  : 0                      NextRetryIn  : 0 sec
-----

```

```

SetupPriori*: 7                               Hold Priori*: 0
Preference   : n/a                             Oper Bw      : 0 Mbps
Bandwidth    : No Reservation                   Class Type   : 0
Hop Limit    : 255                             MainCT Retry: 0
Backup CT    : None                            Limit       :
MainCT Retry: n/a                             Record Label: Record
Rem          :                                 Neg MTU     : 1496
Oper CT      : 0                               Oper Metric : 1000
Record Route: Record                          Exclude Grps:
Oper MTU     : 1496                            None
Adaptive     : Enabled                        CSPF Queries: 3
Include Grps: None                            Failure Node: 10.20.1.3
None
Path Trans   : 1
Failure Code: badNode

Oper Values :
Setup Prior*: 7                               Hold Priori*: 0
Record Route: Record                          Record Label: Record
Hop Limit    : 255
Adspec       : Disabled                       CSPFToFirst*: Disabled
CSPF         : Enabled                        FR Node Pro*: Disabled
Least Fill   : Disabled
Prop Adm Grp: Disabled
Include Grps: None                            Exclude Grps:
None

ExplicitHops:
  No Hops Specified
Actual Hops :
  10.10.2.3(10.20.1.3) @ #                    Record Label : N/A
  -> 10.10.1.1(10.20.1.1)                    Record Label : 131071
ComputedHops:
  10.10.2.3(S)                               -> 10.10.2.1(S)
ResigEligib*: False
LastResignal: n/a                             CSPF Metric : 1000
In Prog MBB :
  MBB Type   : GlobalRevert                   NextRetryIn : 0 sec
  Timeout In : 22 sec
  Started At : 08/26/2011 23:59:29           RetryAttempt: 2
  FailureCode: noError                       Failure Node: n/a
  Signaled BW: 0 Mbps

```

=====
 * indicates that the corresponding row element may have been truncated.

```
show router mpls p2mp-lsp p2mp-instance s2l detail
```

```
-----
LSP 2 S2L 2
-----
```

```

LSP Name      : 2                               S2L LSP ID   : 52230
P2MP ID       : 0                               S2L Grp Id   : 4
Adm State     : Up                             Oper State    : Down
S2L State:    : Inactive
S2L Name      : 2                               To           : 10.20.1.3
S2L Admin     : Up                             S2L Oper     : In Progress
OutInterface  : n/a                            Out Label    : n/a
S2L Up Time   : 0d 00:00:00                    S2L Dn Time  : 0d 00:00:20
RetryAttempt  : 1                             NextRetryIn  : n/a
Timeout In    : 21 sec
S2L Trans     : 6                             CSPF Queries : 27

```

```

Failure Code: noError                               Failure Node: n/a
ExplicitHops:
    10.20.1.2(S)
Actual Hops :
    No Hops Specified
LastResignal: n/a
    
```

```
*A:Dut-C# show router mpls p2mp-lsp
```

```
=====
MPLS P2MP LSPs (Originating)
=====
```

LSP Name	Adm	Opr
ipmsi-1001-73728	Up	Up
ipmsi-4000-73729	Up	Up
ipmsi-65535-73730	Up	Up

```
-----
LSPs : 3
=====
```

```
*A:Dut-C# show router mpls p2mp-lsp detail
```

```
=====
MPLS P2MP LSPs (Originating) (Detail)
=====
```

```
-----
Type : Originating
-----
```

```

LSP Name      : ipmsi-1001-73728
LSP Type      : P2mpAutoLsp
From          : 10.20.1.3
Adm State     : Up
LSP Up Time   : 6d 21:08:37
Transitions   : 1
Retry Limit    : 0
Signaling     : RSVP
Hop Limit     : 255
Adaptive      : Enabled
FastReroute   : Enabled
FR Method     : Facility
FR Bandwidth  : 0 Mbps
FR Object     : Enabled
CSPF          : Enabled
Metric        : Disabled
Include Grps  :
None
Least Fill    : Disabled
Auto BW       : Disabled
LdpOverRsvp  : Disabled
IGP Shortcut  : Disabled
BGPTransTun  : Disabled
Oper Metric   : Disabled
Prop Adm Grp : Disabled
P2MPInstance : 1001
S2L Cfg Cou* : 4
S2L-Name     : path_ipmsi
S2L-Name     : path_ipmsi
S2L-Name     : path_ipmsi
S2L-Name     : path_ipmsi

LSP Tunnel ID : 61440
Oper State     : Up
LSP Down Time  : 0d 00:00:00
Path Changes   : 1
Retry Timer    : 30 sec
Resv. Style    : SE
Negotiated MTU : n/a
ClassType      : 0
Oper FR        : Enabled
FR Hop Limit   : 16
FR Node Protect : Disabled
ADSPEC         : Disabled
Use TE metric  : Disabled
Exclude Grps   :
None
VprnAutoBind  : Disabled
BGP Shortcut   : Disabled
CSPFFirstLoose : Disabled
P2MP-Inst-type : Primary
S2L Oper Count* : 4
To             : 10.20.1.1
To             : 10.20.1.2
To             : 10.20.1.4
To             : 10.20.1.5
    
```

```
-----
Type : Originating
-----
```

```

LSP Name       : ipmsi-4000-73729
LSP Type       : P2mpAutoLsp
From           : 10.20.1.3
Adm State      : Up
LSP Up Time    : 6d 21:08:38
Transitions    : 1
Retry Limit    : 0
Signaling      : RSVP
Hop Limit      : 255
Adaptive       : Enabled
FastReroute    : Enabled
FR Method      : Facility
FR Bandwidth   : 0 Mbps
FR Object      : Enabled
CSPF           : Enabled
Metric         : Disabled
Include Grps   :
None
Least Fill    : Disabled
Auto BW       : Disabled
LdpOverRsvp   : Disabled
IGP Shortcut   : Disabled
BGPTransTun   : Disabled
Oper Metric    : Disabled
Prop Adm Grp  : Disabled
P2MPInstance  : 4000
S2L Cfg Cou*  : 2
S2L-Name      : path_ipmsi
S2L-Name      : path_ipmsi

LSP Tunnel ID  : 61441
Oper State     : Up
LSP Down Time  : 0d 00:00:00
Path Changes   : 1
Retry Timer    : 30 sec
Resv. Style    : SE
Negotiated MTU : n/a
ClassType     : 0
Oper FR        : Enabled
FR Hop Limit   : 16
FR Node Protect : Disabled

ADSPEC         : Disabled
Use TE metric  : Disabled
Exclude Grps   :
None

VprnAutoBind  : Disabled
BGP Shortcut   : Disabled

CSPFFirstLoose : Disabled
P2MP-Inst-type : Primary
S2L Oper Count* : 2
To             : 10.20.1.1
To             : 10.20.1.4

```

Type : Originating

```

LSP Name       : ipmsi-65535-73730
LSP Type       : P2mpAutoLsp
From           : 10.20.1.3
Adm State      : Up
LSP Up Time    : 6d 21:08:39
Transitions    : 1
Retry Limit    : 0
Signaling      : RSVP
Hop Limit      : 255
Adaptive       : Enabled
FastReroute    : Enabled
FR Method      : Facility
FR Bandwidth   : 0 Mbps
FR Object      : Enabled
CSPF           : Enabled
Metric         : Disabled
Include Grps   :
None
Least Fill    : Disabled
Auto BW       : Disabled
LdpOverRsvp   : Disabled
IGP Shortcut   : Disabled
BGPTransTun   : Disabled
Oper Metric    : Disabled
Prop Adm Grp  : Disabled
P2MPInstance  : 65535
S2L Cfg Cou*  : 2
S2L-Name      : path_ipmsi
S2L-Name      : path_ipmsi

LSP Tunnel ID  : 61442
Oper State     : Up
LSP Down Time  : 0d 00:00:00
Path Changes   : 1
Retry Timer    : 30 sec
Resv. Style    : SE
Negotiated MTU : n/a
ClassType     : 0
Oper FR        : Enabled
FR Hop Limit   : 16
FR Node Protect : Disabled

ADSPEC         : Disabled
Use TE metric  : Disabled
Exclude Grps   :
None

VprnAutoBind  : Disabled
BGP Shortcut   : Disabled

CSPFFirstLoose : Disabled
P2MP-Inst-type : Primary
S2L Oper Count* : 2
To             : 10.20.1.1
To             : 10.20.1.4

```

```

=====
* indicates that the corresponding row element may have been truncated.

*A:Dut-C#
*A:siml>config>router>mpls>lsp$ show router mpls p2mp-lsp p2mp-instance s2l detail
=====
MPLS LSP S2L (Detail)
=====
Legend :
  @ - Detour Available           # - Detour In Use
  b - Bandwidth Protected       n - Node Protected
  S - Strict                    L - Loose
  s - Soft Preemption

=====
-----
LSP 1 S2L 1
-----
LSP Name      : 1                S2L LSP ID   : 26624
P2MP ID       : 0                S2L Grp Id  : 0
Adm State     : Up              Oper State   : Down
S2L State     : Inactive
S2L Name      : 1                To           : 10.20.1.3
S2L Admin     : Up              S2L Oper    : Down
OutInterface  : n/a             Out Label    : n/a
S2L Up Time   : 0d 00:00:00     S2L Dn Time  : 0d 00:00:01
RetryAttempt  : 0                NextRetryIn  : 9 sec (Fast)
S2L Trans     : 8                CSPF Queries : 4
Failure Code  : noError         Failure Node  : n/a
ExplicitHops  :
    10.20.1.2(S)
Actual Hops   :
    No Hops Specified
ComputedHops  :
    No Hops Specified
LastResignal : n/a
=====

A:ALU-25# show router mpls p2mp-lsp lsp_1
=====
MPLS LSPs (Originating)
=====
LSP Name      To/P2MP ID      Fastfail      Adm   Opr
                Config
-----
lsp_1         18                Yes           Up    Up
-----
LSPs : 1
=====

A:ALU-25#
A:ALU-25# show router mpls p2mp-lsp Test_p2mp detail
=====
MPLS P2MP LSPs (Originating) (Detail)
=====
-----
Type : Originating
-----
LSP Name      : lsp_1                LSP Tunnel ID : 1
From          : 10.10.1.1            P2MP ID       : 18
Adm State     : Up                  Oper State     : Down
LSP Up Time   : 0d 00:00:00         LSP Down Time : 0d 20:39:48

```

```

Transitions : 0
Retry Limit : 0
Signaling : RSVP
Hop Limit : 255
FastReroute : Disabled
FR Method : Facility
FR Bandwidth: 0 Mbps
FR Object : Enabled
CSPF : Disabled
Metric : 1
Include Grps:
None

Path Changes : 0
Retry Timer : 30 sec
Resv. Style : FF
Adaptive : Enabled
Oper FR : Disabled
FR Hop Limit : 45
FR Node Protect: Disabled

ADSPEC : Disabled
Use TE metric : Disabled
Exclude Grps :
None

P2MPinstance:Test_p2mp p2mp-inst-type : primary

S2L Name :Test-s211 To : 10.20.1.6
S2L Name :Test-s212 To : 10.20.1.5
S2L Name :Test-s213 To : 10.20.1.4

```

A:ALU-25#

A:ALU-25# show router mpls p2mp-lsp Test_p2mp

=====

MPLS P2MP Instance (Originating)

Type : Originating

```

-----
LSP Name : lsp_1 LSP Tunnel ID : 1
P2MP ID : 18 Path LSP ID : 18
Adm State : Up Oper State : Down

P2MPinstance:Test_p2mp p2mp-inst-type : primary
Inst Name : lsp_1 P2MP Inst ID : 1
Adm State : Up Oper State : Down
Inst Up Time: 0d 00:00:00 Inst Down Time : 0d 20:39:48
Hop Limit : 255 Adaptive : Enabled
Record Route: Record Record Label : Record
Include Grps: Exclude Grps :
None None
Bandwidth : 0 Mbps Oper Bw : 0 Mbps

S2L Name :Test-s211 To : 10.20.1.6
S2L Name :Test-s212 To : 10.20.1.5
S2L Name :Test-s213 To : 10.20.1.4

```

A:ALU-25#

Note that the normal output is in detailed format only. There is no separate detail format.

A:ALU-52# show router mpls p2mp-lsp [p2mp-lsp-name] p2mp-instance [p2mp-inst-name]

=====

MPLS P2MP Instance (Originating)

Type : Originating

```

-----
LSP Name : lsp_1 LSP Tunnel ID : 1
P2MP ID : 18 Path LSP ID : 18
Adm State : Up Oper State : Down

```

```

P2MPinstance:Test_p2mp
Inst Name      : lsp_1
Adm State      : Up
Inst Up Time   : 0d 00:00:00
Hop Limit      : 255
Record Route   : Record
Include Grps   :
None
Bandwidth      : 0 Mbps

p2mp-inst-type : primary
P2MP Inst ID   : 1
Oper State     : Down
Inst Down Time : 0d 20:39:48
Adaptive       : Enabled
Record Label   : Record
Exclude Grps   :
None
Oper Bw        : 0 Mbps

S2L Name       :Test-s211      To      : 10.20.1.6
S2L Name       :Test-s212      To      : 10.20.1.5
S2L Name       :Test-s213      To      : 10.20.1.4

```

A:ALU-52#

```

A:ALU-52# show router mpls p2mp-lsp [p2mp-lsp-name] p2mp-instance [p2mp-inst-name]
mbb

```

```

=====
MPLS P2MP Instance (Originating)
=====

```

```

-----
Type : Originating
-----

```

```

LSP Name       : lsp_1
P2MP ID        : 18
Adm State      : Up

LSP Tunnel ID  : 1
Path LSP ID    : 18
Oper State     : Down

```

```

P2MPinstance:Test_p2mp
Inst Name      : lsp_1
Adm State      : Up
Inst Up Time   : 0d 00:00:00
Hop Limit      : 255
Record Route   : Record
Include Grps   :
None
Bandwidth      : 0 Mbps
Last MBB       :
MBB type       :
ended at       :
In Prog MBB    :
MBB type       :
Started at     :
Failure code    :

p2mp-inst-type : primary
P2MP Inst ID   : 1
Oper State     : Down
Inst Down Time : 0d 20:39:48
Adaptive       : Enabled
Record Label   : Record
Exclude Grps   :
None
Oper Bw        : 0 Mbps

Mbb State      :
Old Metric     :

Next Retry In  :
Retry Attempt  :
Failure Node   :

```

```

S2L Name       :Test-s211      To      : 10.20.1.6
S2L Admin      :
Failure code    :
S2L Oper       :
Failure Node    : 10.12.1.1

```

```

S2L Name       :Test-s211      To      : 10.20.1.6
S2L Admin      :
Failure code    :
S2L Oper       :
Failure Node    : 10.12.1.1

```

A:ALU-52#

```

A:ALU-52# show router mpls p2mp-lsp [p2mp-lsp-name] p2mp-instance [p2mp-inst-name]
s2l [s2l-name]
=====

```

```

MPLS S2Ls (Originating)
=====
S2L Name                               To           Next Hop     Adm   Opr
-----
Test-s2l1                               10.20.1.6   10.10.1.2   Up    Up
-----

```

LSPs : 1

A:ALU-52#

```

A:ALU-52# show router mpls p2mp-lsp [p2mp-lsp-name] p2mp-instance [p2mp-inst-name]
s2l [s2l-name] detail

```

```

=====
MPLS S2Ls (Originating) (Detail)
=====

```

```

-----
Type : Originating
-----

```

```

LSP Name       : lsp_1                               LSP Tunnel ID : 1
P2MP ID        : 18                                   Path LSP ID   : 18
Adm State      : Up                                   Oper State    : Down

```

P2MP Primary Instance:

```

Inst Name      : lsp_1                               P2MP Inst ID  : 1
Adm State      : Up                                   Oper State    : Down

```

```

S2L Name       : Test-s2l1                           To             : 10.20.1.6
Adm State      : Up                                   Oper State     : Down
OutInterface   : 1/1/1                               Out Label      : 131071
S2L Up Time    : 0d 00:00:00                         S2L Down Time : 0d 20:39:48
Transitions    : 0                                    Path Changes   : 0
Retry Limit    : 0                                    Retry Timer    : 30 sec
RetryAttempt   : 0                                    NextRetryIn   : 0 sec
Bandwidth      : No Reservation                       Oper Bw        : 0 Mbps
Hop Limit      : 255                                  Adaptive       : Enabled
Record Route   : Record                              Record Label   : Record
Oper MTU       : 1496                                 Neg MTU        : 1496
FastReroute    : Disabled                             Oper FR        : Disabled
FR Method      : Facility                             FR Hop Limit   : 45
FR Bandwidth   : 0 Mbps                               FR Node Protect: Disabled
FR Object      : Enabled
CSPF           : Disabled                             ADSPEC         : Disabled
Metric         : 1                                    Use TE metric  : Disabled
Include Grps   :                                     Exclude Grps   :
None
CSPF Queries  : 9
Failure Code   : noError                             Failure Node   : n/a
ExplicitHops   :
    No Hops Specified
Actual Hops    :
    10.10.1.1(10.20.1.1) @
    -> 10.10.1.2(10.20.1.2)
ComputedHops   :
    10.10.1.1 -> 10.10.1.2
LastResignal   : n/a                                 CSPF Metric    : 1000
-----

```

A:ALU-52#

```

*A:Dut-C# show router mpls p2mp-lsp "ipmsi-1-73752" detail

```



```

=====
MPLS P2MP LSPs (Originating) (Detail)
=====
-----
Type : Originating
-----
LSP Name       : ipmsi-1-73752
LSP Type       : P2mpAutoLsp
From           : 10.20.1.3
Adm State      : Up
LSP Up Time    : 0d 00:00:51
Transitions    : 3
Retry Limit    : 0
Signaling      : RSVP
Hop Limit      : 255
Adaptive       : Enabled
FastReroute    : Enabled
FR Method      : Facility
FR Node Pro*   : Disabled
FR Object      : Enabled
Egress Stats   : Enabled
CSPF           : Enabled
Metric         : Disabled
Include Grps   :
None
Least Fill    : Disabled

Auto BW        : Disabled
LdpOverRsvp    : Enabled
IGP Shortcut   : Enabled
IGP LFA        : Disabled
BGPTransTun    : Enabled
Oper Metric    : Disabled
Prop Adm Grp   : Disabled

LSP Tunnel ID  : 61445
Oper State     : Up
LSP Down Time  : 0d 00:00:00
Path Changes   : 3
Retry Timer    : 30 sec
Resv. Style    : SE
Negotiated MTU : n/a
ClassType      : 0
Oper FR        : Enabled
FR Hop Limit   : 16
FR Prop Adm Grp : Disabled

Egress Oper St* : Out-of-resource
ADSPEC         : Disabled
Use TE metric  : Disabled
Exclude Grps   :
None

VprnAutoBind   : Enabled
BGP Shortcut   : Enabled
IGP Rel Metric : Disabled

P2MPInstance: 1
S2L Cfg Cou* : 4
S2L-Name      : path_ipmsi
S2L-Name      : path_ipmsi
S2L-Name      : path_ipmsi
S2L-Name      : path_ipmsi
P2MP-Inst-type : Primary
S2L Oper Count* : 4
To             : 10.20.1.1
To             : 10.20.1.2
To             : 10.20.1.5
To             : 10.20.1.6
=====

```

srlg-database

- Syntax** `srlg-database [router-id ip-address] [interface ip-address]`
- Context** `show>router>mpls`
- Description** This command displays MPLS SRLG database information.
- Parameters** `router-id ip-address` — Specifies a 32-bit integer uniquely identifying the router in the Autonomous System. By convention to ensure uniqueness, this may default to the value of one of the router's IPv4 host addresses, represented as a 32-bit unsigned integer, if IPv4 is configured on the router. The router-id can be either the local one or some remote router.

interface *ip-address* — Specifies the IP address of the interface.

path

- Syntax** `path [path-name] [lsp-binding]`
- Context** `show>router>mpls`
- Description** This command displays MPLS paths.
- Parameters** *path-name* — The unique name label for the LSP path.
lsp-binding — Keyword to display binding information.
- Output** **MPLS Path Output** — The following table describes MPLS Path output fields.

Label	Description
Path Name	The unique name label for the LSP path.
Adm	Down — The path is administratively disabled. Up — The path is administratively enabled.
Hop Index	The value used to order the hops in a path.
IP Address	The IP address of the hop that the LSP should traverse on the way to the egress router.
Strict/Loose	Strict — The LSP must take a direct path from the previous hop router to the next router. Loose — The route taken by the LSP from the previous hop to the next hop can traverse through other routers.
LSP Name	The name of the LSP used in the path.
Binding	Primary — The preferred path for the LSP. Secondary — The standby path for the LSP.
Paths	Total number of paths configured.

Sample Output

```
*A:SRU4>config>router>mpls# show router mpls path
=====
MPLS Path:
=====
Path Name                Adm  Hop Index  IP Address  Strict/Loose
-----
to_110_20_1_1            Up   no hops    n/a         n/a
to_110_20_1_2            Up   no hops    n/a         n/a
to_110_20_1_3            Up   no hops    n/a         n/a
to_110_20_1_4            Up   no hops    n/a         n/a
to_110_20_1_5            Up   no hops    n/a         n/a
```

```

to_110_20_1_6                Up    no hops    n/a        n/a
to_110_20_1_110             Up    no hops    n/a        n/a
to_10_8_100_15              Up    no hops    n/a        n/a
to_10_20_1_20               Up    no hops    n/a        n/a
to_10_20_1_22               Up    no hops    n/a        n/a
to_10_100_1_1               Up    no hops    n/a        n/a
-----

```

Paths : 11

```

=====
*A:SRU4>config>router>mpls#

```

```

*A:SRU4>config>router>mpls# show router mpls path lsp-binding
=====

```

MPLS Path:

```

=====
Path Name                    Opr  LSP Name                    Binding
-----
to_110_20_1_1                Up   to_110_20_1_1_cspf         Primary
                               Up   to_110_20_1_1_cspf_2       Primary
                               Up   to_110_20_1_1_cspf_3       Primary
Up  to_110_20_1_1_cspf_16    Up   to_110_20_1_1_cspf_17       Primary
                               Up   to_110_20_1_1_cspf_18       Primary
                               Up   to_110_20_1_1_cspf_19       Primary
                               Up   to_110_20_1_1_cspf_20       Primary
to_110_20_1_2                Up   to_110_20_1_2_cspf         Primary
                               Up   to_110_20_1_2_cspf_2       Primary
                               Up   to_110_20_1_2_cspf_3       Primary
                               Up   to_110_20_1_2_cspf_4       Primary
                               Up   to_110_20_1_2_cspf_5       Primary
...
to_10_100_1_1                Down to_10_100_1_1_cspf         Primary
                               Down to_10_100_1_1_cspf_2       Primary
                               Down to_10_100_1_1_cspf_3       Primary
                               Down to_10_100_1_1_cspf_4       Primary
                               Down to_10_100_1_1_cspf_5       Primary
                               Down to_10_100_1_1_cspf_6       Primary
Down to_10_100_1_1_cspf_13    Down to_10_100_1_1_cspf_14       Primary
                               Down to_10_100_1_1_cspf_15       Primary
                               Down to_10_100_1_1_cspf_16       Primary
                               Down to_10_100_1_1_cspf_17       Primary
                               Down to_10_100_1_1_cspf_18       Primary
                               Down to_10_100_1_1_cspf_19       Primary
                               Down to_10_100_1_1_cspf_20       Primary
-----

```

Paths : 11

```

=====
*A:SRU4>config>router>mpls#

```

srlg-group

- Syntax** `srlg-group [group-name]`
- Context** `show>router>mpls`
- Description** This command displays MPLS SRLG groups
- Parameters** *group-name* — Specifies the name of the SRLG group within a virtual router instance.
- Output** **MPLS SRLG Group Output** — The following table describes MPLS SRLG group output fields

Label	Description
Group Name	Displays the name of the SRLG group within a virtual router instance.
Group Value	Displays the group value associated with this SRLG group.
Interface	Displays the interface where the SRLG groups is associated.
No. of Groups	Displays the total number of SRLG groups associated with the output.

Sample Output

```
*A:SRU4>config>router>mpls# show router mpls srlg-group
=====
MPLS Srlg Groups
=====
Group Name                Group Value  Interfaces
-----
1432                      1432        srl-1
1433                      1433        srl-3
1434                      1434        aps-8
1435                      1435        aps-9
2410                      2410        srr-1
2411                      2411        srr-2
2412                      2412        srr-3
3410                      3410        aps-1
3420                      3420        aps-2
3430                      3430        aps-3
3440                      3440        sr4-1
41.80                    4180        g7600
41104                    41104       germ-1
415.70                   41570       gsr1
420.40                   42040       m160
422.60                   42260       gsr2
44.200                   44200       hubA
45100                    45100       ess-7-1
45110                    45110       ess-7-2
45120                    45120       ess-7-3
4651                    4651        src-1.1
4652                    4652        src-1.2
4653                    4653        src-1.3
4654                    4654        src-1.4
4655                    4655        src-1.5
4656                    4656        src-1.6
4657                    4657        src-1.7
4658                    4658        src-1.8
```

```

4659                               4659          src-1.9
4660                               4660          src-1.10
-----
No. of Groups: 30
=====
*A:SRU4>config>router>mpls#

*A:SRU4>config>router>mpls# show router mpls srlg-group "1432"
=====
MPLS Srlg Groups
-----
Group Name                          Group Value  Interfaces
-----
1432                                 1432        srl-1
-----
No. of Groups: 1
=====
*A:SRU4>config>router>mpls#

```

static-lsp

Syntax **static-lsp** [*/sp-name*]
static-lsp {**transit** | **terminate**}
static-lsp **count**

Context show>router>mpls

Description This command displays MPLS static LSP information.

Output **MPLS Static LSP Output** — The following table describes MPLS static LSP output fields.

Label	Description
Lsp Name	The name of the LSP used in the path.
To	The system IP address of the egress router for the LSP.
Next Hop	The system IP address of the next hop in the LSP path.
In I/F	The ingress interface.
Out Label	The egress interface.
Out I/F	The egress interface.
Adm	Down — The path is administratively disabled. Up — The path is administratively enabled.
Opr	Down — The path is operationally down. Up — The path is operationally up.
LSPs	The total number of static LSPs.

Sample Output

```
A:ALA-12# show router mpls static-lsp
=====
MPLS Static LSPs (Originating)
=====
Lsp Name          To          Next Hop      Out Label  Out I/F    Adm  Opr
-----
NYC_SJC_customer2 100.20.1.10 10.10.1.4     1020      1/1/1     Up   Up

LSPs : 1
=====
A:ALA-12#
```

```
*A:SRU4>config>router>mpls# show router mpls static-lsp transit
=====
MPLS Static LSPs (Transit)
=====
In Label   In Port    Out Label    Out Port    Next Hop          Adm  Opr
-----
240        aps-1      440          1/1/10      11.22.11.3       Up   Up
241        aps-1      441          1/1/10      11.22.11.3       Up   Up
242        aps-1      442          1/1/10      11.22.11.3       Up   Up
243        aps-1      443          1/1/10      11.22.11.3       Up   Up
244        aps-1      444          1/1/10      11.22.11.3       Up   Up
245        aps-1      445          1/1/10      11.22.11.3       Up   Up
246        aps-1      446          1/1/10      11.22.11.3       Up   Up
247        aps-1      447          1/1/10      11.22.11.3       Up   Up
248        aps-1      448          1/1/10      11.22.11.3       Up   Up
249        aps-1      449          1/1/10      11.22.11.3       Up   Up
250        aps-1      450          1/1/10      11.22.11.3       Up   Up
251        aps-1      451          1/1/10      11.22.11.3       Up   Up
252        aps-1      452          1/1/10      11.22.11.3       Up   Up
253        aps-1      453          1/1/10      11.22.11.3       Up   Up
...
207        3/2/8     407          1/1/9       11.22.10.3       Up   Up
208        3/2/8     408          1/1/9       11.22.10.3       Up   Up
209        3/2/8     409          1/1/9       11.22.10.3       Up   Up

LSPs : 256
=====
*A:SRU4>config>router>mpls#
```

```
A:ALA-12# show router mpls static-lsp terminate
=====
MPLS Static LSPs (Terminate)
=====
In Label   In I/F     Out Label    Out I/F     Next Hop          Adm  Opr
-----
1021       1/1/1     n/a          n/a         n/a               Up   Up

LSPs : 1
=====
A:ALA-12#
```

statistics-summary

Syntax `statistics-summary`

Context `show>router>mpls>statistics-summary`

Description This command displays the number of LSP statistics configured.

Sample Output

```
*A:SRU4>config>router>mpls# show router mpls statistics-summary
=====
Statistics Summary
=====
LSP egress statistics           : 0
LSP ingress statistics          : 0
=====
*A:SRU4>config>router>mpls#
```


status

Syntax	status
Context	show>router>mpls
Description	This command displays MPLS operation information.
Output	MPLS Status Output — The following table describes MPLS status output fields.

Label	Description
Admin Status	Down — MPLS is administratively disabled. Up — MPLS is administratively enabled.
Oper Status	Down — MPLS is operationally down. Up — MPLS is operationally up.
LSP Counts	Static LSPs — Displays the count of static LSPs that originate, transit, and terminate on or through the router. Dynamic LSPs — Displays the count of dynamic LSPs that originate, transit, and terminate on or through the router. Detour LSPs — Displays the count of detour LSPs that originate, transit, and terminate on or through the router.
FR Object	Enabled — Specifies that Fast reroute object is signaled for the LSP. Disabled — Specifies that Fast reroute object is not signaled for the LSP.
Resignal Timer	Enabled — Specifies that the resignal timer is enabled for the LSP. Disabled — Specifies that the resignal timer is disabled for the LSP.
Hold Timer	Displays the amount of time that the ingress node holds before programming its data plane and declaring the LSP up to the service module.

Sample Output

```
*A:SR-1/Dut-A# /show router mpls status
```

```
=====
MPLS Status
=====
Admin Status      : Up           Oper Status      : Up
Oper Down Reason  : n/a
FR Object         : Enabled      Resignal Timer   : Disabled
Hold Timer        : 1 seconds   Next Resignal    : N/A
Srlg Frr          : Disabled    Srlg Frr Strict  : Disabled
```

```

Admin Group Frr      : Disabled
Dynamic Bypass      : Enabled
BypassResignalTimer: Disabled
Least Fill Min Thd.: 5 percent
Short. TTL Prop Lo*: Enabled
AB Sample Multipli*: 1
Exp Backoff Retry   : Disabled
Lsp Init RetryTime*: 30 seconds
Bundl*: Disabled
RetryIgpOverload    : Disabled

User Srlg Database : Disabled
BypassNextResignal : N/A
LeastFill ReoptiThd: 10 percent
Short. TTL Prop Tr*: Enabled
AB Adjust Multipli*: 288
CSPF On Loose Hop  : Disabled
MBB Pref Current H*: Enabled   Logger Event

```

```

P2mp Resignal Timer: Disabled
Sec FastRetryTimer : Disabled
P2P Max Bypass Ass*: 1000
P2PActPathFastRetry: Disabled
In Maintenance Mode: No
MplsTp              : Disabled

P2mp Next Resignal : N/A
Static LSP FR Timer: 30 seconds

P2MP S2L Fast Retry: Disabled

```

LSP Counts	Originate	Transit	Terminate
Static LSPs	0	0	0
Dynamic LSPs	501	0	0
Detour LSPs	0	0	0
P2MP S2Ls	0	0	0
MPLS-TP LSPs	0	0	0
Mesh-P2P LSPs	0	0	0
One Hop-P2P LSPs	0	0	0

```
*A:bksim3107>show>router>mpls# status
```

```
=====
MPLS Status
=====
```

```

Admin Status      : Down
Oper Down Reason  : adminDown
FR Object         : Enabled
Hold Timer        : 1 seconds
Srlg Frr          : Disabled
Admin Group Frr   : Disabled
Dynamic Bypass    : Enabled
Least Fill Min Thd.: 5 percent
Short. TTL Prop Lo*: Enabled
AB Sample Multipli*: 1
Exp Backoff Retry : Disabled
Lsp Init RetryTime*: 30 seconds
Logger Event Bundl*: Disabled

Oper Status       : Down
Resignal Timer    : Disabled
Next Resignal     : N/A
Srlg Frr Strict   : Disabled
User Srlg Database : Disabled
LeastFill ReoptiThd: 10 percent
Short. TTL Prop Tr*: Enabled
AB Adjust Multipli*: 288
CSPF On Loose Hop : Disabled

```

```

P2mp Resignal Timer: Disabled
Sec FastRetryTimer : Disabled
P2P Max Bypass Ass*: 1000
P2PActPathFastRetry: Disabled
In Maintenance Mode: No
MplsTp              : Disabled

P2mp Next Resignal : N/A
Static LSP FR Timer: 30 seconds

P2MP S2L Fast Retry: Disabled

```

LSP Counts	Originate	Transit	Terminate
Static LSPs	0	0	0
Dynamic LSPs	0	0	0
Detour LSPs	0	0	0

```
P2MP S2Ls          0          0          0
MPLS-TP LSPs      0          0          0
=====
```

*A:Dut-C# show router mpls status

=====

MPLS Status

=====

```
Admin Status      : Down          Oper Status       : Down
Oper Down Reason  : adminDown
FR Object         : Enabled       Resignal Timer    : Disabled
Hold Timer        : 1 seconds     Next Resignal     : N/A
Srlg Frr          : Disabled      Srlg Frr Strict   : Disabled
Dynamic Bypass    : Enabled       User Srlg Database : Disabled
Least Fill Min Thd.: 5 percent   LeastFill ReoptiThd: 10 percent
Short. TTL Prop Lo*: Enabled     Short. TTL Prop Tr*: Enabled
AB Sample Multipli*: 1          AB Adjust Multipli*: 288
Exp Backoff Retry : Disabled      CSPF On Loose Hop : Disabled
```

```
P2mp Resignal Timer: Disabled      P2mp Next Resignal : N/A
Sec FastRetryTimer : Disabled      Static LSP FR Timer: 30 seconds
P2P Max Bypass Ass*: 1000
P2P Active Path Fa*: 10           P2MP S2l Fast Retr*: 10
ActiveFastRetryTime: Disabled
```

LSP Counts	Originate	Transit	Terminate
Static LSPs	0	0	0
Dynamic LSPs	0	0	0
Detour LSPs	0	0	0
P2MP S2Ls	0	0	0

=====

* indicates that the corresponding row element may have been truncated.

*A:Dut-C# /show router mpls status

=====

MPLS Status

=====

```
Admin Status      : Up           Oper Status       : Up
Oper Down Reason  : n/a
FR Object         : Enabled       Resignal Timer    : Disabled
Hold Timer        : 1 seconds     Next Resignal     : N/A
Srlg Frr          : Disabled      Srlg Frr Strict   : Disabled
Dynamic Bypass    : Disabled      User Srlg Database : Disabled
Least Fill Min Thd.: 5 percent   LeastFill ReoptiThd: 10 percent
Short. TTL Prop Lo*: Enabled     Short. TTL Prop Tr*: Enabled
AB Sample Multipli*: 1          AB Adjust Multipli*: 288
Exp Backoff Retry : Disabled      CSPF On Loose Hop : Disabled
Lsp Init RetryTime*: 30 seconds
```

```
P2mp Resignal Timer: Disabled      P2mp Next Resignal : N/A
Sec FastRetryTimer : Disabled      Static LSP FR Timer: 30 seconds
P2P Max Bypass Ass*: 1000
P2PActPathFastRetry: Disabled     P2MP S2L Fast Retry: Disabled
```

LSP Counts	Originate	Transit	Terminate
Static LSPs	0	0	0
Dynamic LSPs	3	0	2

```

Detour LSPs          0          0          0
P2MP S2Ls           0          0          0
=====

```

* indicates that the corresponding row element may have been truncated.

```
show router mpls status
```

```
=====
MPLS Status
=====
```

```

Admin Status       : Down          Oper Status       : Down
Oper Down Reason   : adminDown
FR Object          : Enabled       Resignal Timer    : Disabled
Hold Timer        : 1 seconds     Next Resignal     : N/A
Srlg Frr          : Disabled      Srlg Frr Strict   : Disabled
Dynamic Bypass    : Enabled       User Srlg Database : Disabled
Least Fill Min Thd.: 5 percent   LeastFill ReoptiThd: 10 percent
Short. TTL Prop Lo*: Enabled     Short. TTL Prop Tr*: Enabled
AB Sample Multipli*: 1           AB Adjust Multipli*: 288
Exp Backoff Retry : Disabled      CSPF On Loose Hop : Disabled

```

```

P2mp Resignal Timer: Disabled      P2mp Next Resignal : N/A
Sec FastRetryTimer : Disabled      Static LSP FR Timer: 30 seconds
P2P Max Bypass Ass*: 1000
P2P Active Path Fa*: 10           P2MP S2l Fast Retr*: 10
ActiveFastRetryTime: Disabled

```

```

LSP Counts          Originate          Transit          Terminate
-----
Static LSPs         0              0              0
Dynamic LSPs        0              0              0
Detour LSPs         0              0              0
P2MP S2Ls           0              0              0
=====

```

* indicates that the corresponding row element may have been truncated.

```
*A:SRU4>config>router>mpls# show router mpls status
```

```
=====
MPLS Status
=====
```

```

Admin Status       : Up           Oper Status       : Up
Oper Down Reason   : n/a
FR Object          : Enabled       Resignal Timer    : 30 minutes
Hold Timer        : 1 seconds     Next Resignal     : 13 minutes
Srlg Frr          : Enabled      Srlg Frr Strict   : Enabled
Dynamic Bypass    : Enabled       User Srlg Database : Disabled
Least Fill Min Thd.: 5 percent   LeastFill ReoptiThd: 10 percent
Short. TTL Prop Lo*: Enabled     Short. TTL Prop Tr*: Enabled

```

```

P2mp Resignal Timer: Disabled      P2mp Next Resignal : N/A
Sec FastRetryTimer : Disabled      Static LSP FR Timer: 30 seconds

```

```

LSP Counts          Originate          Transit          Terminate
-----
Static LSPs         0              136            0
Dynamic LSPs        140            499            1626
Detour LSPs         0              0              0
P2MP S2Ls           0              0              0
=====

```

* indicates that the corresponding row element may have been truncated.

```
*A:SRU4>config>router>mpls#
```

tp-lsp

Syntax **tp-lsp** [*lsp-name*] [**status** {**up** | **down**}] [**from** *ip-address* | **to** *ip-address*] [**detail**]
tp-lsp [*lsp-name*] **path** [**protecting** | **working**] [**detail**]
tp-lsp [*lsp-name*] **protection**

Context show>router>mpls

Parameters *lsp-name* — Specifies the LSP name up to 32 characters; accepts * as a wildcard character
path — Displays LSP path information.
protection — Displays LSP protection information.
up | **down** — Specifies the state of the LSP.

Output

```
*A:mlstp-dutA# show router mpls tp-lsp
path
protection
to <a.b.c.d>
<lsp-name>
  "lsp-32" "lsp-33" "lsp-34" "lsp-35" "lsp-36" "lsp-37" "lsp-38" "lsp-39"
  "lsp-40" "lsp-41"
status {up|down}
from <ip-address>
detail

*A:mlstp-dutA# show router mpls tp-lsp "lsp-
"lsp-32" "lsp-33" "lsp-34" "lsp-35" "lsp-36" "lsp-37" "lsp-38" "lsp-39"
"lsp-40" "lsp-41"
*A:mlstp-dutA# show router mpls tp-lsp "lsp-32"

=====
MPLS MPLS-TP LSPs (Originating)
=====
LSP Name                               To                               Tun   Protect  Adm  Opr
Id                                       Id                                Path
-----
lsp-32                                0.0.3.234                        32    No        Up   Up
-----
LSPs : 1
=====
*A:mlstp-dutA# show router mpls tp-lsp "lsp-32" detail

=====
MPLS MPLS-TP LSPs (Originating) (Detail)
=====
Type : Originating
-----
LSP Name      : lsp-32
LSP Type      : MplsTp
From Node Id  : 0.0.3.233+
Adm State     : Up
LSP Up Time   : 0d 04:50:47
Transitions   : 1
DestGlobalId  : 42
LSP Tunnel ID : 32
To Node Id    : 0.0.3.234
Oper State    : Up
LSP Down Time : 0d 00:00:00
Path Changes  : 2
DestTunnelNum : 32
=====
```

*A:mlstp-dutA# show router mpls tp-lsp path

```
=====
MPLS-TP LSP Path Information
=====
LSP Name      : lsp-32                      To          : 0.0.3.234
Admin State   : Up                          Oper State   : Up

-----
Path          NextHop          InLabel    OutLabel    Out I/F      Admin  Oper
-----
Working                               32         32         AtoB_1       Up     Down
Protect                               2080      2080      AtoC_1       Up     Up
=====
LSP Name      : lsp-33                      To          : 0.0.3.234
Admin State   : Up                          Oper State   : Up

-----
Path          NextHop          InLabel    OutLabel    Out I/F      Admin  Oper
-----
Working                               33         33         AtoB_1       Up     Down
Protect                               2082      2082      AtoC_1       Up     Up
=====
LSP Name      : lsp-34                      To          : 0.0.3.234
Admin State   : Up                          Oper State   : Up

-----
Path          NextHop          InLabel    OutLabel    Out I/F      Admin  Oper
-----
Working                               34         34         AtoB_1       Up     Down
Protect                               2084      2084      AtoC_1       Up     Up
=====
LSP Name      : lsp-35                      To          : 0.0.3.234
Admin State   : Up                          Oper State   : Up

-----
Path          NextHop          InLabel    OutLabel    Out I/F      Admin  Oper
-----
Working                               35         35         AtoB_1       Up     Down
Protect                               2086      2086      AtoC_1       Up     Up
=====
LSP Name      : lsp-36                      To          : 0.0.3.234
Admin State   : Up                          Oper State   : Up

-----
Path          NextHop          InLabel    OutLabel    Out I/F      Admin  Oper
-----
Working                               36         36         AtoB_1       Up     Down
Protect                               2088      2088      AtoC_1       Up     Up
=====
LSP Name      : lsp-37                      To          : 0.0.3.234
Admin State   : Up                          Oper State   : Up

-----
Path          NextHop          InLabel    OutLabel    Out I/F      Admin  Oper
-----
Working                               37         37         AtoB_1       Up     Down
Protect                               2090      2090      AtoC_1       Up     Up
=====
LSP Name      : lsp-38                      To          : 0.0.3.234
Admin State   : Up                          Oper State   : Up
=====
```

```

-----
Path          NextHop          InLabel  OutLabel  Out I/F          Admin  Oper
-----
Working                               38       38       AtoB_1          Up    Down
Protect                               2092     2092     AtoC_1          Up    Up
=====

```

```

LSP Name      : lsp-39
Admin State   : Up
To            : 0.0.3.234
Oper State    : Up

```

```

-----
Path          NextHop          InLabel  OutLabel  Out I/F          Admin  Oper
-----
Working                               39       39       AtoB_1          Up    Down
Protect                               2094     2094     AtoC_1          Up    Up
=====

```

```

LSP Name      : lsp-40
Admin State   : Up
To            : 0.0.3.234
Oper State    : Up

```

```

-----
Path          NextHop          InLabel  OutLabel  Out I/F          Admin  Oper
-----
Working                               40       40       AtoB_1          Up    Down
Protect                               2096     2096     AtoC_1          Up    Up
=====

```

```

LSP Name      : lsp-41
Admin State   : Up
To            : 0.0.3.234
Oper State    : Up

```

```

-----
Path          NextHop          InLabel  OutLabel  Out I/F          Admin  Oper
-----
Working                               41       41       AtoB_1          Up    Down
Protect                               2098     2098     AtoC_1          Up    Up

```

```

*A:mlstp-dutA# show router mpls tp-lsp "lsp-32" path working

```

```

=====
MPLS-TP LSP Working Path Information
LSP: "lsp-32"

```

```

LSP Name      : lsp-32
Admin State   : Up
To            : 0.0.3.234
Oper State    : Up

```

```

-----
Path          NextHop          InLabel  OutLabel  Out I/F          Admin  Oper
-----
Working                               32       32       AtoB_1          Up    Down

```

```

*A:mlstp-dutA# show router mpls tp-lsp "lsp-32" path protect

```

```

=====
MPLS-TP LSP Protect Path Information
LSP: "lsp-32"

```

```

LSP Name      : lsp-32
Admin State   : Up
To            : 0.0.3.234
Oper State    : Up

```

```

-----
Path          NextHop          InLabel  OutLabel  Out I/F          Admin  Oper
-----
Protect                               2080     2080     AtoC_1          Up    Up

```

```
=====
*A:mlstp-dutA# show router mpls tp-lsp "lsp-32" path protect detail
=====
```

```
MPLS-TP LSP Protect Path Information
LSP: "lsp-32" (Detail)
=====
```

```
LSP Name      : lsp-32                      To           : 0.0.3.234
Admin State   : Up                          Oper State    : Up
```

```
Protect path information
-----
```

```
Path Type     : Protect                      LSP Num      : 2
Path Admin    : Up                          Path Oper     : Up
Out Interface : AtoC_1                       Next Hop Addr : n/a
In Label      : 2080                          Out Label     : 2080
Path Up Time  : 0d 04:52:17                   Path Dn Time  : 0d 00:00:00
Active Path   : Yes                           Active Time   : 0d 00:52:56
```

```
MEP information
MEP State     : Up                          BFD           : cc
OAM Templ     : privatebed-oam-template      CC Status     : inService
                                                    CV Status     : unknown
Protect Templ : privatebed-protection-template WTR Count Down: 0 seconds
RX PDU        : SF (1,1)                    TX PDU        : SF (1,1)
Defects       :
```

```
=====
*A:mlstp-dutA# show router mpls tp-lsp "lsp-32" path working detail
=====
```

```
MPLS-TP LSP Working Path Information
LSP: "lsp-32" (Detail)
=====
```

```
LSP Name      : lsp-32                      To           : 0.0.3.234
Admin State   : Up                          Oper State    : Up
```

```
Working path information
-----
```

```
Path Type     : Working                      LSP Num      : 1
Path Admin    : Up                          Path Oper     : Down
Down Reason   : ccFault ifDn
Out Interface : AtoB_1                       Next Hop Addr : n/a
In Label      : 32                          Out Label     : 32
Path Up Time  : 0d 00:00:00                   Path Dn Time  : 0d 00:53:01
Active Path   : No                           Active Time   : n/a
```

```
MEP information
MEP State     : Up                          BFD           : cc
OAM Templ     : privatebed-oam-template      CC Status     : outOfService
                                                    CV Status     : unknown
```

```
=====
*A:mlstp-dutA#
```

```
*A:mlstp-dutA# show router mpls tp-lsp protection
=====
```

```
MPLS-TP LSP Protection Information
Legend: W-Working, P-Protect,
=====
```

```
LSP Name      Admin Oper Path   Ingr/Egr   Act. Rx PDU
```


	State	State	State	Label	Path	Tx	PDU
lsp-32	Up	Up	W	Down 32/32	No	SF	(1,1)
			P	Up 2080/2080	Yes	SF	(1,1)
lsp-33	Up	Up	W	Down 33/33	No	SF	(1,1)
			P	Up 2082/2082	Yes	SF	(1,1)
lsp-34	Up	Up	W	Down 34/34	No	SF	(1,1)
			P	Up 2084/2084	Yes	SF	(1,1)
lsp-35	Up	Up	W	Down 35/35	No	SF	(1,1)
			P	Up 2086/2086	Yes	SF	(1,1)
lsp-36	Up	Up	W	Down 36/36	No	SF	(1,1)
			P	Up 2088/2088	Yes	SF	(1,1)
lsp-37	Up	Up	W	Down 37/37	No	SF	(1,1)
			P	Up 2090/2090	Yes	SF	(1,1)
lsp-38	Up	Up	W	Down 38/38	No	SF	(1,1)
			P	Up 2092/2092	Yes	SF	(1,1)
lsp-39	Up	Up	W	Down 39/39	No	SF	(1,1)
			P	Up 2094/2094	Yes	SF	(1,1)
lsp-40	Up	Up	W	Down 40/40	No	SF	(1,1)
			P	Up 2096/2096	Yes	SF	(1,1)
lsp-41	Up	Up	W	Down 41/41	No	SF	(1,1)
			P	Up 2098/2098	Yes	SF	(1,1)

No. of MPLS-TP LSPs: 10

=====

Show RSVP Commands

interface

Syntax `interface [ip-int-name | ip-address] statistics[detail]`

Context show>router>rsvp

Description This command shows RSVP interfaces.

ip-int-name — The name of the network IP interface. An interface name cannot be in the form of an IP address. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

ip-address — The system or network interface IP address.

statistics — Displays the RSVP interface name and counts of various RSVP packets sent and received on the interface.

detail — Displays detailed information.

Output **RSVP Interface Output** — The following table describes RSVP interface output fields.

Label	Description
Interface	The name of the IP interface.
Total Sessions	The total number of RSVP sessions on this interface. This count includes sessions that are active as well as sessions that have been signaled but a response has not yet been received.
Active Sessions	The total number of active RSVP sessions on this interface.
Total BW (Mbps)	The amount of bandwidth in megabits per second (Mbps) available to be reserved for the RSVP protocol on the interface.
Resv BW (Mbps)	The amount of bandwidth in mega-bits per seconds (Mbps) reserved on this interface. A value of zero (0) indicates that no bandwidth is reserved.
Adm	Down — The RSVP interface is administratively disabled. Up — The RSVP interface is administratively enabled.
Bfd	Yes — BFD is enabled on the RSVP interface. No — BFD is disabled on the RSVP interface.
Opr	Down — The RSVP interface is operationally down. Up — The RSVP interface is operationally up.
Port ID	Specifies the physical port bound to the interface.

Label	Description (Continued)
Active Resvs	The total number of active RSVP sessions that have reserved bandwidth.
Subscription	Specifies the percentage of the link bandwidth that RSVP can use for reservation. When the value is zero (0), no new sessions are permitted on this interface.
Port Speed	Specifies the speed for the interface.
Unreserved BW	Specifies the amount of unreserved bandwidth.
Reserved BW	Specifies the amount of bandwidth in megabits per second (Mbps) reserved by the RSVP session on this interface. A value of zero (0) indicates that no bandwidth is reserved.
Total BW	Specifies the amount of bandwidth in megabits per second (Mbps) available to be reserved for the RSVP protocol on this interface.
Aggregate	Aggregate messages are used to pack multiple RSVP messages into a single packet to reduce the network overhead. When the value is true, RSVP negotiates with each neighbor and gets consensus before sending aggregate messages.
Hello Interval	Specifies the length of time, in seconds, between the hello packets that the router sends on the interface. This value must be the same for all routers attached to a common network. When the value is zero (0), the sending of hello messages is disabled.
Refresh Time	Specifies the interval between the successive Path and Resv refresh messages. RSVP declares the session down after it misses $((\text{keep-multiplier} + 0.5) * 1.5 * \text{refresh-time})$ consecutive refresh messages.
Hello Timeouts	The total number of hello messages that timed out on this RSVP interface.
Neighbors	The IP address of the RSVP neighbor.
Sent	The total number of error free RSVP packets that have been transmitted on the RSVP interface.
Recd	The total number of error free RSVP packets received on the RSVP interface.
Total Packets	The total number of RSVP packets, including errors, received on the RSVP interface.
Bad Packets	The total number of RSVP packets with errors transmitted on the RSVP interface.
Paths	The total number of RSVP PATH messages received on the RSVP interface.
Path Errors	The total number of RSVP PATH ERROR messages transmitted on the RSVP interface.

Label	Description (Continued)
Path Tears	The total number of RSVP PATH TEAR messages received on the RSVP interface.
Resvs	The total number of RSVP RESV messages received on the RSVP interface.
Resv Confirms	The total number of RSVP RESV CONFIRM messages received on the RSVP interface.
Resv Errors	Total RSVP RESV ERROR messages received on RSVP interface.
Resv Tears	Total RSVP RESV TEAR messages received on RSVP interface.
Refresh Summaries	Total RSVP RESV summary refresh messages received on interface.
Refresh Acks	Total RSVP RESV acknowledgement messages received when refresh reduction is enabled on the RSVP interface.
Bundle Packets	Total RSVP RESV bundled packets received on the RSVP interface.
Hellos	Total RSVP RESV HELLO REQ messages received on the interface.

Sample Output

```
*A:Dut-A>config>router>mpls>lsp$ /show router rsvp interface "ip-10.10.1.1" detail

=====
RSVP Interface (Detailed) : ip-10.10.1.1
=====
-----
Interface : ip-10.10.1.1
-----
Interface          : ip-10.10.1.1
Port ID           : 1/1/1
Admin State       : Up
Oper State        : Up
Active Sessions   : 1
Active Resvs      : 0
Total Sessions    : 1
Subscription      : 100 %
Port Speed        : 100 Mbps
Total BW          : 100 Mbps
Aggregate         : Dsabl
Hello Interval    : n/a
Hello Timeouts    : n/a
Authentication    : Disabled
Auth Rx Seq Num   : n/a
Auth Key Id       : n/a
Auth Tx Seq Num   : n/a
Auth Win Size     : n/a
Refresh Reduc.    : Disabled
Reliable Deli.    : Disabled
Bfd Enabled       : n/a
Graceful Shut.    : Disabled
ImplicitNullLabel : Disabled*
GR helper         : n/a

Percent Link Bandwidth for Class Types*
Link Bw CT0       : 100
Link Bw CT1       : 0
Link Bw CT2       : 0
Link Bw CT3       : 0
Link Bw CT4       : 0
Link Bw CT5       : 0
Link Bw CT6       : 0
Link Bw CT7       : 0

Bandwidth Constraints for Class Types (Kbps)
BC0               : 100000
BC4               : 0
```

```
BC1          : 0          BC5          : 0
BC2          : 0          BC6          : 0
BC3          : 0          BC7          : 0
```

Bandwidth for TE Class Types (Kbps)

```
TE0-> Resv. Bw : 0          Unresv. Bw : 100000
TE1-> Resv. Bw : 0          Unresv. Bw : 100000
TE2-> Resv. Bw : 0          Unresv. Bw : 100000
TE3-> Resv. Bw : 0          Unresv. Bw : 100000
TE4-> Resv. Bw : 0          Unresv. Bw : 100000
TE5-> Resv. Bw : 0          Unresv. Bw : 100000
TE6-> Resv. Bw : 0          Unresv. Bw : 100000
TE7-> Resv. Bw : 0          Unresv. Bw : 100000
```

IGP Update

```
Up Thresholds(%) : 0 15 30 45 60 75 80 85 90 95 96 97 98 99 100 *
Down Thresholds(%) : 100 99 98 97 96 95 90 85 80 75 60 45 30 15 0 *
IGP Update Pending : No
Next Update       : N/A
Neighbors        : 10.20.1.2
* indicates inherited values
```

```
===== *A:Dut-A>con-
fig>router>mpls>lsp$
```

```
*A:SRU4>show>router>rsvp# interface
```

RSVP Interfaces

Interface	Total Sessions	Active Sessions	Total BW (Mbps)	Resv BW (Mbps)	Adm	Opr
system	-	-	-	-	Up	Up
aps-1	0	0	6012	0	Up	Up
aps-2	0	0	6010	0	Up	Up
aps-3	0	0	6010	0	Up	Up
sr4-1	0	0	6010	0	Up	Up
ess-7-1	9	9	100	0	Up	Up
ess-7-2	7	7	100	0	Up	Up
ess-7-3	4	4	100	0	Up	Up
ess-7-4	0	0	800	0	Up	Up
ess-7-5	0	0	800	0	Up	Up
ess-7-6	0	0	800	0	Up	Up
hubA	0	0	100	0	Up	Up
germ-1	0	0	1000	0	Up	Up
src-1.1	3	3	100	0	Up	Up
src-1.2	2	2	100	0	Up	Up
src-1.3	3	3	100	0	Up	Up
src-1.4	5	5	100	0	Up	Up
...						
g7600	0	0	1000	0	Up	Up
m160	481	481	1000	82	Up	Up

```
Interfaces : 35
```

```
*A:SRU4>show>router>rsvp#
```

```
*A:SRU4>show>router>rsvp# interface statistics
```

```
RSVP Interface (statistics)
```

```

=====
-----
Interface system
-----
Interface                : Up
Total Packets            (Sent) : 0                (Recd.): 0
Bad Packets              (Sent) : 0                (Recd.): 0
Paths                    (Sent) : 0                (Recd.): 0
Path Errors              (Sent) : 0                (Recd.): 0
Path Tears               (Sent) : 0                (Recd.): 0
Resvs                    (Sent) : 0                (Recd.): 0
Resv Confirms            (Sent) : 0                (Recd.): 0
Resv Errors              (Sent) : 0                (Recd.): 0
Resv Tears               (Sent) : 0                (Recd.): 0
Refresh Summaries        (Sent) : 0                (Recd.): 0
Refresh Acks             (Sent) : 0                (Recd.): 0
Bundle Packets           (Sent) : 0                (Recd.): 0
Hellos                   (Sent) : 0                (Recd.): 0
Auth Errors              (Sent) : 0                (Recd.): 0
-----
...
-----
Interface ml60
-----
Interface                : Up
Total Packets            (Sent) : 883643           (Recd.): 3052503
Bad Packets              (Sent) : 0                (Recd.): 0
Paths                    (Sent) : 592153           (Recd.): 373610
Path Errors              (Sent) : 464              (Recd.): 30716
Path Tears               (Sent) : 29563             (Recd.): 3480
Resvs                    (Sent) : 93970            (Recd.): 2518660
Resv Confirms            (Sent) : 0                (Recd.): 0
Resv Errors              (Sent) : 136815            (Recd.): 54115
Resv Tears               (Sent) : 13338             (Recd.): 71922
Refresh Summaries        (Sent) : 0                (Recd.): 0
Refresh Acks             (Sent) : 0                (Recd.): 0
Bundle Packets           (Sent) : 0                (Recd.): 0
Hellos                   (Sent) : 17340             (Recd.): 0
Auth Errors              (Sent) : 0                (Recd.): 0
=====
*A:SRU4>show>router>rsvp#

*A:SRU4>show>router>rsvp# interface "sr4-1" statistics
=====
RSVP Interface : sr4-1 (statistics)
=====
-----
Interface sr4-1
-----
Interface                : Up
Total Packets            (Sent) : 33100            (Recd.): 20405
Bad Packets              (Sent) : 0                (Recd.): 0
Paths                    (Sent) : 0                (Recd.): 1833
Path Errors              (Sent) : 1783             (Recd.): 9
Path Tears               (Sent) : 0                (Recd.): 1157
Resvs                    (Sent) : 76              (Recd.): 0
Resv Confirms            (Sent) : 0                (Recd.): 0
Resv Errors              (Sent) : 0                (Recd.): 0
Resv Tears               (Sent) : 1                (Recd.): 0
Refresh Summaries        (Sent) : 4                (Recd.): 33

```

```

Refresh Acks          (Sent) : 1520                (Recd.): 4
Bundle Packets       (Sent) : 0                  (Recd.): 0
Hellos               (Sent) : 29716             (Recd.): 17369
Auth Errors          (Sent) : 0                  (Recd.): 0
=====
*A:SRU4>show>router>rsvp#

*A:SRU4>show>router>rsvp#  interface detail
=====
RSVP Interfaces (Detailed)
=====
-----
Interface : system
-----
Interface          : system
Port ID            : system
Admin State        : Up                        Oper State         : Up
Active Sessions    : 0                        Active Resvs       : 0
Total Sessions     : 0
Subscription       : 100 %                    Port Speed         : 0 Mbps
Total BW           : 0 Mbps                    Aggregate          : Dsabl
Hello Interval     : 3000 ms                   Hello Timeouts     : 0
Authentication     : Disabled
Auth Rx Seq Num    : n/a                       Auth Key Id        : n/a
Auth Tx Seq Num    : n/a                       Auth Win Size      : n/a
Refresh Reduc.     : Enabled                    Reliable Deli.     : Disabled
Bfd Enabled        : No                        Graceful Shut.     : Disabled
ImplicitNullLabel  : Disabled*

Percent Link Bandwidth for Class Types*
Link Bw CT0        : 100                       Link Bw CT4        : 0
Link Bw CT1        : 0                         Link Bw CT5        : 0
Link Bw CT2        : 0                         Link Bw CT6        : 0
Link Bw CT3        : 0                         Link Bw CT7        : 0

Bandwidth Constraints for Class Types (Kbps)
BC0                 : 0                         BC4                 : 0
BC1                 : 0                         BC5                 : 0
BC2                 : 0                         BC6                 : 0
BC3                 : 0                         BC7                 : 0

Bandwidth for TE Class Types (Kbps)
TE0-> Resv. Bw      : 0                         Unresv. Bw         : 0
TE1-> Resv. Bw      : 0                         Unresv. Bw         : 0
TE2-> Resv. Bw      : 0                         Unresv. Bw         : 0
TE3-> Resv. Bw      : 0                         Unresv. Bw         : 0
TE4-> Resv. Bw      : 0                         Unresv. Bw         : 0
TE5-> Resv. Bw      : 0                         Unresv. Bw         : 0
TE6-> Resv. Bw      : 0                         Unresv. Bw         : 0
TE7-> Resv. Bw      : 0                         Unresv. Bw         : 0

IGP Update
Up Thresholds(%)   : 0 15 30 45 60 75 80 85 90 95 96 97 98 99 100 *
Down Thresholds(%) : 100 99 98 97 96 95 90 85 80 75 60 45 30 15 0 *
IGP Update Pending : No
Next Update        : N/A
No Neighbors.
-----
Interface : m160
-----
Interface          : m160

```

```

Port ID          : 3/2/1
Admin State     : Up
Active Sessions : 218
Total Sessions  : 517
Subscription    : 1000 %
Total BW        : 1000 Mbps
Hello Interval  : 3000 ms
Authentication   : Disabled
Auth Rx Seq Num : n/a
Auth Tx Seq Num : n/a
Refresh Reduc.  : Enabled
Bfd Enabled     : No
ImplicitNullLabel : Disabled*

Oper State      : Up
Active Resvs    : 0
Port Speed     : 100 Mbps
Aggregate       : Dsabl
Hello Timeouts  : 0
Auth Key Id     : n/a
Auth Win Size   : n/a
Reliable Deli.  : Disabled
Graceful Shut.  : Disabled

```

```

Percent Link Bandwidth for Class Types*
Link Bw CT0      : 100
Link Bw CT1      : 0
Link Bw CT2      : 0
Link Bw CT3      : 0
Link Bw CT4      : 0
Link Bw CT5      : 0
Link Bw CT6      : 0
Link Bw CT7      : 0

```

```

Bandwidth Constraints for Class Types (Kbps)
BC0              : 1000000
BC1              : 0
BC2              : 0
BC3              : 0
BC4              : 0
BC5              : 0
BC6              : 0
BC7              : 0

```

```

Bandwidth for TE Class Types (Kbps)
TE0-> Resv. Bw   : 0
TE1-> Resv. Bw   : 0
TE2-> Resv. Bw   : 0
TE3-> Resv. Bw   : 0
TE4-> Resv. Bw   : 0
TE5-> Resv. Bw   : 0
TE6-> Resv. Bw   : 0
TE7-> Resv. Bw   : 0
Unresv. Bw      : 1000000
Unresv. Bw      : 1000000
Unresv. Bw      : 1000000
Unresv. Bw      : 1000000
Unresv. Bw      : 1000000
Unresv. Bw      : 1000000
Unresv. Bw      : 1000000
Unresv. Bw      : 1000000

```

```

IGP Update
Up Thresholds(%) : 0 15 30 45 60 75 80 85 90 95 96 97 98 99 100 *
Down Thresholds(%) : 100 99 98 97 96 95 90 85 80 75 60 45 30 15 0 *
IGP Update Pending : No
Next Update       : N/A
Neighbors         : 10.100.30.20
* indicates inherited values

```

```

=====
*A:SRU4>show>router>rsvp#

```

neighbor

- Syntax** `neighbor [ip-address] [detail]`
- Context** `show>router>rsvp`
- Description** This command shows neighbor information.
- Parameters** *ip-address* — Displays RSVP information about the specified IP address.

detail — Displays detailed information. **Sample Output**

```
*A:Dut-A>config>router>mpls>lsp$ /show router rsvp neighbor
=====
RSVP Neighbors
=====
Legend :
  LR - Local Refresh Reduction          RR - Remote Refresh Reduction
  LD - Local Reliable Delivery          RM - Remote Node supports Message ID
  LG - Local Graceful Restart          RG - Remote Graceful Restart
=====
Neighbor      Interface      Hello  Last Oper  Flags
                        Change
=====
10.20.1.2     ip-10.10.1.1  N/A   0d 00:00:44
10.20.1.3     ip-10.10.2.1  N/A   0d 00:00:44
-----
Neighbors : 2
-----
*A:Dut-A>config>router>mpls>lsp$
```

```
*A:SR1# show router rsvp neighbor detail
=====
RSVP Neighbors (Detailed)
=====
Legend :
  LR - Local Refresh Reduction          RR - Remote Refresh Reduction
  LD - Local Reliable Delivery          RM - Remote Node supports Message ID
  LG - Local Graceful Restart          RG - Remote Graceful Restart
=====
Neighbor : 30.30.30.2
-----
Interface      : int_SR1_SR3      Hello State      : Up
Last Oper Change : 0d 00:01:02      Flags           :
Source Instance  : 0x6c8b7          Dst. Instance   : 0x530f8e0
Hello Refresh Time : 2 secs          Hello Timeout Time : 8 secs
Hello Timeout Cnt : 0              Inst. Mismatch Cnt : 0
Srefresh Time Rem. : 0 secs          Epoch Num Rx     : 0
Max Msg Id Rx    : 0              Out of order Msgs : 0
Retransmitted Msgs : 0          GR Helper        : Disabled
GR Proc Invoked Cnt : 0          GR Helper State  : None
GR Helper Time Rem : N/A          GR Nbr Restart Cap : N/A
GR Nbr Restart Time : N/A          GR Nbr Recvry Time : N/A
=====
```

```
*B:edge13# show router rsvp neighbor
=====
RSVP Neighbors
=====
Legend :
  LR - Local Refresh Reduction          RR - Remote Refresh Reduction
  LD - Local Reliable Delivery          RM - Remote Node supports Message ID
  LG - Local Graceful Restart          RG - Remote Graceful Restart
=====
Neighbor      Interface      Hello  Last Oper  Flags Change
=====
10.11.101.2   e13c2_1       Up    1d 00:52:56  LR RR LD RM
-----
LG RG
```

```

10.11.102.2      e13c2_2          Up    1d 00:52:56    LR RR LD RM
10.11.103.3      e13s1_1          Up    1d 00:52:54    LR RR LD RM

LG
10.11.104.3      e13s1_2          Up    1d 00:52:56
10.11.105.4      e13s2_1          Up    1d 00:52:56
10.11.106.4      e13s2_2          Up    1d 00:52:56

```

Neighbors : 6

session

Syntax `session session-type [from ip-address | to ip-address] lsp-name name [status {up | down}] [detail]`

Context show>router>rsvp

Description This command shows RSVP session information.

Parameters `session session-type` — Specifies the session type.

Values originate, transit, terminate, detour, detour-transit, detour-terminate, bypass-tunnel

from ip-address — Specifies the IP address of the originating router.

to ip-address — Specifies the IP address of the egress router.

lsp-name name — Specifies the name of the LSP used in the path.

status up — Specifies to display a session that is operationally up.

status down — Specifies to display a session that is operationally down.

detail — Displays detailed information.

Output **RSVP Session Output** — The following table describes RSVP session output fields.

Label	Description
From	The IP address of the originating router.
To	The IP address of the egress router.
Tunnel ID	The IP address of the tunnel's ingress node supporting this RSVP session.
LSP ID	The ID assigned by the agent to this RSVP session.
Name	The administrative name assigned to the RSVP session by the agent.
State	Down — The operational state of this RSVP session is down. Up — The operational state of this RSVP session is up.

Sample Output

```
*A:SRU4>show>router>rsvp# session
```

```

=====
RSVP Sessions
=====
From          To            Tunnel LSP   Name                               State
              ID            ID
-----
110.20.1.5    110.20.1.4    18     27648 b4-1::b4-1                        Up
110.20.1.5    110.20.1.4    1       37902 gsr::gsr                          Up
110.20.1.5    10.20.1.22    11     53760 to_10_20_1_22_cspf::to_10_2*    Up
110.20.1.4    10.20.1.20    146    17920 to_10_20_1_20_cspf_3::to_10*    Up
110.20.1.4    10.20.1.20    145    34816 to_10_20_1_20_cspf_2::to_10*    Up
110.20.1.4    10.20.1.20    147    45056 to_10_20_1_20_cspf_4::to_10*    Up
110.20.1.4    10.20.1.20    148    6656  to_10_20_1_20_cspf_5::to_10*    Up
110.20.1.4    10.20.1.20    149    58880 to_10_20_1_20_cspf_6::to_10*    Up
110.20.1.4    10.20.1.20    150    13312 to_10_20_1_20_cspf_7::to_10*    Up
110.20.1.4    10.20.1.20    152    40448 to_10_20_1_20_cspf_9::to_10*    Up
110.20.1.4    10.20.1.20    154    27648 to_10_20_1_20_cspf_11::to_1*    Up
110.20.1.4    10.20.1.20    155    12288 to_10_20_1_20_cspf_12::to_1*    Up
110.20.1.4    10.20.1.20    151    46080 to_10_20_1_20_cspf_8::to_10*    Up
110.20.1.4    10.20.1.20    153    512   to_10_20_1_20_cspf_10::to_1*    Up
110.20.1.4    10.20.1.22    164    62464 to_10_20_1_22_cspf_2::to_10*    Up
110.20.1.4    10.20.1.20    156    37888 to_10_20_1_20_cspf_13::to_1*    Up
110.20.1.4    10.20.1.20    157    24064 to_10_20_1_20_cspf_14::to_1*    Up
110.20.1.4    10.20.1.20    158    19968 to_10_20_1_20_cspf_15::to_1*    Up
110.20.1.4    10.20.1.20    161    59904 to_10_20_1_20_cspf_18::to_1*    Up
...
110.20.1.3    110.20.1.4    54     23088 to_110_20_1_4_cspf_4::to_11*    Up
-----

```

Sessions : 1976

* indicates that the corresponding row element may have been truncated.

*A:SRU4>show>router>rsvp#

A:ALA-12# show router rsvp session lsp-name A_C_2::A_C_2 status up

```

=====
RSVP Sessions
=====
From          To            Tunnel LSP   Name                               State
              ID            ID
-----
10.20.1.1     10.20.1.3     2       40   A_C_2::A_C_2                        Up
-----

```

Sessions : 1

A:ALA-12#

*A:SRU4>show>router>rsvp# session detail

RSVP Sessions (Detailed)

LSP : b4-1::b4-1

```

-----
From          : 110.20.1.5           To            : 110.20.1.4
Tunnel ID     : 18                     LSP ID       : 27648
Style        : FF                          State        : Up
Session Type  : Terminate
In Interface  : 3/2/1                Out Interface : n/a
In Label     : 131071           Out Label    : n/a
-----

```

Previous Hop : 10.100.30.20 Next Hop : n/a
SetupPriority : 7 Hold Priority : 0
Class Type : 0
SubGrpOrig ID : 0 SubGrpOrig Addr: 0.0.0.0
P2MP ID : 0

Path Recd : 7497 Path Sent : 0
Resv Recd : 0 Resv Sent : 1757

Summary messages:

SPath Recd : 0 SPath Sent : 0
SResv Recd : 0 SResv Sent : 0

LSP : gsr::gsr

From : 110.20.1.5 To : 110.20.1.4
Tunnel ID : 1 LSP ID : 37902
Style : FF State : Up
Session Type : Terminate
In Interface : 3/2/7 Out Interface : n/a
In Label : 128547 Out Label : n/a
Previous Hop : 160.60.60.2 Next Hop : n/a
SetupPriority : 7 Hold Priority : 0
Class Type : 0
SubGrpOrig ID : 0 SubGrpOrig Addr: 0.0.0.0
P2MP ID : 0

Path Recd : 5225 Path Sent : 0
Resv Recd : 0 Resv Sent : 1741

Summary messages:

SPath Recd : 0 SPath Sent : 0
SResv Recd : 0 SResv Sent : 0

...

From : 110.20.1.3 To : 110.20.1.4
Tunnel ID : 54 LSP ID : 23088
Style : SE State : Up
Session Type : Terminate
In Interface : aps-1 Out Interface : n/a
In Label : 130409 Out Label : n/a
Previous Hop : 104.104.0.3 Next Hop : n/a
SetupPriority : 7 Hold Priority : 0
Class Type : 0
SubGrpOrig ID : 0 SubGrpOrig Addr: 0.0.0.0
P2MP ID : 0

Path Recd : 1 Path Sent : 0
Resv Recd : 0 Resv Sent : 1

Summary messages:

SPath Recd : 840 SPath Sent : 0
SResv Recd : 0 SResv Sent : 850

=====
*A:SRU4>show>router

*A:Dut-B# show router rsvp session detour detail

=====
RSVP Sessions (Detailed)

```

=====
LSP : tof919::l_detour
-----
From          : 10.20.1.2          To          : 10.20.1.4
Tunnel ID     : 919                LSP ID     : 15441
Style        : SE                  State       : Up
Session Type  : Originate (Detour)
In Interface  : n/a                Out Interface : 1/1/2:1
In Label     : n/a                Out Label   : 129865
Previous Hop  : n/a                Next Hop    : 10.10.101.4
SetupPriority : 4                  Hold Priority : 4
Class Type   : 5
SugGrpOrig ID : 0                  SubGrpOrig Addr: 0.0.0.0
P2MP ID     : 0

Path Recd    : 0                    Path Sent   : 106
Resv Recd   : 113                  Resv Sent   : 0

Summary messages:
SPath Recd  : 0                    SPath Sent  : 0
SResv Recd  : 0                    SResv Sent  : 0
=====

```

```
*A:Dut-B#
```

```
*A:Dut-C# show router mpls lsp transit detail
```

```
=====
MPLS LSPs (Transit) (Detail)
-----
```

```
LSP tof1::sec2
-----
```

```

From          : 10.20.1.2          To          : 10.20.1.4
State         : Up
SetupPriority  : 5                  Hold Priority : 5
Class Type    : 5
In Interface  : lag-1:0            In Label    : 131068
Out Interface : 2/1/2              Out Label   : 131068
Previous Hop  : 10.10.12.2         Next Hop    : 10.10.11.4
Reserved BW   : 1000 Kbps
=====

```

```
*A:Dut-C#
```

```
*A:Dut-B# show router rsvp session detour-terminate detail
```

```
=====
RSVP Sessions (Detailed)
=====
```

```
LSP : tof878::l_detour
-----
```

```

From          : 10.20.1.2          To          : 10.20.1.4
Tunnel ID     : 878                LSP ID     : 14929
Style        : SE                  State       : Up
Session Type  : Terminate (Detour)
In Interface  : lag-1:0            Out Interface : 1/1/2:8
In Label     : 131069              Out Label   : 127951
Previous Hop  : 10.10.12.3         Next Hop    : 10.10.108.4
SetupPriority : 4                  Hold Priority : 4
Class Type   : 5
SugGrpOrig ID : 0                  SubGrpOrig Addr: 0.0.0.0
P2MP ID     : 0

```

```
Path Recd      : 128          Path Sent      : 0
Resv Recd     : 125          Resv Sent     : 124
```

Summary messages:

```
SPath Recd    : 0           SPath Sent    : 0
SResv Recd    : 0           SResv Sent    : 0
```

=====
*A:Dut-B#

*A:Dut-B# show router rsvp session bypass-tunnel detail

=====
RSVP Sessions (Detailed)
=====

LSP : bypass-link10.10.108.4

```
From           : 10.20.1.2          To           : 10.10.109.4
Tunnel ID      : 4003              LSP ID      : 6
Style          : FF                State        : Up
Session Type   : Bypass Tunnel
In Interface   : n/a              Out Interface : 1/1/2:9
In Label       : n/a              Out Label    : 124069
Previous Hop   : n/a              Next Hop     : 10.10.109.4
SetupPriority  : 7                 Hold Priority : 0
Class Type     : 0
SugGrpOrig ID : 0                 SubGrpOrig Addr: 0.0.0.0
P2MP ID       : 0
```

```
Path Recd      : 0          Path Sent      : 3
Resv Recd     : 4          Resv Sent     : 0
```

Summary messages:

```
SPath Recd    : 0           SPath Sent    : 0
SResv Recd    : 0           SResv Sent    : 0
```

=====
*A:Dut-B#

*A:Dut-B# show router rsvp session detour detail

=====
RSVP Sessions (Detailed)
=====

LSP : tof919::l_detour

```
From           : 10.20.1.2          To           : 10.20.1.4
Tunnel ID      : 919              LSP ID      : 15441
Style          : SE                State        : Up
Session Type   : Originate (Detour)
In Interface   : n/a              Out Interface : 1/1/2:1
In Label       : n/a              Out Label    : 129865
Previous Hop   : n/a              Next Hop     : 10.10.101.4
SetupPriority  : 4                 Hold Priority : 4
Class Type     : 5
SugGrpOrig ID : 0                 SubGrpOrig Addr: 0.0.0.0
P2MP ID       : 0
```

```
Path Recd      : 0          Path Sent      : 106
Resv Recd     : 113         Resv Sent     : 0
```

Summary messages:

```
SPath Recd    : 0           SPath Sent    : 0
SResv Recd    : 0           SResv Sent    : 0
```

```

=====
*A:Dut-B#

*A:Dut-B# show router rsvp session detour-transit detail
=====
RSVP Sessions (Detailed)
-----
LSP : tof919::l_detour
-----
From          : 10.20.1.2          To          : 10.20.1.4
Tunnel ID     : 919                LSP ID     : 15441
Style        : SE                  State      : Up
Session Type  : Transit (Detour)
In Interface  : lag-1:0            Out Interface : 1/1/2:6
In Label     : 131071             Out Label   : 127952
Previous Hop  : 10.10.12.3        Next Hop    : 10.10.106.4
SetupPriority : 4                  Hold Priority : 4
Class Type   : 5
SugGrpOrig ID : 0                  SubGrpOrig Addr: 0.0.0.0
P2MP ID      : 0

Path Recd    : 119                Path Sent   : 123
Resv Recd    : 121                Resv Sent   : 120

Summary messages:
SPath Recd   : 0                  SPath Sent  : 0
SResv Recd   : 0                  SResv Sent  : 0
=====
*A:Dut-B#

*A:Dut-B# show router rsvp session detour-terminate detail
=====
RSVP Sessions (Detailed)
-----
LSP : tof878::l_detour
-----
From          : 10.20.1.2          To          : 10.20.1.4
Tunnel ID     : 878                LSP ID     : 14929
Style        : SE                  State      : Up
Session Type  : Terminate (Detour)
In Interface  : lag-1:0            Out Interface : 1/1/2:8
In Label     : 131069             Out Label   : 127951
Previous Hop  : 10.10.12.3        Next Hop    : 10.10.108.4
SetupPriority : 4                  Hold Priority : 4
Class Type   : 5
SugGrpOrig ID : 0                  SubGrpOrig Addr: 0.0.0.0
P2MP ID      : 0

Path Recd    : 128                Path Sent   : 0
Resv Recd    : 125                Resv Sent   : 124

Summary messages:
SPath Recd   : 0                  SPath Sent  : 0
SResv Recd   : 0                  SResv Sent  : 0
=====
*A:Dut-B#

*A:Dut-B# show router rsvp session bypass-tunnel detail

```

```

=====
RSVP Sessions (Detailed)
-----
LSP : bypass-link10.10.108.4
-----
From          : 10.20.1.2          To          : 10.10.109.4
Tunnel ID     : 4003              LSP ID     : 6
Style        : FF                State       : Up
Session Type  : Bypass Tunnel
In Interface  : n/a              Out Interface : 1/1/2:9
In Label     : n/a              Out Label   : 124069
Previous Hop  : n/a              Next Hop    : 10.10.109.4
SetupPriority : 7                Hold Priority : 0
Class Type   : 0
SugGrpOrig ID : 0                SubGrpOrig Addr: 0.0.0.0
P2MP ID     : 0

Path Recd    : 0                Path Sent   : 3
Resv Recd   : 4                Resv Sent   : 0

Summary messages:
SPath Recd  : 0                SPath Sent  : 0
SResv Recd  : 0                SResv Sent  : 0
=====
*A:Dut-B#

```

statistics

- Syntax** **statistics**
- Context** show>router>rsvp
- Description** This command displays global statistics in the RSVP instance.
- Output** **RSVP Statistics Output** — The following table describes RSVP statistics output fields.

Label	Description
PATH Timeouts	The total number of path timeouts.
RESV Timeouts	The total number of RESV timeouts.

Sample Output

```

*A:SR1# /show router rsvp statistics
=====
RSVP Global Statistics
=====
PATH Timeouts      : 0                RESV Timeouts      : 0
GR Helper PATH Tim*: 0                GR Helper RESV Tim*: 0
=====
* indicates that the corresponding row element may have been truncated.

*A:SRU4>show>router>rsvp# statistics
=====

```



```

RSVP Global Statistics
=====
PATH Timeouts      : 1026                RESV Timeouts      : 182
=====
*A:SRU4>show>router>rsvp#

```

status

Syntax `rsvp status`

Context `show>router>rsvp`

Description This command displays RSVP status.

Output **RSVP Status** — The following table describes RSVP status output fields.

Label	Description
Admin Status	Down — RSVP is administratively disabled. Up — RSVP is administratively enabled.
Oper Status	Down — RSVP is operationally down. Up — RSVP is operationally up.
Keep Multiplier	Displays the keep-multiplier <i>number</i> used by RSVP to declare that a reservation is down or the neighbor is down.
Refresh Time	Displays the refresh-time interval, in seconds, between the successive Path and Resv refresh messages.
Message Pacing	Enabled — RSVP messages, specified in the max-burst command, are sent in a configured interval, specified in the period command. Disabled — Message pacing is disabled. RSVP message transmission is not regulated.
Pacing Period	Displays the time interval, in milliseconds, when the router can send the specified number of RSVP messages specified in the rsvp max-burst command.
Max Packet Burst	Displays the maximum number of RSVP messages that are sent in the specified period under normal operating conditions.
Soft Preemption Timer	Displays the time, in seconds, a node holds on to a reservation for which it has triggered the soft preemption procedure.
Rapid Retransmit	Displays the value of the rapid retransmission interval.
Rapid Retry Limit	Displays the rapid retry limit.
Graceful Shutdown	Specifies whether graceful shutdown of the RSVP node is enabled.

Sample Output

```
B:# show router rsvp status
=====
RSVP Status
=====
Admin Status      : Down                Oper Status      : Down
Keep Multiplier   : 3                    Refresh Time     : 30 sec
Message Pacing    : Disabled              Pacing Period    : 100 msec
Max Packet Burst  : 650 msgs             Refresh Bypass   : Disabled
Rapid Retransmit  : 5 hmsec             Rapid Retry Limit : 3
Graceful Shutdown : Disabled              SoftPreemptionTimer: 300 sec
Implicit Null Label: Disabled          Node-id in RRO   : Exclude
P2P Merge Point Ab*: 10                P2MP Merge Point A*: 10
DiffServTE AdmModel: Basic
Percent Link Bw CT0: 100                Percent Link Bw CT4: 0
Percent Link Bw CT1: 0                  Percent Link Bw CT5: 0
Percent Link Bw CT2: 0                  Percent Link Bw CT6: 0
Percent Link Bw CT3: 0                  Percent Link Bw CT7: 0
TE0 -> Class Type : 0                    Priority          : 0
TE1 -> Class Type : 0                    Priority          : 1
TE2 -> Class Type : 0                    Priority          : 2
TE3 -> Class Type : 0                    Priority          : 3
TE4 -> Class Type : 0                    Priority          : 4
TE5 -> Class Type : 0                    Priority          : 5
TE6 -> Class Type : 0                    Priority          : 6
TE7 -> Class Type : 0                    Priority          : 7
IgpThresholdUpdate : Disabled
Up Thresholds(%)   : 0 15 30 45 60 75 80 85 90 95 96 97 98 99 100
Down Thresholds(%) : 100 99 98 97 96 95 90 85 80 75 60 45 30 15 0
Update Timer       : N/A
Update on CAC Fail : Disabled
=====
* indicates that the corresponding row element may have been truncated.
=====
RSVP Status
=====
Admin Status      : Down                Oper Status      : Down
Keep Multiplier   : 3                    Refresh Time     : 30 sec
Message Pacing    : Disabled              Pacing Period    : 100 msec
Max Packet Burst  : 650 msgs             Refresh Bypass   : Disabled
Rapid Retransmit  : 5 hmsec             Rapid Retry Limit : 3
Graceful Shutdown : Disabled              SoftPreemptionTimer: 300 sec
GR Max Recovery   : 300 sec              GR Max Restart   : 120 sec
Implicit Null Label: Disabled          Node-id in RRO   : Exclude
P2P Merge Point Ab*: Disabled          P2MP Merge Point A*: Disabled
DiffServTE AdmModel: Basic
Percent Link Bw CT0: 100                Percent Link Bw CT4: 0
Percent Link Bw CT1: 0                  Percent Link Bw CT5: 0
Percent Link Bw CT2: 0                  Percent Link Bw CT6: 0
Percent Link Bw CT3: 0                  Percent Link Bw CT7: 0
TE0 -> Class Type : 0                    Priority          : 0
TE1 -> Class Type : 0                    Priority          : 1
TE2 -> Class Type : 0                    Priority          : 2
TE3 -> Class Type : 0                    Priority          : 3
TE4 -> Class Type : 0                    Priority          : 4
TE5 -> Class Type : 0                    Priority          : 5
TE6 -> Class Type : 0                    Priority          : 6
TE7 -> Class Type : 0                    Priority          : 7
IgpThresholdUpdate : Disabled
```

```
Up Thresholds(%) : 0 15 30 45 60 75 80 85 90 95 96 97 98 99 100
Down Thresholds(%) : 100 99 98 97 96 95 90 85 80 75 60 45 30 15 0
Update Timer      : N/A
Update on CAC Fail : Disabled
```

```
=====
* indicates that the corresponding row element may have been truncated.
```

```
show router rsvp status
```

```
=====
RSVP Status
```

```
=====
Admin Status      : Down                Oper Status       : Down
Keep Multiplier   : 3                  Refresh Time      : 30 sec
Message Pacing    : Disabled            Pacing Period    : 100 msec
Max Packet Burst  : 650 msgs            Refresh Bypass    : Disabled
Rapid Retransmit  : 5 hmsec             Rapid Retry Limit : 3
Graceful Shutdown : Disabled                SoftPreemptionTimer: 300 sec
Implicit Null Label: Disabled            Node-id in RRO    : Exclude
P2P Merge Point Ab*: 10                 P2MP Merge Point A*: 10
DiffServTE AdmModel: Basic
Percent Link Bw CT0: 100                 Percent Link Bw CT4: 0
Percent Link Bw CT1: 0                   Percent Link Bw CT5: 0
Percent Link Bw CT2: 0                   Percent Link Bw CT6: 0
Percent Link Bw CT3: 0                   Percent Link Bw CT7: 0
TE0 -> Class Type : 0                   Priority           : 0
TE1 -> Class Type : 0                   Priority           : 1
TE2 -> Class Type : 0                   Priority           : 2
TE3 -> Class Type : 0                   Priority           : 3
TE4 -> Class Type : 0                   Priority           : 4
TE5 -> Class Type : 0                   Priority           : 5
TE6 -> Class Type : 0                   Priority           : 6
TE7 -> Class Type : 0                   Priority           : 7
IgpThresholdUpdate : Disabled
Up Thresholds(%)   : 0 15 30 45 60 75 80 85 90 95 96 97 98 99 100
Down Thresholds(%) : 100 99 98 97 96 95 90 85 80 75 60 45 30 15 0
Update Timer       : N/A
Update on CAC Fail : Disabled
```

```
=====
* indicates that the corresponding row element may have been truncated.
*A:SRU4>show>router>rsvp# status
```

```
=====
RSVP Status
```

```
=====
Admin Status      : Up                  Oper Status       : Up
Keep Multiplier   : 3                  Refresh Time      : 30 sec
Message Pacing    : Disabled            Pacing Period    : 100 msec
Max Packet Burst  : 650 msgs            Refresh Bypass    : Disabled
Rapid Retransmit  : 100 hmsec           Rapid Retry Limit : 3
Graceful Shutdown : Disabled                SoftPreemptionTimer: 300 sec
Implicit Null Label: Disabled
DiffServTE AdmModel: Basic
Percent Link Bw CT0: 100                 Percent Link Bw CT4: 0
Percent Link Bw CT1: 0                   Percent Link Bw CT5: 0
Percent Link Bw CT2: 0                   Percent Link Bw CT6: 0
Percent Link Bw CT3: 0                   Percent Link Bw CT7: 0
TE0 -> Class Type : 0                   Priority           : 0
TE1 -> Class Type : 0                   Priority           : 1
TE2 -> Class Type : 0                   Priority           : 2
TE3 -> Class Type : 0                   Priority           : 3
TE4 -> Class Type : 0                   Priority           : 4
TE5 -> Class Type : 0                   Priority           : 5
```

```
TE6 -> Class Type : 0          Priority          : 6
TE7 -> Class Type : 0          Priority          : 7
IgpThresholdUpdate : Disabled
Up Thresholds(%)   : 0 15 30 45 60 75 80 85 90 95 96 97 98 99 100
Down Thresholds(%) : 100 99 98 97 96 95 90 85 80 75 60 45 30 15 0
Update Timer       : N/A
Update on CAC Fail : Disabled
=====
*A:SRU4>show>router>rsvp#
```

Tools Commands

lspinfo

Syntax	lspinfo
Context	tools>dump>router>mpls
Description	This command dumps LSP information for MPLS.

Sample Output

```
A:PC6-192.168.2.104># /tools dump router mpls lspinfo
LSP "1" LspIdx 1 LspType Dynamic State LSPS_ON_PRIMARY Flags 0x2000
NumPaths 2 NumSdps 0 NumCBFSdps 0
HoldTimeRemaining 0sec ClassType 0 Metric 0 OperMetric 65535
LDPoRsvp Include VprnAutoBind Include IgpShortCut Include BgpShortCut Include
BgpTransTunnel Include IpShCutTtlPropLocal TRUE IpShCutTtlPropTans TRUE
  Path "1" LspId 54322 PathType Primary ActivePath Yes
    Hop No: 1 IngIp 10.254.1.3 EgrIp 0.0.0.0 rtrId 0.0.0.0 HopType 1 Flag 0x0
    Hop No: 2 IngIp 10.254.1.1 EgrIp 0.0.0.0 rtrId 0.0.0.0 HopType 1 Flag 0x1
  LspPath FsmState LSP_PATHS_UP Flags 0x0
  RetryAttempts 0 RetryRemaining -1 RetryInterval 30 NextRetryIn 0secs
  Class Type 0 SetupPri 7 HoldPri 0 Pref 0 HopLimit 20 TotIgpCost 0 Metric 65535
  Oper Values:
    Class Type 0 SetupPri 7 HoldPri 0 HopLimit 20 record route no record label
    Metric 65535 [TE] include 0x25 exclude 0x80
  Last MBB -
    Type Config Change State Successful CspfFailures 0
    Started 2011/03/30 12:21:23.25 UTC
    Ended 2011/03/30 12:21:24.13 UTC
    Pre-MBB IGP Cost 0
```

ftn

Syntax	ftn
Context	tools>dump>router>mpls
Description	This command dumps FTN information for MPLS.

ilm

Syntax	ilm
Context	tools>dump>router>mpls
Description	This command dumps ILM information for MPLS.

memory-usage

Syntax	memory-usage
Context	tools>dump>router>mpls
Description	This command dumps memory usage information for MPLS.

adjust-autobandwidth

Syntax	adjust-autobandwidth [lsp <i>lsp-name</i> [force [bandwidth <i>mbps</i>]]]								
Context	tools>perform>router>mpls								
Description	<p>This command initiates an immediate auto-bandwidth adjustment attempt for either one specific LSP or all active LSPs. If an LSP is not specified then the system assumes the command applies to all LSPs. The optional force parameter, which is available only when an LSP is referenced, determines whether adjust-up and adjust-down threshold checks are applied. If force is not specified then the maximum average data rate must differ from the current reservation by more than the adjust-up or adjust-down thresholds, otherwise no bandwidth adjustment occurs. If the force option is specified then, bandwidth adjustment ignores the configured thresholds. If a bandwidth is specified as part of the force option then the bandwidth of the LSP is changed to this specific value, otherwise the bandwidth is changed to the maximum average data rate that has been measured by the system in the current adjust interval.</p> <p>The adjust-count and maximum average data rate are not reset by the manual auto-bandwidth command, whether or not the bandwidth adjustment succeeds or fails. The overflow count is reset only if the manual auto-bandwidth attempt is successful.</p>								
Default	none								
Parameters	<p><i>lsp-name</i> — The name of the LSP to which this command applies. If this parameter is not supplied the command applies to all active LSPs.</p> <table><tr><td>Values</td><td>String (32 chars max)</td></tr><tr><td>Default</td><td>none</td></tr></table> <p><i>mbps</i> — The bandwidth that the LSP should be immediately resized to.</p> <table><tr><td>Values</td><td>0—100000</td></tr><tr><td>Default</td><td>none</td></tr></table>	Values	String (32 chars max)	Default	none	Values	0—100000	Default	none
Values	String (32 chars max)								
Default	none								
Values	0—100000								
Default	none								

cspf

Syntax `cspf to ip-address`

Context `tools>perform>router>mpls`

Sample Output

```
*A:Dut-C# /tools perform router mpls cspf to 10.20.1.6
Req CSPF for all ECMP paths
  from: this node to: 10.20.1.6 w/(no Diffserv) class: 0 , setup Priority 7, Hold
Priority 0 TE Class: 7

CSPF Path
To      : 10.20.1.6
Path 1  : (cost 2000)
  Src:   10.20.1.3   (= Rtr)
  Egr:   unnumbered lnkId 4      -> Ingr:   unnumbered lnkId 2
Rtr:    10.20.1.5    (met 1000)
  Egr:   unnumbered lnkId 3      -> Ingr:   unnumbered lnkId 3
Rtr:    10.20.1.6    (met 1000)
  Dst:   10.20.1.6   (= Rtr)

Path 2  : (cost 2000)
  Src:   10.20.1.3   (= Rtr)
  Egr:   unnumbered lnkId 5      -> Ingr:   unnumbered lnkId 5
Rtr:    10.20.1.4    (met 1000)
  Egr:   unnumbered lnkId 3      -> Ingr:   unnumbered lnkId 2
Rtr:    10.20.1.6    (met 1000)
  Dst:   10.20.1.6   (= Rtr)

*A:Dut-C#
```

force-switch-path

Syntax `force-switch-path [lsp lsp-name] [path path-name]`

Context `tools>perform>router>mpls`

Description Use this command to move from a standby path to any other standby path regardless of priority.

The **no** form of the command reverts to priority path.

Parameters *lsp-name* — Specifies an existing LSP name to move.

path-name — Specifies the path name to which to move the specified LSP.

plr

Syntax plr

Context tools>dump>router>mpls>bypass-tunnel

Description

Sample Output

```
tools dump router mpls bypass-tunnel plr
=====
MPLS Bypass Tunnels
=====
Legend :  m - Manual      d - Dynamic      p - P2mp
=====
To          State  Out I/F      Out Label    Reserved   Protected   Type
                BW (Kbps)   LSP Count
-----
10.10.12.1   Up    1/1/4        124181       0          369         d

To          : 10.10.12.1      State          : Up
Out I/F     : 1/1/4            Out Label     : 124181
Up Time    : 0d 19:24:13   Active Time    : n/a
Reserved BW : 0 Kbps        Protected LSP Count : 369
Type       : Dynamic
SetupPriority : 7              Hold Priority   : 0
Class Type  : 0              Tunnel Id     : 63697
Actual Hops :
    10.10.12.2(S)      -> 10.10.12.1(S)

    Plr List: (Last PlrIdx 2)
    -----
        PLR List Index = 1
        PLR current State = PLRS_CONNECTED
        NextNodeSysId = 8.8.8.8
        AvoidNodeId = 2.2.2.2
        NodeProtect = 2 (Node Protect)
        LSP Count = 197
        PLR List Index = 2
        PLR current State = PLRS_BackupInUse
        NextNodeSysId = 8.8.8.8
        AvoidNodeId = 2.2.2.2
        NodeProtect = 2 (Node Protect)
        LSP Count = 203
```

cspf

Syntax **cspf to ip-addr [from ip-addr] [bandwidth bandwidth] [include-bitmap bitmap] [exclude-bitmap bitmap] [hop-limit limit] [exclude-address excl-addr [excl-addr...(up to 8 max)]] [use-te-metric] [strict-srlg] [srlg-group grp-id...(up to 8 max)] [exclude-node excl-node-id]**

[*excl-node-id*..(up to 8 max)] [**skip-interface** *interface-name*] [**ds-class-type** *class-type*]
 [**cspf-reqtype** *req-type*] [**least-fill-min-thd** *thd*] [**setup-priority** *val*] [**hold-priority** *val*]

Context	tools>perform>router>mpls
Description	This command computes a CSPF path with specified user constraints.
Default	none
Parameters	<p>to <i>ip-addr</i> — Specify the destination IP address.</p> <p>from <i>ip-addr</i> — Specify the originating IP address.</p> <p>bandwidth <i>bandwidth</i> — Specifies the amount of bandwidth in mega-bits per second (Mbps) to be reserved.</p> <p>include-bitmap <i>bitmap</i> — Specifies to include a bit-map that specifies a list of admin groups that should be included during setup.</p> <p>exclude-bitmap <i>bitmap</i> — Specifies to exclude a bit-map that specifies a list of admin groups that should be included during setup.</p> <p>hop-limit <i>limit</i> — Specifies the total number of hops a detour LSP can take before merging back onto the main LSP path.</p> <p>exclude-address <i>ip-addr</i> — Specifies IP addresses, up to 8, that should be included during setup.</p> <p>use-te-metric — Specifies the use of the traffic engineering metric used on the interface.</p> <p>strict-srlg — Specifies whether to associate the LSP with a bypass or signal a detour if a bypass or detour satisfies all other constraints except the SRLG constraints.</p> <p>srlg-group <i>grp-id</i> — Specifies up to 8 Shared Risk Loss Groups (SRLGs). An SRLG group represents a set of interfaces which could be subject to the same failures or defects and thus share the same risk of failing.</p> <p>Values 0 — 4294967295</p> <p>exclude-node <i>excl-node-id</i> — specifies a list of address that should be excluded when this LSP is setup.</p> <p>skip-interface <i>interface-name</i> — Specifies an interface name that should be skipped during setup.</p> <p>ds-class-type <i>class-type</i> — Specifies the class type (CT) associated with this LSP.</p> <p>Values 0 — 7</p> <p>cspf-reqtype <i>req-type</i> — Specifies the req. type.</p> <p>Values all, random, least-fill</p> <p>least-fill-min-thd <i>thd</i> — Specifies whether the use of the least-fill path selection method for the computation of the path of this LSP is enabled.</p> <p>Values 1 — 100</p> <p>setup-priority <i>val</i> — Specifies the setup priority to use when insufficient bandwidth is available to setup an LSP.</p> <p>Values 0 — 7</p>

hold-priority *val* — Specifies the hold priority value to use when insufficient bandwidth is available to setup an LSP.

Values 0 — 7

resignal

Syntax	resignal { lsp <i>lsp-name</i> path <i>path-name</i> delay <i>minutes</i> } resignal { p2mp-lsp <i>p2mp-lsp-name</i> p2mp-instance <i>p2mp-instance-name</i> p2mp-delay <i>p2mp-minutes</i> }
Context	tools>perform>router>mpls
Description	This command resignals a specific LSP path. The <i>minutes</i> parameter configures the global timer or all LSPs for resignal. If only <i>lsp-name</i> and <i>path-name</i> are provided, the LSP will be resignaled immediately.
Parameters	<i>lsp-name</i> — Specifies an existing LSP name to resignal. <i>path-name</i> — Specifies an existing path name to resignal. delay <i>minutes</i> — Configures the global timer or all LSPs to resignal. p2mp-lsp <i>p2mp-lsp-name</i> — Specifies an existing point-to-multipoint LSP name. p2mp-instance <i>p2mp-instance-name</i> — Specifies a name that identifies the P2MP LSP instance p2mp-delay <i>p2mp-minutes</i> — Specifies the delay time, in minutes.
Values	0 — 60

resignal-bypass

Syntax	resignal-bypass { lsp <i>bypass-lsp-name</i> [force] delay <i>minutes</i> }
Context	tools>perform>router>mpls
Description	<p>This command performs a manual re-optimization of a specific dynamic or manual bypass LSP, or of all dynamic bypass LSPs.</p> <p>The name of a manual bypass LSP is the one provided by the user at configuration time. The name of a dynamic bypass LSP is shown in the output of “show>router>mpls>bypass-tunnel dynamic detail”.</p> <p>The delay option triggers the global re-optimization of all dynamic bypass LSPs at the expiry of the specified delay. In essence, this option forces the global bypass resignal timer to expire after an amount of time equal to the value of the delay parameter. This option has no affect on a manual bypass LSP.</p> <p>However, when a specific bypass LSP name is specified, the named dynamic or manual bypass LSP is signaled and the associations are evaluated even if the new bypass LSP path has the same cost as the current one.</p> <p>In the specific case where the name corresponds to a manual bypass LSP, the LSP is torn down and re-signaled using the new path provided by CSPF if no PSB associations exist. If there is one or more</p>

PSB association but no PLR is active, the command is failed and the user is asked to explicitly enter the **force** option. In this case, the manual bypass LSP is torn down and re-signaled, leaving temporarily the associated LSP primary paths unprotected. Finally, if one or more PLRs associated with the manual bypass LSP is active, the command is failed.

- Parameters**
- lsp** *bypass-lsp-name* [**force**] — Specifies the name of the dynamic or manual bypass LSP. The force option is required when the name corresponds to a manual bypass LSP and the LSP has PSB associations.
- delay** *minutes* — Specifies the time, in minutes, MPLS waits before attempting to re-signal dynamic bypass LSP paths originated on the system.
- Values** 0 — 30

switch-path

switch-path [**lsp** *lsp-name*] [**path** *path-name*]

Context tools>perform>router>mpls

Use this command to move from a standby (or an active secondary) to another standby of the same priority. If a new standby path with a higher priority or a primary path comes up after the **tools perform** command is executed, the path re-evaluation command runs and the path is moved to the path specified by the outcome of the re-evaluation.

- Parameters**
- lsp-name* — Specifies an existing LSP name to move.
- path-name* — Specifies the path name to which to move the specified LSP.

te-lspinfo

Syntax **te-lspinfo** [**endpoint** *ip-address*] [**sender** *ip-address*] [**lspid** *lsp-id*] [**detail**] [**p2p** | *p2p-tid tunnel-id*]
te-lspinfo [**endpoint** *ip-address*] [**sender** *ip-address*] [**lspid** *lsp-id*] [**detail**] [**p2p** | *p2p-tid tunnel-id*]{ [**phops**] [**nhops**] [**s2l** *ip-address*] }

Context tools>dump>router>mpls

Description This command displays TE LSP information for MPLS.

Default none

Sample Output

```
B:Dut-R# tools dump router mpls te-lspinfo
Key P2P: Session(10.10.3.2, 201, 3.3.3.3) Sender(3.3.3.3, 2) PHOP(10.10.3.1), Flags
0x0

Key P2P: Session(10.10.3.1, 1035, 4.4.4.4) Sender(4.4.4.4, 22) PHOP(10.10.11.2),
Flags 0x0

Key P2MP: Session(0.0.0.0, 1, 4.4.4.4) Sender(4.4.4.4, 52226) PHOP(0.0.0.0) Flags
```

```

0x10
  S2L [1] Key: endPoint to 2.2.2.2 subGroupId - 1 subGroupOrigId - 4.4.4.4
  S2L [2] Key: endPoint to 10.10.2.2 subGroupId - 3 subGroupOrigId - 4.4.4.4
  S2L [3] Key: endPoint to 10.10.13.2 subGroupId - 4 subGroupOrigId - 4.4.4.4

Key P2MP: Session(0.0.0.0, 2, 4.4.4.4) Sender(4.4.4.4, 51714) PHOP(0.0.0.0) Flags
0x10
  S2L [1] Key: endPoint to 2.2.2.2 subGroupId - 1 subGroupOrigId - 4.4.4.4
  S2L [2] Key: endPoint to 10.10.2.2 subGroupId - 3 subGroupOrigId - 4.4.4.4
  S2L [3] Key: endPoint to 10.10.13.2 subGroupId - 4 subGroupOrigId - 4.4.4.4

Key P2MP: Session(0.0.0.0, 3, 4.4.4.4) Sender(4.4.4.4, 53250) PHOP(0.0.0.0) Flags
0x10

*A:Dut-T# tools dump router mpls te-lspinfo p2mp-tid 102 nhops

  Key P2MP: Session(0.0.0.0, 102, 4.4.4.4) Sender(4.4.4.4, 3074) PHOP(0.0.0.0) Flags
0x10
  -----
  List of NEXT HOPS
  -----

  NextHop [1] =>
  Key: Nhop - isFrr 0, outIf 0, NextHop 0.0.0.0 label - 128843 global Instance 0 is
  Leaf node
  -----
  Primary NHLFE => outLabel - 0 and NextHop - 0.0.0.0, outIf 0 (0)
    Port(NONE) NhIdx 0, ProtNhIdx 0, NumS2L 1
    ProtectInstance - 0, ProtectGroup 0
  POP
  No Backup NHLFEs for this Ltn entry
  Mid List : 3428 numS2Ls - 1 (Primary MID),

  NextHop [2] =>
  Key: Nhop - isFrr 0, outIf 3, NextHop 10.10.13.2 label - 128806 global Instance -
  48747
  -----
  Primary NHLFE => outLabel - 128806 and NextHop - 10.10.13.2, outIf 3 (126)
    Port(9/1/1) NhIdx 4322, ProtNhIdx 2275, NumS2L 1
    ProtectInstance - 1, ProtectGroup 126
  SWAP
  Backup NHLFE => outLabel - 130223 and NextHop - 10.10.3.2, outIf 5 (124)
    Port(9/2/3) outPushLabel 128806, NhIdx 5469, ProtNhIdx 0, NumS2L 1
  Mid List : 3428 numS2Ls - 1 (Primary MID),

  NextHop [3] =>
  Key: Nhop - isFrr 0, outIf 4, NextHop 10.10.2.2 label - 128836 global Instance -
  48974
  -----
  Primary NHLFE => outLabel - 128836 and NextHop - 10.10.2.2, outIf 4 (125)
    Port(lag-1) NhIdx 4292, ProtNhIdx 2245, NumS2L 2
    ProtectInstance - 1, ProtectGroup 125
  SWAP
  Backup NHLFE => outLabel - 130223 and NextHop - 10.10.3.2, outIf 5 (124)
    Port(9/2/3) outPushLabel 128836, NhIdx 5659, ProtNhIdx 0, NumS2L 2
  Mid List : 3428 numS2Ls - 1 (Primary MID), 3471 numS2Ls - 1 (Backup MID),

  S2L [1] Key: endPoint to 2.2.2.2 subGroupId - 1 subGroupOrigId - 4.4.4.4
  S2L [2] Key: endPoint to 3.3.3.3 subGroupId - 2 subGroupOrigId - 4.4.4.4
  S2L [3] Key: endPoint to 10.10.2.2 subGroupId - 3 subGroupOrigId - 4.4.4.4
  S2L [4] Key: endPoint to 10.10.13.2 subGroupId - 4 subGroupOrigId - 4.4.4.4

```

```
Total TeLspInfo Count : 1
```

psb

Syntax psb

Context tools>dump>router>rsvp

Sample Output

```
*A:Dut-A>config>router>mpls>lsp$ /tools dump router rsvp psb detail
-----
PSB:
  P2P: Session (To: 10.20.1.4 - 61441 - 10.20.1.1), Sender (10.20.1.1 - 2) PHop
  255.255.255.255

PSB CurrState: BACKUPS_CONNECTED PrevState: BACKUPS_INIT Flags: 0x0
LocalLabel 0 OutLabel 131070
Incoming IfIndex: Interface: Local API(-1)
Refresh interval 0, Send Path refresh in 3 secs, Path Refresh timeout 0 secs
PrevHop: Ctype 1 Addr 255.255.255.255, LIH 0
DnStream Nbr: Addr-> 10.20.1.3 IfIndex ip-10.10.2.1(3)
UpStream Neighbor is NULLP
Session Attribute:
  Session Name: bypass-node10.20.1.2
  HoldPri: 0 SetupPri: 7 Flags: 0x2
  Ctype: 7, IncludeGroup: 0x0 IncludeAllGroup: 0x0 ExcludeGroup: 0x0
ClassType: Absent
TSpec: Flags 0x8000 QOSC 0, PDR (infinity), PBS 0.000 bps, CDR (0.000 bps) MTU: 0
CSPF Hop List: ->
  (1) UnnumIfId 3 RtrId 10.20.1.1 EgrAdmGrp 0x0 (Strict)
  (2) UnnumIfId 2 RtrId 10.20.1.3 EgrAdmGrp 0x0 (Strict)
  (3) UnnumIfId 5 RtrId 10.20.1.4 EgrAdmGrp 0x0 (Strict)
PSB RRO : ->
  (1) * Flags : 0x0 : U
  (1) * UnInf : 10.20.1.1, 3
PSB SENT RRO : ->
  (1) * Flags : 0x0 : U
  (1) * UnInf : 10.20.1.1, 3
PSB FILTERSPEC RRO : ->
  (1) * Flags : 0x0 : U
  (1) * UnInf : 10.20.1.3, 2
  (2) * Flags : 0x1 : Global
  (2) * Label : 131070
  (3) * Flags : 0x0 : U
  (3) * UnInf : 10.20.1.4, 5
  (4) * Flags : 0x1 : Global
  (4) * Label : 131070
PSB ERO : ->
  (1) Unnumbered RouterId 10.20.1.1, LinkId 3, Strict
  (2) Unnumbered RouterId 10.20.1.3, LinkId 2, Strict
  (3) Unnumbered RouterId 10.20.1.4, LinkId 5, Strict
PSB SENT ERO : ->
  (1) Unnumbered RouterId 10.20.1.3, LinkId 2, Strict
  (2) Unnumbered RouterId 10.20.1.4, LinkId 5, Strict
SendTempl: Sender:10.20.1.1_2
```

```

AdSpec Present - Flags: 0x0
  AdSpec General
    - Service Break bit           : 0x0
    - IS Hop Count                 : 0x0
    - Path Bandwidth Estimate      : 0x0
    - Minimum Path latency         : 0x0
    - Composed path MTU            : 0

Num Paths Received   :0
Num Paths Transmitted:5
Num Resvs Received   :8
Num Resvs Transmitted:0

Num Summary Paths Received   :0
Num Summary Paths Transmitted:0
Num Summary Resvs Received   :0
Num Summary Resvs Transmitted:0
Created at 91359 (26 secs back)
-----
PSB:
  P2P: Session (To: 10.20.1.6 - 1 - 10.20.1.1), Sender (10.20.1.1 - 30208) PHop
  0.0.0.0

PSB CurrState: PRIMARYS_CONNECTED PrevState: PRIMARYS_INIT Flags: 0x8
LocalLabel 0 OutLabel 131071
Incoming IfIndex: Interface: Local API(-1)
Refresh interval 5, Send Path refresh in 4 secs, Path Refresh timeout 0 secs
PrevHop: Ctype 1 Addr 0.0.0.0, LIH 0
DnStream Nbr: Addr-> 10.20.1.2 IfIndex ip-10.10.1.1(2)
UpStream Neighbor is NULLP
Session Attribute:
  Session Name: 1::1
  HoldPri: 0 SetupPri: 7 Flags: 0x17
  Ctype: 7, IncludeGroup: 0x0 IncludeAllGroup: 0x0 ExcludeGroup: 0x0
ClassType: Absent
TSpec: Flags 0x8000 QOSC 1, PDR (infinity), PBS 0.000 bps, CDR (0.000 bps) MTU: 0
CSPF Hop List: ->
  (1) UnnumIfId 2 RtrId 10.20.1.1 EgrAdmGrp 0x0 (Strict)
  (2) UnnumIfId 2 RtrId 10.20.1.2 EgrAdmGrp 0x0 (Strict)
  (3) UnnumIfId 2 RtrId 10.20.1.4 EgrAdmGrp 0x0 (Strict)
  (4) UnnumIfId 2 RtrId 10.20.1.6 EgrAdmGrp 0x0 (Strict)
PSB RRO : ->
  (1) * Flags : 0x9 :      U LP_AVAIL NODE
  (1) * UnInf : 10.20.1.1, 2
PSB SENT RRO : ->
  (1) * Flags : 0x0 :      U
  (1) * UnInf : 10.20.1.1, 2
PSB FILTERSPEC RRO : ->
  (1) * Flags : 0x9 :      U LP_AVAIL NODE
  (1) * UnInf : 10.20.1.2, 2
  (2) * Flags : 0x1 :      Global
  (2) * Label : 131071
  (3) * Flags : 0x1 :      U LP_AVAIL
  (3) * UnInf : 10.20.1.4, 2
  (4) * Flags : 0x1 :      Global
  (4) * Label : 131071
  (5) * Flags : 0x0 :      U
  (5) * UnInf : 10.20.1.6, 2
  (6) * Flags : 0x1 :      Global
  (6) * Label : 131071

```

```

PSB ERO : ->
  (1) Unnumbered RouterId 10.20.1.2, LinkId 2, Strict
  (2) Unnumbered RouterId 10.20.1.4, LinkId 2, Strict
  (3) Unnumbered RouterId 10.20.1.6, LinkId 2, Strict
PSB SENT ERO : ->
  (1) Unnumbered RouterId 10.20.1.2, LinkId 2, Strict
  (2) Unnumbered RouterId 10.20.1.4, LinkId 2, Strict
  (3) Unnumbered RouterId 10.20.1.6, LinkId 2, Strict
SendTempl: Sender:10.20.1.1_30208
AdSpec not present
FRR: Flags 0x2 HopLimit 16 SetupPri 7 HoldPri 0 IncludeAny 0x0 ExcludeAny 0x0
IncludeAll 0x0
PLR: Flag (0x166) State PLRS_BYPASS_UP AvoidNodeId 10.20.1.2 inIntf -1 inLabel 0
PLR: FRRRequestCount: 1 CSPFFailures: 0 ProtectionType: NodeProtect

Num Paths Received    :0
Num Paths Transmitted:5
Num Resvs Received    :5
  Num Resvs Transmitted:0

Num Summmary Paths Received    :0
Num Summmary Paths Transmitted:0
Num Summmary Resvs Received    :0
Num Summmary Resvs Transmitted:0
Created at 91359 (28 secs back)
-----
Total PSB Count      : 2

```

rsb

Syntax **rsb****Context** tools>dump>router>rsvp**Sample Output**

```

4) *A:Dut-A>config>router>mpls>lsp$ /tools dump router rsvp rsb detail
-----
RSB:
  EndPt 10.20.1.4 Tid 61441 XTid 10.20.1.1 Sndr 10.20.1.1 LspId 2 ifIndex 3 NHop
  20.20.1.3
  Style FF, refresh in 0 secs
  RSVP NextHop 20.20.1.3, LIH 3 (TLV: RtrId 10.20.1.3 IntfId 2)
  CT Shared Reservation Info:
  No Reservation:
  FlowSpec :Flags 0x8000 QOSC 1, PDR (infinity), PBS 0.000 bps, CDR (0.000 bps)
             CBS 0, EBS 0, RSpecR 0, RSpecS 0 MTU 1500 MPU 20
  FwdFlowspec :Flags 0x0 QOSC 0, PDR (0.000 bps), PBS 0.000 bps, CDR (0.000 bps)
             CBS 0, EBS 0, RSpecR 0, RSpecS 0 MPU 0
  FilterSpec:
  Timeout in : 26 secs, LocLabel: 0 Sender: 10.20.1.1 lspId: 2 OutIfId: 0
  RRO :
  (1) * Flags : 0x0 :      U
  (1) * UnInf : 10.20.1.3, 2
  (2) * Flags : 0x1 :      Global
  (2) * Label : 131070
  (3) * Flags : 0x0 :      U

```

```

(3) * UnInf : 10.20.1.4, 5
(4) * Flags : 0x1 :      Global
(4) * Label : 131070
-----
RSB:
  EndPt 10.20.1.6  Tid 1  XTid 10.20.1.1  Sndr 0.0.0.0  LspId 0  ifIndex 2  NHop
20.20.1.2
Style SE, refresh in 0 secs
RSVP NextHop 20.20.1.2, LIH 2 (TLV: RtrId 10.20.1.2 IntfId 2)
CT Shared Reservation Info:
No Reservation:
FlowSpec :Flags 0x8000 QOSC 1, PDR (infinity), PBS 0.000 bps, CDR (0.000 bps)
          CBS 0, EBS 0, RSpecR 0, RSpecS 0 MTU 1496 MPU 20
FwdFlowspec :Flags 0x0 QOSC 0, PDR (0.000 bps), PBS 0.000 bps, CDR (0.000 bps)
            CBS 0, EBS 0, RSpecR 0, RSpecS 0 MPU 0
FilterSpec:
Timeout in : 21 secs, LocLabel: 0  Sender: 10.20.1.1 lspId: 30208 OutIfId: 0
RRO :
(1) * Flags : 0x9 :      U LP_AVAIL NODE
(1) * UnInf : 10.20.1.2, 2
(2) * Flags : 0x1 :      Global
(2) * Label : 131071
(3) * Flags : 0x1 :      U LP_AVAIL
(3) * UnInf : 10.20.1.4, 2
(4) * Flags : 0x1 :      Global
(4) * Label : 131071
(5) * Flags : 0x0 :      U
(5) * UnInf : 10.20.1.6, 2
(6) * Flags : 0x1 :      Global
(6) * Label : 131071
-----
Total RSB Count   : 2

```

trap-suppress

- Syntax** `trap-suppress number-of-traps time-interval`
- Context** `tools>perform>router>mpls`
- Description** This command modifies thresholds for trap suppression. The *time-interval* parameter is used to suppress traps after a certain number of traps have been raised within a period. By executing this command, there will be no more than *number-of-traps* within *time-interval*.
- Parameters** *number-of-traps* — Specifies to suppress the number of traps raised within a period.
- Values** 100 — 1000, in multiples of 100
- time-interval* — Specifies to suppress a certain number of traps raised within a period.
- Values** 1 — 300

tunnel-interface

Syntax	[no] tunnel-interface rsvp-p2mp <i>lsp-name</i> [sender <i>sender-address</i>]
Context	config>router config>router>igmp
Description	<p>This command creates a tunnel interface associated with an RSVP P2MP LSP. IPv4 multicast packets are forwarded over the P2MP LSP at the ingress LER based on a static join configuration of the multicast group against the tunnel interface associated with the originating P2MP LSP. At the egress LER, packets of a multicast group are received from the P2MP LSP via a static assignment of the specific <S,G> to the tunnel interface associated with a terminating LSP.</p> <p>At ingress LER, the tunnel interface identifier consists of a string of characters representing the LSP name for the RSVP P2MP LSP. The user can create one or more tunnel interfaces and associate each to a different RSVP P2MP LSP.</p> <p>At egress LER, the tunnel interface identifier consists of a couple of string of characters representing the LSP name for the RSVP P2MP LSP followed by the system address of the ingress LER. The LSP name must correspond to a P2MP LSP name configured by the user at the ingress LER. The LSP name string must not contain “:” (two :s) nor contain a “.” (single “.”) at the end of the LSP name. However, a “.” (single “.”) can appear anywhere in the string except at the end of the name.</p>
Default	none
Parameters	<p>rsvp-p2mp <i>lsp-name</i> — Specifies the LSP. The LSP name can be up to 32 characters long and must be unique.</p> <p>sender <i>sender-address</i> — Specifies system address of the ingress LER for the P2MP RSVP LSP.</p>

update-path

Syntax	update-path { lsp <i>lsp-name</i> path <i>current-path-name</i> new-path <i>new-path-name</i> }
Context	tools>perform>router>mpls
Description	<p>This command enables you to instruct MPLS to replace the path of a primary or secondary LSP. The primary or secondary LSP path is indirectly identified via the <i>current-path-name</i> value. The same path name cannot be used more than once in a given LSP name.</p> <p>This command applies to both CSPF LSP and to a non-CSPF LSP. This command will only work when the specified <i>current-path-name</i> has the adaptive option enabled. The adaptive option can be enabled at the LSP level or the path level.</p> <p>The new path must have been configured in the CLI or provided via SNMP. The CLI command for entering the path is</p> <pre>configure router mpls path <i>path-name</i></pre> <p>The command fails if any of the following conditions exist:</p> <ul style="list-style-type: none"> • The specified <i>current-path-name</i> of this LSP does not have the adaptive option enabled. • The specified <i>new-path-name</i> value does not correspond to a previously defined path.

- The specified *new-path-name* value exists but is being used by any path of the same LSP, including this one.

When you execute this command, MPLS performs the following procedures:

- MPLS performs a single MBB attempt to move the LSP path to the new path.
- If the MBB is successful, MPLS updates the new path
 - MPLS writes the corresponding NHLFE in the data path if this path is the current backup path for the primary.
 - If the current path is the active LSP path, it will update the path, write the new NHLFE in the data path that will cause traffic to switch to the new path.
- If the MBB is not successful, the path retains its current value.
- The update-path MBB has the same priority as the manual re-signal MBB.

Clear Commands

interface

- Syntax** `interface ip-int-name [statistics]`
- Context** `clear>router>mpls`
- Description** This command resets or clears statistics for MPLS interfaces.
- Parameters** *ip-int-name* — The name of an existing IP interface. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.
- statistics** — This parameter clears only statistics.

lsp

- Syntax** `lsp lsp-name`
- Context** `clear>router>mpls`
- Description** This command resets and restarts an LSP.
- Parameters** *lsp-name* — The name of the LSP to clear up to 64 characters in length.

lsp-autobandwidth

- Syntax** `lsp-autobandwidth [lsp-name]`
- Context** `clear>router>mpls>lsp`
- Description** This command clears the following counters/timers, as follows:
- The sample count is reset to zero, and the average data rate of the current sample interval is discarded.
 - The adjust count is reset to zero.
 - The maximum average data rate is zeroed.
 - The overflow count is zeroed.

ingress-stats

Syntax	ingress-statistics
Context	clear>router>mpls
Description	This command provides the context for the user to enter the LSP names for the purpose of enabling ingress data path statistics at the terminating node of the LSP (for example, egress LER).
Default	none

lsp-egress-stats

Syntax	lsp-egress-stats lsp-egress-stats <i>lsp-name</i>
Context	clear>router>mpls
Description	This command clears MPLS LSP egress statistics information.

lsp-ingress-stats

Syntax	lsp-ingress-stats lsp-ingress-stats <i>ip-address</i> lsp <i>lsp-name</i> lsp-ingress-stats <i>sender-address:lsp-name</i>
Context	clear>router>mpls
Description	This command clears MPLS LSP ingress statistics information.

interface

Syntax	interface <i>ip-int-name</i> statistics
Context	clear>router>rsvp
Description	This command resets or clears statistics for an RSVP interface.
Parameters	<i>ip-int-name</i> — The name of the IP interface to clear. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes. statistics — This parameter clears only statistics.

statistics

Syntax	statistics
Context	clear>router>rsvp
Description	This command clears global statistics for the RSVP instance, for example, clears path and resv time-out counters.

Debug Commands

mpls

Syntax	mpls [lsp <i>lsp-name</i>] [sender <i>source-address</i>] [endpoint <i>endpoint-address</i>] [tunnel-id <i>tunnel-id</i>] [lsp-id <i>lsp-id</i>] [interface <i>ip-int-name</i>] no mpls
Context	debug>router
Description	This command enables and configures debugging for MPLS.
Parameters	lsp <i>lsp-name</i> — Name that identifies the LSP. The LSP name can be up to 32 characters long and must be unique. sender <i>source-address</i> — The system IP address of the sender. endpoint <i>endpoint-address</i> — The far-end system IP address. tunnel-id <i>tunnel-id</i> — The MPLS SDP ID. Values 0 — 4294967295 lsp-id <i>lsp-id</i> — The LSP ID. Values 1 — 65535 interface <i>ip-int-name</i> — Name that identifies the interface. The interface name can be up to 32 characters long and must be unique. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

event

Syntax	[no] event
Context	debug>router>mpls debug>router>rsvp
Description	This command enables debugging for specific events. The no form of the command disables the debugging.

all

Syntax	all [detail] no all
Context	debug>router>mpls>event debug>router>rsvp>event
Description	This command debugs all events. The no form of the command disables the debugging.
Parameters	detail — Displays detailed information about all events.

frr

Syntax	frr [detail] no frr
Context	debug>router>mpls>event
Description	This command debugs fast re-route events. The no form of the command disables the debugging.
Parameters	detail — Displays detailed information about re-route events.

iom

Syntax	iom [detail] no iom
Context	debug>router>mpls>event
Description	This command reports MPLS debug events originating from the XMA. The no form of the command disables the debugging.
Parameters	detail — Displays detailed information about MPLS events originating from the XMA.

lsp-setup

Syntax	lsp-setup [detail] no lsp-setup
Context	debug>router>mpls>event
Description	This command debugs LSP setup events. The no form of the command disables the debugging.
Parameters	detail — Displays detailed information about LSP setup events.

mbb

Syntax	mbb [detail] no mbb
Context	debug>router>mpls>event
Description	This command debugs the state of the most recent invocation of the make-before-break (MBB) functionality. The no form of the command disables the debugging.
Parameters	detail — Displays detailed information about MBB events.

misc

Syntax	misc [detail] no misc
Context	debug>router>mpls>event debug>router>rsvp>event
Description	This command debugs miscellaneous events. The no form of the command disables the debugging.
Parameters	detail — Displays detailed information about miscellaneous events.

xc

Syntax	xc [detail] no xc
Context	debug>router>mpls>event
Description	This command debugs cross connect events. The no form of the command disables the debugging.
Parameters	detail — Displays detailed information about cross connect events.

rsvp

Syntax	[lsp lsp-name] [sender source-address] [endpoint endpoint-address] [tunnel-id tunnel-id] [lsp-id lsp-id] [interface ip-int-name] no rsvp
Context	debug>router
Description	This command enables and configures debugging for RSVP.

Parameters **lsp** *lsp-name* — Name that identifies the LSP. The LSP name can be up to 32 characters long and must be unique.

sender *source-address* — The system IP address of the sender.

endpoint *endpoint-address* — The far-end system IP address.

tunnel-id *tunnel-id* — The RSVP tunnel ID.

Values 0 — 4294967295

lsp-id *lsp-id* — The LSP ID.

Values 1 — 65535

interface *ip-int-name* — The interface name. The interface name can be up to 32 characters long and must be unique. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

nbr

Syntax **nbr** [**detail**]
no nbr

Context debug>router>rsvp>event

Description This command debugs neighbor events.
The **no** form of the command disables the debugging.

Parameters **detail** — Displays detailed information about neighbor events.

path

Syntax **path** [**detail**]
no path

Context debug>router>rsvp>event

Description This command debugs path-related events.
The **no** form of the command disables the debugging.

Parameters **detail** — Displays detailed information about path-related events.

resv

Syntax **resv** [**detail**]
no resv

Context debug>router>rsvp>event

Description This command debugs RSVP reservation events.

The **no** form of the command disables the debugging.

Parameters **detail** — Displays detailed information about RSVP reservation events.

te-threshold-update

Syntax **te-threshold-update**
no te-threshold-update

Context debug>router>rsvp>event
debug>router>rsvp>ip-int-name>event

Description This command debugs the te-threshold-update events.
The **no** form of this command disables the debugging

packet

Syntax [**no**] **packet**

Context debug>router>rsvp>

Description This command enters the syntax to debug packets.

all

Syntax **all [detail]**
no all

Context debug>router>rsvp>packet

Description This command debugs all packets.
The **no** form of the command disables the debugging.

Parameters **detail** — Displays detailed information about all RSVP packets.

hello

Syntax **hello [detail]**
no hello

Context debug>router>rsvp>packet

Description This command debugs hello packets.
The **no** form of the command disables the debugging.

Parameters **detail** — Displays detailed information about hello packets.

path

Syntax	path [detail] no path
Context	debug>router>rsvp>packet
Description	This command enables debugging for RSVP path packets. The no form of the command disables the debugging.
Parameters	detail — Displays detailed information about path-related events.

patherr

Syntax	patherr [detail] no patherr
Context	debug>router>rsvp>packet
Description	This command debugs path error packets. The no form of the command disables the debugging.
Parameters	detail — Displays detailed information about path error packets.

pathtear

Syntax	pathtear [detail] no pathtear
Context	debug>router>rsvp>packet
Description	This command debugs path tear packets. The no form of the command disables the debugging.
Parameters	detail — Displays detailed information about path tear packets.

resv

Syntax	resv [detail] no resv
Context	debug>router>rsvp>packet
Description	This command enables debugging for RSVP resv packets. The no form of the command disables the debugging.
Parameters	detail — Displays detailed information about RSVP Resv events.

resvterr

Syntax	resvterr [detail] no resvterr
Context	debug>router>rsvp>packet
Description	This command debugs ResvErr packets. The no form of the command disables the debugging.
Parameters	detail — Displays detailed information about ResvErr packets.

resvttear

Syntax	resvttear [detail] no resvttear
Context	debug>router>rsvp>packet
Description	This command debugs ResvTear packets. The no form of the command disables the debugging.
Parameters	detail — Displays detailed information about ResvTear packets.

Label Distribution Protocol

In This Chapter

This chapter provides information to enable Label Distribution Protocol (LDP):

- [Label Distribution Protocol on page 430](#)
 - [LDP and MPLS on page 430](#)
 - [LDP Architecture on page 431](#)
 - [Subsystem Interrelationships on page 432](#)
 - [Execution Flow on page 434](#)
 - [Label Exchange on page 436](#)
 - [Global LDP Filters on page 438](#)
 - [Global LDP Filters on page 438](#)
 - [Configuring Multiple LDP LSR ID on page 439](#)
 - [TTL Security for BGP and LDP on page 443](#)
 - [ECMP Support for LDP on page 444](#)
 - [Operation of LDP over an Unnumbered IP Interface on page 446](#)
 - [LDP over RSVP Tunnels on page 449](#)
 - [LDP over RSVP Without Area Boundary on page 454](#)
 - [LDP ECMP Uniform Failover on page 456](#)
 - [LDP Fast-Reroute for IS-IS and OSPF Prefixes on page 458](#)
 - [LDP FEC to BGP Label Route Stitching on page 469](#)
 - [Automatic Creation of a Targeted Hello Adjacency and LDP Session on page 473](#)
 - [Multicast P2MP LDP for GRT on page 478](#)
 - [LDP P2MP Support on page 479](#)
 - [Multicast LDP Fast Upstream Switchover on page 480](#)
 - [Multi-Area and Multi-Instance Extensions to LDP on page 485](#)
- [LDP Process Overview on page 503](#)

Label Distribution Protocol

Label Distribution Protocol (LDP) is a protocol used to distribute labels in non-traffic-engineered applications. LDP allows routers to establish label switched paths (LSPs) through a network by mapping network-layer routing information directly to data link layer-switched paths.

An LSP is defined by the set of labels from the ingress Label Switching Router (LSR) to the egress LSR. LDP associates a Forwarding Equivalence Class (FEC) with each LSP it creates. A FEC is a collection of common actions associated with a class of packets. When an LSR assigns a label to a FEC, it must let other LSRs in the path know about the label. LDP helps to establish the LSP by providing a set of procedures that LSRs can use to distribute labels.

The FEC associated with an LSP specifies which packets are mapped to that LSP. LSPs are extended through a network as each LSR splices incoming labels for a FEC to the outgoing label assigned to the next hop for the given FEC. The next-hop for a FEC prefix is resolved in the routing table. LDP can only resolve FECs for IGP and static prefixes. LDP does not support resolving FECs of a BGP prefix.

LDP allows an LSR to request a label from a downstream LSR so it can bind the label to a specific FEC. The downstream LSR responds to the request from the upstream LSR by sending the requested label.

LSRs can distribute a FEC label binding in response to an explicit request from another LSR. This is known as Downstream On Demand (DOD) label distribution. LSRs can also distribute label bindings to LSRs that have not explicitly requested them. This is called Downstream Unsolicited (DU).

LDP and MPLS

LDP performs the label distribution only in MPLS environments. The LDP operation begins with a hello discovery process to find LDP peers in the network. LDP peers are two LSRs that use LDP to exchange label/FEC mapping information. An LDP session is created between LDP peers. A single LDP session allows each peer to learn the other's label mappings (LDP is bi-directional) and to exchange label binding information.

LDP signaling works with the MPLS label manager to manage the relationships between labels and the corresponding FEC. For service-based FECs, LDP works in tandem with the Service Manager to identify the virtual leased lines (VLLs) and Virtual Private LAN Services (VPLSs) to signal.

An MPLS label identifies a set of actions that the forwarding plane performs on an incoming packet before discarding it. The FEC is identified through the signaling protocol (in this case, LDP) and allocated a label. The mapping between the label and the FEC is communicated to the

forwarding plane. In order for this processing on the packet to occur at high speeds, optimized tables are maintained in the forwarding plane that enable fast access and packet identification.

When an unlabeled packet ingresses the router, classification policies associate it with a FEC. The appropriate label is imposed on the packet, and the packet is forwarded. Other actions that can take place before a packet is forwarded are imposing additional labels, other encapsulations, learning actions, etc. When all actions associated with the packet are completed, the packet is forwarded.

When a labeled packet ingresses the router, the label or stack of labels indicates the set of actions associated with the FEC for that label or label stack. The actions are preformed on the packet and then the packet is forwarded.

The LDP implementation provides DOD, DU, ordered control, liberal label retention mode support.

LDP Architecture

LDP comprises a few processes that handle the protocol PDU transmission, timer-related issues, and protocol state machine. The number of processes is kept to a minimum to simplify the architecture and to allow for scalability. Scheduling within each process prevents starvation of any particular LDP session, while buffering alleviates TCP-related congestion issues.

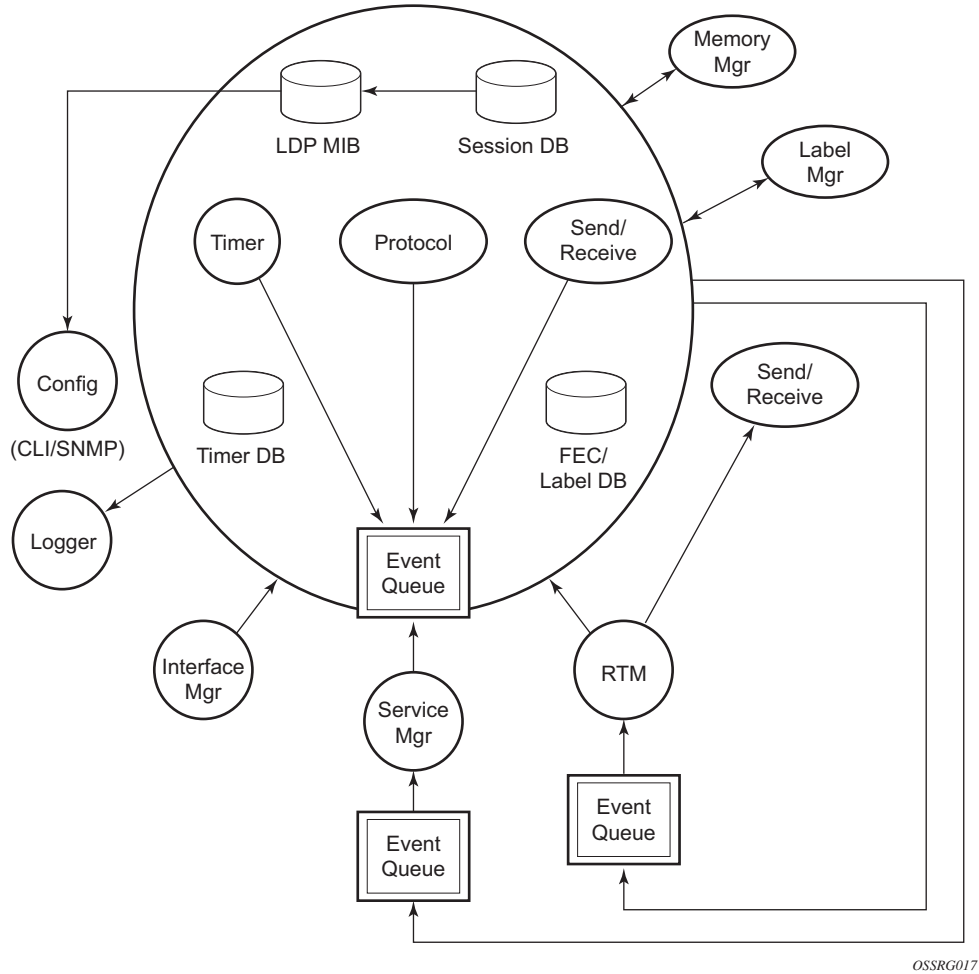
The LDP subsystems and their relationships to other subsystems are illustrated in [Figure 35](#). This illustration shows the interaction of the LDP subsystem with other subsystems, including memory management, label management, service management, SNMP, interface management, and RTM. In addition, debugging capabilities are provided through the logger.

Communication within LDP tasks is typically done by inter-process communication through the event queue, as well as through updates to the various data structures. The primary data structures that LDP maintains are:

- FEC/label database — This database contains all the FEC to label mappings that include, both sent and received. It also contains both address FECs (prefixes and host addresses) as well as service FECs (L2 VLLs and VPLS).
- Timer database — This database contains all the timers for maintaining sessions and adjacencies.
- Session database — This database contains all the session and adjacency records, and serves as a repository for the LDP MIB objects.

Subsystem Interrelationships

The sections below describe how LDP and the other subsystems work to provide services.



OSSRG017

Figure 35: Subsystem Interrelationships

Memory Manager and LDP

LDP does not use any memory until it is instantiated. It pre-allocates some amount of fixed memory so that initial startup actions can be performed. Memory allocation for LDP comes out of a pool reserved for LDP that can grow dynamically as needed. Fragmentation is minimized by allocating memory in larger chunks and managing the memory internally to LDP. When LDP is shut down, it releases all memory allocated to it.

Label Manager

LDP assumes that the label manager is up and running. LDP will abort initialization if the label manager is not running. The label manager is initialized at system boot-up; hence, anything that causes it to fail will likely imply that the system is not functional. The router uses a dynamic label range from values 18,432 through 262,143 (131,071 in chassis modes lower than D) to allocate all dynamic labels, including RSVP and BGP allocated labels and VC labels.

LDP Configuration

The router uses a single consistent interface to configure all protocols and services. CLI commands are translated to SNMP requests and are handled through an agent-LDP interface. LDP can be instantiated or deleted through SNMP. Also, LDP targeted sessions can be set up to specific endpoints. Targeted-session parameters are configurable.

Logger

LDP uses the logger interface to generate debug information relating to session setup and teardown, LDP events, label exchanges, and packet dumps. Per-session tracing can be performed.

Service Manager

All interaction occurs between LDP and the service manager, since LDP is used primarily to exchange labels for Layer 2 services. In this context, the service manager informs LDP when an LDP session is to be set up or torn down, and when labels are to be exchanged or withdrawn. In turn, LDP informs service manager of relevant LDP events, such as connection setups and failures, timeouts, labels signaled/withdrawn.

Execution Flow

LDP activity in the operating system is limited to service-related signaling. Therefore, the configurable parameters are restricted to system-wide parameters, such as hello and keepalive timeouts.

Initialization

LDP makes sure that the various prerequisites, such as ensuring the system IP interface is operational, the label manager is operational, and there is memory available, are met. It then allocates itself a pool of memory and initializes its databases.

Session Lifetime

In order for a targeted LDP (T-LDP) session to be established, an adjacency must be created. The LDP extended discovery mechanism requires hello messages to be exchanged between two peers for session establishment. After the adjacency establishment, session setup is attempted.

Adjacency Establishment

In the router, the adjacency management is done through the establishment of a Service Distribution Path (SDP) object, which is a service entity in the TiMetra Systems service model.

The TiMetra Systems service model uses logical entities that interact to provide a service. The service model requires the service provider to create configurations for four main entities:

- Customers
- Services
- Service Access Paths (SAPs) on the local routers
- Service Distribution Points (SDPs) that connect to one or more remote routers.

An SDP is the network-side termination point for a tunnel to a remote router. An SDP defines a local entity that includes the system IP address of the remote routers and a path type. Each SDP comprises:

- The SDP ID
- The transport encapsulation type, either MPLS or GRE
- The far-end system IP address

If the SDP is identified as using LDP signaling, then an LDP extended hello adjacency is attempted.

If another SDP is created to the same remote destination, and if LDP signaling is enabled, no further action is taken, since only one adjacency and one LDP session exists between the pair of nodes.

An SDP is a uni-directional object, so a pair of SDPs pointing at each other must be configured in order for an LDP adjacency to be established. Once an adjacency is established, it is maintained through periodic hello messages.

Session Establishment

When the LDP adjacency is established, the session setup follows as per the LDP specification. Initialization and keepalive messages complete the session setup, followed by address messages to exchange all interface IP addresses. Periodic keepalives or other session messages maintain the session liveliness.

Since TCP is back-pressured by the receiver, it is necessary to be able to push that back-pressure all the way into the protocol. Packets that cannot be sent are buffered on the session object and re-attempted as the back-pressure eases.

Label Exchange

Label exchange is initiated by the service manager. When an SDP is attached to a service (for example, the service gets a transport tunnel), a message is sent from the service manager to LDP. This causes a label mapping message to be sent. Additionally, when the SDP binding is removed from the service, the VC label is withdrawn. The peer must send a label release to confirm that the label is not in use.

Other Reasons for Label Actions

Other reasons for label actions include:

- **MTU changes:** LDP withdraws the previously assigned label, and re-signals the FEC with the new MTU in the interface parameter.
 - **Clear labels:** When a service manager command is issued to clear the labels, the labels are withdrawn, and new label mappings are issued.
 - **SDP down:** When an SDP goes administratively down, the VC label associated with that SDP for each service is withdrawn.
 - **Memory allocation failure:** If there is no memory to store a received label, it is released.
 - **VC type unsupported:** When an unsupported VC type is received, the received label is released.
-

Cleanup

LDP closes all sockets, frees all memory, and shuts down all its tasks when it is deleted, so its memory usage is 0 when it is not running.

Configuring Implicit Null Label

The implicit null label option allows an egress LER to receive MPLS packets from the previous hop without the outer LSP label. The user can configure to signal the implicit operation of the previous hop is referred to as penultimate hop popping (PHP). This option is signaled by the egress LER to the previous hop during the FEC signaling by the LDP control protocol.

Enable the use of the implicit null option, for all LDP FECs for which this node is the egress LER, using the following command:

```
config>router>ldp>implicit-null-label
```

When the user changes the implicit null configuration option, LDP withdraws all the FECs and re-advertises them using the new label value.

Global LDP Filters

Both inbound and outbound LDP label binding filtering are supported.

Inbound filtering is performed by way of the configuration of an import policy to control the label bindings an LSR accepts from its peers. Label bindings can be filtered based on:

- Prefix-list: Match on bindings with the specified prefix/prefixes.

The default import policy is to accept all FECs received from peers.

Outbound filtering is performed by way of the configuration of an export policy. The Global LDP export policy can be used to explicitly originate label bindings for local interfaces. The Global LDP export policy does not filter out or stop propagation of any FEC received from neighbors. Use the LDP peer export prefix policy for this purpose. It must also be noted that the system IP address AND static FECs cannot be blocked using an export policy.

Export policy enables configuration of a policy to advertise label bindings based on:

- Direct: All local subnets.
- Prefix-list: Match on bindings with the specified prefix or prefixes.

The default export policy is to originate label bindings for system address only and to propagate all FECs received from other LDP peers.

Finally, it must be noted that the 'neighbor' statement inside a global import or export policy is not considered by LDP. Use the LDP peer import or export prefix policy for this purpose.

Per LDP Peer FEC Import and Export Policies

The FEC prefix export policy provides a way to control which FEC prefixes received from prefixes received from other LDP and T-LDP peers are re-distributed to this LDP peer.

The user configures the FEC prefix export policy using the following command:

```
config>router>ldp>peer-parameters>peer>export-prefixes policy-name
```

By default, all FEC prefixes are exported to this peer.

The FEC prefix import policy provides a mean of controlling which FEC prefixes received from this LDP peer are imported and installed by LDP on this node. If resolved these FEC prefixes are then re-distributed to other LDP and T-LDP peers.

The user configures the FEC prefix export policy using the following command:

```
config>router>ldp>peer-parameters>peer>import-prefixes policy-name
```

By default, all FEC prefixes are imported from this peer.

Configuring Multiple LDP LSR ID

The multiple LDP LSR-ID feature provides the ability to configure and initiate multiple Targeted LDP (T-LDP) sessions on the same system using different LDP LSR-IDs. In the current implementation, all T-LDP sessions must have the LSR-ID match the system interface address. This feature continues to allow the use of the system interface by default, but also any other network interface, including a loopback, address on a per T-LDP session basis. Note that LDP control plane will not allow more than a single T-LDP session with different local LSR ID values to the same LSR-ID in a remote node.

An SDP of type LDP can use a provisioned targeted session with the local LSR-ID set to any network IP for the T-LDP session to the peer matching the SDP far-end address. If, however, no targeted session has been explicitly pre-provisioned to the far-end node under LDP, then the SDP will auto-establish one but will use the system interface address as the local LSR-ID.

An SDP of type RSVP must use an RSVP LSP with the destination address matching the remote node LDP LSR-ID. An SDP of type GRE can only use a T-LDP session with a local LSR-ID set to the system interface.

The multiple LDP LSR-ID feature also provides the ability to use the address of the local LDP interface, or any other network IP interface configured on the system, as the LSR-ID to establish link LDP Hello adjacency and LDP session with directly connected LDP peers. The network interface can be a loopback or not.

Link LDP sessions to all peers discovered over a given LDP interface share the same local LSR-ID. However, LDP sessions on different LDP interfaces can use different network interface addresses as their local LSR-ID.

By default, the link and targeted LDP sessions to a peer use the system interface address as the LSR-ID unless explicitly configured using this feature. Note, however, that the system interface must always be configured on the router or the LDP protocol will not come up on the node. There is no requirement to include it in any routing protocol.

Note that when an interface other than system is used as the LSR-ID, the transport connection (TCP) for the link or targeted LDP session will also use the address of that interface as the transport address.

T-LDP hello reduction

This feature implements a new mechanism to suppress the transmission of the Hello messages following the establishment of a Targeted LDP session between two LDP peers. The Hello adjacency of the targeted session does not require periodic transmission of Hello messages as in the case of a link LDP session. In link LDP, one or more peers can be discovered over a given network IP interface and as such, the periodic transmission of Hello messages is required to discover new peers in addition to the periodic Keep-Alive message transmission to maintain the existing LDP sessions. A Targeted LDP session is established to a single peer. Thus, once the Hello Adjacency is established and the LDP session is brought up over a TCP connection, Keep-Alive messages are sufficient to maintain the LDP session.

When this feature is enabled, the targeted Hello adjacency is brought up by advertising the Hold-Time value the user configured in the Hello timeout parameter for the targeted session. The LSR node will then start advertising an exponentially increasing Hold-Time value in the Hello message as soon as the targeted LDP session to the peer is up. Each new incremented Hold-Time value is sent in a number of Hello messages equal to the value of the Hello reduction factor before the next exponential value is advertised. This provides time for the two peers to settle on the new value. When the Hold-Time reaches the maximum value of 0xffff (binary 65535), the two peers will send Hello messages at a frequency of every $[(65535-1)/\text{local helloFactor}]$ seconds for the lifetime of the targeted-LDP session (for example, if the local Hello Factor is three (3), then Hello messages will be sent every 21844 seconds).

Both LDP peers must be configured with this feature to bring gradually their advertised Hold-Time up to the maximum value. If one of the LDP peers does not, the frequency of the Hello messages of the targeted Hello adjacency will continue to be governed by the smaller of the two Hold-Time values. This feature complies to *draft-pdutta-mpls-tldp-hello-reduce*.

Tracking a T-LDP Peer with BFD

BFD tracking of an LDP session associated with a T-LDP adjacency allows for faster detection of the liveliness of the session by registering the transport address of a LDP session with a BFD session.

By enabling BFD for a selected targeted session, the state of that session is tied to the state of the underneath BFD session between the two nodes. The parameters used for the BFD are set with the BFD command under the IP interface.

Tracking a Link LDP Peer with BFD

Tracking of the Hello adjacency to an LDP peer using BFD is supported.

Hello adjacency tracking with BFD is enabled by enabling BFD on an LDP interface:

- `config>router>ldp>interface-parameters>interface>enable-bfd`

The parameters used for the BFD session, for example, `transmit-interval`, `receive-interval`, and `multiplier`, are those configured under the IP interface in the existing `config>router>interface>bfd` context.

When this command is enabled on an LDP interface, LDP registers with BFD and starts tracking the LSR-id of all peers it forms Hello adjacencies with over that LDP interface.

The parameters used for the BFD session, for example, `transmit-interval`, `receive-interval`, and `multiplier`, are those configured under the IP interface in the existing `config>router>interface>bfd` context.

When enabled, the LDP hello mechanism is used to determine the remote address to be used for the BFD session. If a BFD session fails, then the associated LDP adjacency is also declared down and LDP will immediately begin its reconvergence.

LDP LSP Statistics

RSVP-TE LSP statistics is extended to LDP to provide the following counters:

- Per-forwarding-class forwarded in-profile packet count
- Per-forwarding-class forwarded in-profile byte count
- Per-forwarding-class forwarded out-of-profile packet count
- Per-forwarding-class forwarded out-of-profile byte count

The counters are available for the egress data path of an LDP FEC at ingress LER and at LSR. Because an ingress LER is also potentially an LSR for an LDP FEC, combined egress data path statistics will be provided whenever applicable.

TTL Security for BGP and LDP

The BGP TTL Security Hack (BTSH) was originally designed to protect the BGP infrastructure from CPU utilization-based attacks. It is derived from the fact that the vast majority of ISP eBGP peerings are established between adjacent routers. Since TTL spoofing is considered nearly impossible, a mechanism based on an expected TTL value can provide a simple and reasonably robust defense from infrastructure attacks based on forged BGP packets.

While TTL Security Hack (TSH) is most effective in protecting directly connected peers, it can also provide a lower level of protection to multi-hop sessions. When a multi-hop BGP session is required, the expected TTL value can be set to 255 minus the configured range-of-hops. This approach can provide a qualitatively lower degree of security for BGP (such as a DoS attack could, theoretically, be launched by compromising a box in the path). However, BTSH will catch a vast majority of observed distributed DoS (DDoS) attacks against eBGP.

TSH can be used to protect LDP peering sessions as well. For details, see draft-chen-ldp-ttl-xx.txt, *TTL-Based Security Option for LDP Hello Message*.

The TSH implementation supports the ability to configure TTL security per BGP/LDP peer and evaluate (in hardware) the incoming TTL value against the configured TTL value. If the incoming TTL value is less than the configured TTL value, the packets are discarded and a log is generated.

ECMP Support for LDP

ECMP support for LDP performs load balancing for LDP based LSPs by having multiple outgoing next-hops for a given IP prefix on ingress and transit LSRs.

An LSR that has multiple equal cost paths to a given IP prefix can receive an LDP label mapping for this prefix from each of the downstream next-hop peers. As the LDP implementation uses the liberal label retention mode, it retains all the labels for an IP prefix received from multiple next-hop peers.

Without ECMP support for LDP, only one of these next-hop peers will be selected and installed in the forwarding plane. The algorithm used to determine the next-hop peer to be selected involves looking up the route information obtained from the RTM for this prefix and finding the first valid LDP next-hop peer (for example, the first neighbor in the RTM entry from which a label mapping was received). If, for some reason, the outgoing label to the installed next-hop is no longer valid, say the session to the peer is lost or the peer withdraws the label, a new valid LDP next-hop peer will be selected out of the existing next-hop peers and LDP will reprogram the forwarding plane to use the label sent by this peer.

With ECMP support, all the valid LDP next-hop peers, those that sent a label mapping for a given IP prefix, will be installed in the forwarding plane. In both cases, ingress LER and transit LSR, an ingress label will be mapped to the nexthops that are in the RTM and from which a valid mapping label has been received. The forwarding plane will then use an internal hashing algorithm to determine how the traffic will be distributed amongst these multiple next-hops, assigning each “flow” to a particular next-hop.

The hash algorithm at LER and transit LSR are described in the LAG and ECMP Hashing section of the 7750 SR OS Interface Guide.

Label Operations

If an LSR is the ingress for a given IP prefix, LDP programs a push operation for the prefix in the forwarding engine. This creates an LSP ID to the Next Hop Label Forwarding Entry (NHLFE) (LTN) mapping and an LDP tunnel entry in the forwarding plane. LDP will also inform the Tunnel Table Manager (TTM) of this tunnel. Both the LTN entry and the tunnel entry will have a NHLFE for the label mapping that the LSR received from each of its next-hop peers.

If the LSR is to behave as a transit for a given IP prefix, LDP will program a swap operation for the prefix in the forwarding engine. This involves creating an Incoming Label Map (ILM) entry in the forwarding plane. The ILM entry will have to map an incoming label to possibly multiple NHLFEs. If an LSR is an egress for a given IP prefix, LDP will program a POP entry in the forwarding engine. This too will result in an ILM entry being created in the forwarding plane but with no NHLFEs.

When unlabeled packets arrive at the ingress LER, the forwarding plane will consult the LTN entry and will use a hashing algorithm to map the packet to one of the NHLFEs (push label) and forward the packet to the corresponding next-hop peer. For labeled packets arriving at a transit or egress LSR, the forwarding plane will consult the ILM entry and either use a hashing algorithm to map it to one of the NHLFEs if they exist (swap label) or simply route the packet if there are no NHLFEs (pop label).

Static FEC swap will not be activated unless there is a matching route in system route table that also matches the user configured static FEC next-hop.

Unnumbered Interface Support in LDP

This feature allows LDP to establish Hello adjacency and to resolve unicast and multicast FECs over unnumbered LDP interfaces.

This feature also extends the support of `lsp-ping`, `p2mp-lsp-ping`, and `ldp-treetrace` to test an LDP unicast or multicast FEC which is resolved over an unnumbered LDP interface.

Feature Configuration

This feature does not introduce a new CLI command for adding an unnumbered interface into LDP.

Note however that the **fec-originate** command has been extended to specify the interface name since an unnumbered interface will not have an IP address of its own. The user can however specify the interface name for numbered interfaces too.

See the CLI section for the changes to the **fec-originate** command.

Operation of LDP over an Unnumbered IP Interface

Consider the setup shown in [Figure 36](#).




Figure 36: LDP Adjacency and Session over Unnumbered Interface

LSR A and LSR B have the following LDP identifiers respectively:

<LSR Id=A> : <label space id=0>

<LSR Id=B> : <label space id=0>

There are two P2P unnumbered interfaces between LSR A and LSR B. These interfaces are identified on each system with their unique local link identifier. In other words, the combination of {Router-ID, Local Link Identifier} uniquely identifies the interface in OSPF or IS-IS throughout the network.

A borrowed IP address is also assigned to the interface to be used as the source address of IP packets which need to be originated from the interface. The borrowed IP address defaults to the system loopback interface address, A and B respectively in this setup. The user can change the borrowed IP interface to any configured IP interface, loopback or not, by applying the following command:

```
configure> router>interface>unnumbered [<ip-int-name| ip-address>]
```

When the unnumbered interface is added into LDP, it will have the following behavior.

Link LDP

Hello adjacency will be brought up using link Hello packet with source IP address set to the interface borrowed IP address and a destination IP address set to 224.0.0.2.

As a consequence of (1), Hello packets with the same source IP address should be accepted when received over parallel unnumbered interfaces from the same peer LSR-ID. The corresponding Hello adjacencies would be associated with a single LDP session.

The transport address for the TCP connection, which is encoded in the Hello packet, will always be set to the LSR-ID of the node regardless if the user enabled the interface option under **configure>router>ldp>interface-parameters>interface>transport-address**.

The user can configure the local-lsr-id option on the interface and change the value of the LSR-ID to either the local interface or to some other interface name, loopback or not, numbered or not. If the local interface is selected or the provided interface name corresponds to an unnumbered IP interface, the unnumbered interface borrowed IP address will be used as the LSR-ID. In all cases, the transport address for the LDP session will be updated to the new LSR-ID value but the link Hello packets will continue to use the interface borrowed IP address as the source IP address.

The LSR with the highest transport address, i.e., LSR-ID in this case, will bootstrap the TCP connection and LDP session.

Source and destination IP addresses of LDP packets are the transport addresses, i.e., LDP LSR-IDs of systems A and B in this case.

Targeted LDP

Source and destination addresses of targeted Hello packet are the LDP LSR-IDs of systems A and B.

The user can configure the `local-lsr-id` option on the targeted session and change the value of the LSR-ID to either the local interface or to some other interface name, loopback or not, numbered or not. If the local interface is selected or the provided interface name corresponds to an unnumbered IP interface, the unnumbered interface borrowed IP address will be used as the LSR-ID. In all cases, the transport address for the LDP session and the source IP address of targeted Hello message will be updated to the new LSR-ID value.

The LSR with the highest transport address, i.e., LSR-ID in this case, will bootstrap the TCP connection and LDP session.

Source and destination IP addresses of LDP messages are the transport addresses, i.e., LDP LSR-IDs of systems A and B in this case.

FEC Resolution

LDP will advertise/withdraw unnumbered interfaces using the Address/Address-Withdraw message. The borrowed IP address of the interface is used.

A FEC can be resolved to an unnumbered interface in the same way as it is resolved to a numbered interface. The outgoing interface and next-hop are looked up in RTM cache. The next-hop consists of the router-id and link identifier of the interface at the peer LSR.

LDP FEC ECMP next-hops over a mix of unnumbered and numbered interfaces is supported.

All LDP FEC types are supported.

The **fec-originate** command is supported when the next-hop is over an unnumbered interface.

All LDP features are supported except for the following:

- BFD cannot be enabled on an unnumbered LDP interface. This is a consequence of the fact that BFD is not supported on unnumbered IP interface on the 7x50 system.
- As a consequence of (1), LDP FRR procedures will not be triggered via a BFD session timeout but only by physical failures and local interface down events.
- Unnumbered IP interfaces cannot be added into LDP global and peer prefix policies.

LDP over RSVP Tunnels

LDP over RSVP-TE provides end-to-end tunnels that have two important properties, fast reroute and traffic engineering which are not available in LDP. LDP over RSVP-TE is focused at large networks (over 100 nodes in the network). Simply using end-to-end RSVP-TE tunnels will not scale. While an LER may not have that many tunnels, any transit node will potentially have thousands of LSPs, and if each transit node also has to deal with detours or bypass tunnels, this number can make the LSR overly burdened.

LDP over RSVP-TE allows tunneling of user packets using an LDP LSP inside an RSVP LSP. The main application of this feature is for deployment of MPLS based services, for example, VPRN, VLL, and VPLS services, in large scale networks across multiple IGP areas without requiring full mesh of RSVP LSPs between PE routers.

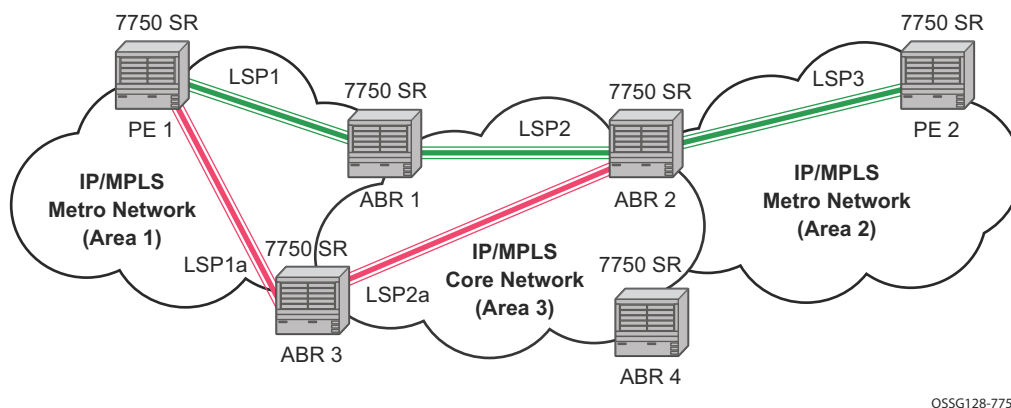


Figure 37: LDP over RSVP Application

The network displayed in [Figure 37](#) consists of two metro areas, Area 1 and 2 respectively, and a core area, Area 3. Each area makes use of TE LSPs to provide connectivity between the edge routers. In order to enable services between PE1 and PE2 across the three areas, LSP1, LSP2, and LSP3 are set up using RSVP-TE. There are in fact 6 LSPs required for bidirectional operation but we will refer to each bi-directional LSP with a single name, for example, LSP1. A targeted LDP (T-LDP) session is associated with each of these bidirectional LSP tunnels. That is, a T-LDP adjacency is created between PE1 and ABR1 and is associated with LSP1 at each end. The same is done for the LSP tunnel between ABR1 and ABR2, and finally between ABR2 and PE2. The loopback address of each of these routers is advertised using T-LDP. Similarly, backup bidirectional LDP over RSVP tunnels, LSP1a and LSP2a, are configured by way of ABR3.

This setup effectively creates an end-to-end LDP connectivity which can be used by all PEs to provision services. The RSVP LSPs are used as a transport vehicle to carry the LDP packets from

one area to another. Note that only the user packets are tunneled over the RSVP LSPs. The T-LDP control messages are still sent unlabeled using the IGP shortest path.

Note that in this application, the bi-directional RSVP LSP tunnels are not treated as IP interfaces and are not advertised back into the IGP. A PE must always rely on the IGP to look up the next hop for a service packet. LDP-over-RSVP introduces a new tunnel type, tunnel-in-tunnel, in addition to the existing LDP tunnel and RSVP tunnel types. If multiple tunnels types match the destination PE FEC lookup, LDP will prefer an LDP tunnel over an LDP-over-RSVP tunnel by default.

The design in [Figure 37](#) allows a service provider to build and expand each area independently without requiring a full mesh of RSVP LSPs between PEs across the three areas.

In order to participate in a VPRN service, PE1 and PE2 perform the autobind to LDP. The LDP label which represents the target PE loopback address is used below the RSVP LSP label. Therefore a 3 label stack is required.

In order to provide a VLL service, PE1 and PE2 are still required to set up a targeted LDP session directly between them. Again a 3 label stack is required, the RSVP LSP label, followed by the LDP label for the loopback address of the destination PE, and finally the pseudowire label (VC label).

This implementation supports a variation of the application in [Figure 37](#), in which area 1 is an LDP area. In that case, PE1 will push a two label stack while ABR1 will swap the LDP label and push the RSVP label as illustrated in [Figure 38](#). LDP-over-RSVP tunnels can also be used as IGP shortcuts.

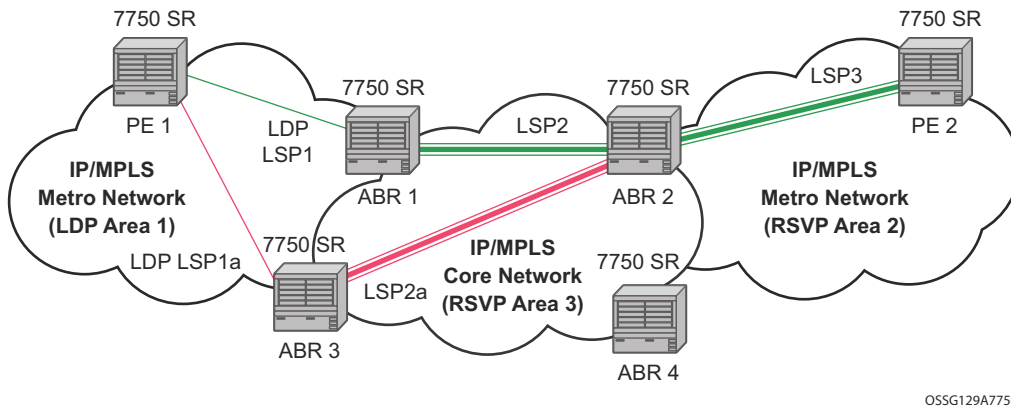


Figure 38: LDP over RSVP Application Variant

Signaling and Operation

- [LDP Label Distribution and FEC Resolution on page 451](#)
 - [Default FEC Resolution Procedure on page 452](#)
-

LDP Label Distribution and FEC Resolution

The user creates a targeted LDP (T-LDP) session to an ABR or the destination PE. This results in LDP hellos being sent between the two routers. These messages are sent unlabeled over the IGP path. Next, the user enables LDP tunneling on this T-LDP session and optionally specifies a list of LSP names to associate with this T-LDP session. By default, all RSVP LSPs which terminate on the T-LDP peer are candidates for LDP-over-RSVP tunnels. At this point in time, the LDP FECs resolving to RSVP LSPs are added into the Tunnel Table Manager as tunnel-in-tunnel type.

Note that if LDP is running on regular interfaces also, then the prefixes LDP learns are going to be distributed over both the T-LDP session as well as regular IGP interfaces. However, only /32 FEC prefixes will be resolved over RSVP LSPs. The policy controls which prefixes go over the T-LDP session, for example, only /32 prefixes, or a particular prefix range.

LDP-over-RSVP works with both OSPF and ISIS. These protocols include the advertising router when adding an entry to the RTM. LDP-over-RSVP tunnels can be used as shortcuts for BGP next-hop resolution.

Default FEC Resolution Procedure

When LDP tries to resolve a prefix received over a T-LDP session, it performs a lookup in the Routing Table Manager (RTM). This lookup returns the next hop to the destination PE and the advertising router (ABR or destination PE itself). If the next-hop router advertised the same FEC over link-level LDP, LDP will prefer the LDP tunnel by default unless the user explicitly changed the default preference using the system wide prefer-tunnel-in-tunnel command. If the LDP tunnel becomes unavailable, LDP will select an LDP-over-RSVP tunnel if available.

When searching for an LDP-over-RSVP tunnel, LDP selects the advertising router(s) with best route. If the advertising router matches the T-LDP peer, LDP then performs a second lookup for the advertising router in the Tunnel Table Manager (TTM) which returns the user configured RSVP LSP with the best metric. If there are more than one configured LSP with the best metric, LDP selects the first available LSP.

If all user configured RSVP LSPs are down, no more action is taken. If the user did not configure any LSPs under the T-LDP session, the lookup in TTM will return the first available RSVP LSP which terminates on the advertising router with the lowest metric.

FEC Resolution Procedure When prefer-tunnel-in-tunnel is Enabled

When LDP tries to resolve a prefix received over a T-LDP session, it performs a lookup in the Routing Table Manager (RTM). This lookup returns the next hop to the destination PE and the advertising router (ABR or destination PE itself).

When searching for an LDP-over-RSVP tunnel, LDP selects the advertising router(s) with best route. If the advertising router matches the targeted LDP peer, LDP then performs a second lookup for the advertising router in the Tunnel Table Manager (TTM) which returns the user configured RSVP LSP with the best metric. If there are more than one configured LSP with the best metric, LDP selects the first available LSP.

If all user configured RSVP LSPs are down, then an LDP tunnel will be selected if available.

If the user did not configure any LSPs under the T-LDP session, a lookup in TTM will return the first available RSVP LSP which terminates on the advertising router. If none are available, then an LDP tunnel will be selected if available.

Rerouting Around Failures

Every failure in the network can be protected against, except for the ingress and egress PEs. All other constructs have protection available. These constructs are LDP-over-RSVP tunnel and ABR.

- [LDP-over-RSVP Tunnel Protection on page 453](#)
 - [ABR Protection on page 453](#)
-

LDP-over-RSVP Tunnel Protection

An RSVP LSP can deal with a failure in two ways.

- If the LSP is a loosely routed LSP, then RSVP will find a new IGP path around the failure, and traffic will follow this new path. This may involve some churn in the network if the LSP comes down and then gets re-routed. The tunnel damping feature was implemented on the LSP so that all the dependent protocols and applications do not flap unnecessarily.
 - If the LSP is a CSPF-computed LSP with the fast reroute option enabled, then RSVP will switch to the detour path very quickly. From that point, a new LSP will be attempted from the head-end (global revertive). When the new LSP is in place, the traffic switches over to the new LSP with make-before-break.
-

ABR Protection

If an ABR fails, then routing around the ABR requires that a new next-hop LDP-over-RSVP tunnel be found to a backup ABR. If an ABR fails, then the T-LDP adjacency fails. Eventually, the backup ABR becomes the new next hop (after SPF converges), and LDP learns of the new next-hop and can reprogram the new path.

LDP over RSVP Without Area Boundary

The LDP over RSVP capability set includes the ability to stitch LDP-over-RSVP tunnels at internal (non-ABR) OSPF and IS-IS routers.

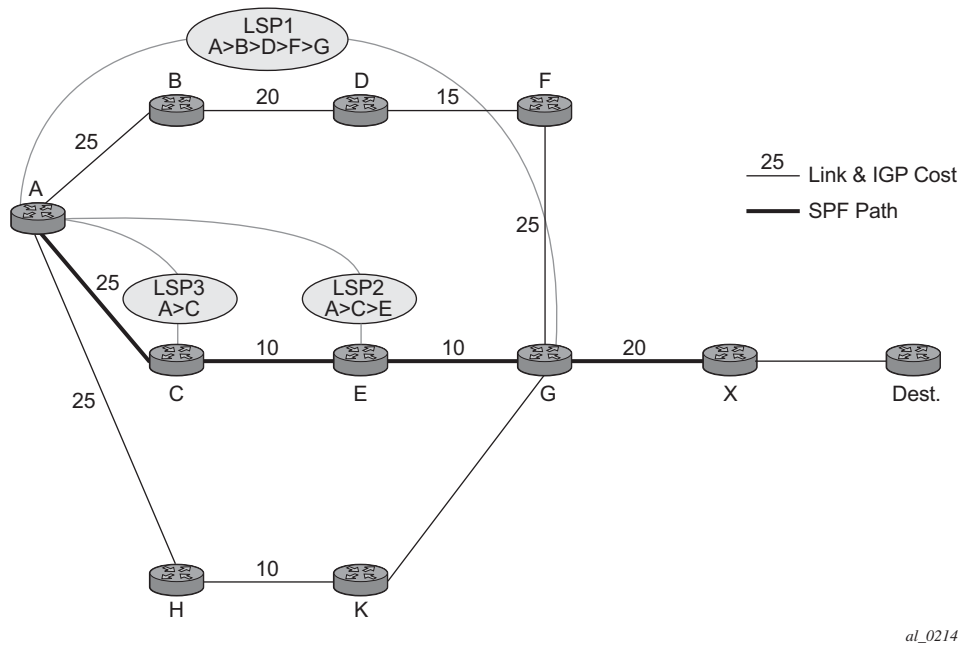


Figure 39: LDP over RSVP Without ABR Stitching Point

In Figure 39, assume that the user wants to use LDP over RSVP between router A and destination “Dest”. The first thing that happens is that either OSPF or IS-IS will perform an SPF calculation resulting in an SPF tree. This tree specifies the lowest possible cost to the destination. In the example shown, the destination “Dest” is reachable at the lowest cost through router X. The SPF tree will have the following path: A>C>E>G>X.

Using this SPF tree, router A will search for the endpoint that is closest (farthest/highest cost from the origin) to “Dest” that is eligible. Assuming that all LSPs in the above diagram are eligible, LSP endpoint G will be selected as it terminates on router G while other LSPs only reach routers C and E, respectively.

IGP and LSP metrics associated with the various LSP are ignored; only tunnel endpoint matters to IGP. The endpoint that terminates closest to “Dest” (highest IGP path cost) will be selected for

further selection of the LDP over RSVP tunnels to that endpoint. Note that the explicit path the tunnel takes may not match the IGP path the SPF computes.

If router A and G have an additional LSP terminating on router G, there would now be two tunnels both terminating on the same router closest to the final destination. For IGP, it does not make any difference on the numbers of LDPs to G, only that there is at least one LSP to G. In this case, the LSP metric will be considered by LDP when deciding which LSP to stitch for the LDP over RSVP connection.

The IGP only passes endpoint information to LDP. LDP looks up the tunnel table for all tunnels to that endpoint and picks up the one with the least tunnel metric. There may be many tunnels with the same least cost. Note that only /32 FEC prefixes will be resolved over RSVP LSPs within an area.

LDP over RSVP and ECMP

ECMP for LDP over RSVP is supported (also see [ECMP Support for LDP on page 444](#)). If ECMP applies, all LSP endpoints found over the ECMP IGP path will be installed in the routing table by the IGP for consideration by LDP. It is important to note that IGP costs to each endpoint may differ because IGP selects the farthest endpoint per ECMP path.

LDP will choose the endpoint that is highest cost in the route entry and will do further tunnel selection over those endpoints. If there are multiple endpoints with equal highest cost, then LDP will consider all of them.

LDP ECMP Uniform Failover

LDP ECMP uniform failover allows the fast re-distribution by the ingress data path of packets forwarded over an LDP FEC next-hop to other next-hops of the same FEC when the currently used next-hop fails. The switchover is performed within a bounded time, which does not depend on the number of impacted LDP ILMs (LSR role) or service records (ingress LER role). The uniform failover time is only supported for a single LDP interface or LDP next-hop failure event.

This feature complements the coverage provided by the LDP Fast-ReRoute (FRR) feature, which provides a Loop-Free Alternate (LFA) backup next-hop with uniform failover time. Prefixes that have one or more ECMP next-hop protection are not programmed with a LFA back-up next-hop, and vice-versa.

The LDP ECMP uniform failover feature builds on the concept of Protect Group ID (PG-ID) introduced in LDP FRR. LDP assigns a unique PG-ID to all FECs that have their primary Next-Hop Label Forwarding Entry (NHLFE) resolved to the same outgoing interface and next-hop.

When an ILM record (LSR role) or LSPid-to-NHLFE (LTN) record (LER role) is created on the IOM, it has the PG-ID of each ECMP NHLFE the FEC is using.

When a packet is received on this ILM/LTN, the hash routine selects one of the up to 32, or the ECMP value configured on the system, whichever is less, ECMP NHLFEs for the FEC based on a hash of the packet's header. If the selected NHLFE has its PG-ID in DOWN state, the hash routine re-computes the hash to select a backup NHLFE among the first 16, or the ECMP value configured on the system, whichever is less, NHLFEs of the FEC, excluding the one that is in DOWN state. Packets of the subset of flows that resolved to the failed NHLFE are thus sprayed among a maximum of 16 NHLFEs.

LDP then re-computes the new ECMP set to exclude the failed path and downloads it into the IOM. At that point, the hash routine will update the computation and begin spraying over the updated set of NHLFEs.

LDP sends the DOWN state update of the PG-ID to the IOM when the outgoing interface or a specific LDP next-hop goes down. This can be the result of any of the following events:

- Interface failure detected directly.
- Failure of the LDP session detected via T-LDP BFD or LDP Keep-Alive.
- Failure of LDP Hello adjacency detected via link LDP BFD or LDP Hello.

In addition, PIP will send an interface down event to the IOM if the interface failure is detected by other means than the LDP control plane or BFD. In that case, all PG-IDs associated with this interface will have their state updated by the IOM.

When tunneling LDP packets over an RSVP LSP, it is the detection of the T-LDP session going down, via BFD or Keep-Alive, which triggers the LDP ECMP uniform failover procedures. If the

RSVP LSP alone fails and the latter is not protected by RSVP FRR, the failure event will trigger the re-resolution of the impacted FECs in the slow path.

When a multicast LDP (mLDP) FEC is resolved over ECMP links to the same downstream LDP LSR, the PG-ID DOWN state will cause packets of the FEC resolved to the failed link to be switched to another link using the linear FRR switchover procedures.

The LDP ECMP uniform failover is not supported in the following forwarding contexts:

- VLL services packets received on a SAP when the LDP ECMP next-hop selection is based on service-id. When HPOL is enabled in the ingress SAP, the spraying of packets based on a hash of the packet's headers is performed and thus the feature is supported.
- VPLS BUM packets.
- Packets forwarded to an IES/VPRN spoke-interface.
- Packets forwarded towards VPLS spoke in routed VPLS.

Finally, note that the LDP ECMP uniform failover is only supported for a single LDP interface, LDP next-hop, or peer failure event.

LDP Fast-Reroute for IS-IS and OSPF Prefixes

LDP Fast Re-Route (FRR) is a feature which allows the user to provide local protection for an LDP FEC by pre-computing and downloading to IOM both a primary and a backup NHLFE for this FEC.

The primary NHLFE corresponds to the label of the FEC received from the primary next-hop as per standard LDP resolution of the FEC prefix in RTM. The backup NHLFE corresponds to the label received for the same FEC from a Loop-Free Alternate (LFA) next-hop.

The LFA next-hop pre-computation by IGP is described in RFC 5286 – “Basic Specification for IP Fast Reroute: Loop-Free Alternates”. LDP FRR relies on using the label-FEC binding received from the LFA next-hop to forward traffic for a given prefix as soon as the primary next-hop is not available. This means that a node resumes forwarding LDP packets to a destination prefix without waiting for the routing convergence. The label-FEC binding is received from the loop-free alternate next-hop ahead of time and is stored in the Label Information Base since LDP on the router operates in the liberal retention mode.

This feature requires that IGP performs the Shortest Path First (SPF) computation of an LFA next-hop, in addition to the primary next-hop, for all prefixes used by LDP to resolve FECs. IGP also populates both routes in the Routing Table Manager (RTM).

LDP FRR Configuration

The user enables Loop-Free Alternate (LFA) computation by SPF under the IS-IS or OSPF routing protocol level:

```
config>router>isis>loopfree-alternate  
config>router>ospf>loopfree-alternate.
```

The above commands instruct the IGP SPF to attempt to pre-compute both a primary next-hop and an LFA next-hop for every learned prefix. When found, the LFA next-hop is populated into the RTM along with the primary next-hop for the prefix.

Next the user enables the use by LDP of the LFA next-hop by configuring the following option:

```
config>router>ldp>fast-reroute
```

When this command is enabled, LDP will use both the primary next-hop and LFA next-hop, when available, for resolving the next-hop of an LDP FEC against the corresponding prefix in the RTM. This will result in LDP programming a primary NHLFE and a backup NHLFE into the IOM for each next-hop of a FEC prefix for the purpose of forwarding packets over the LDP FEC.

Note that because LDP can detect the loss of a neighbor/next-hop independently, it is possible that it switches to the LFA next-hop while IGP is still using the primary next-hop. In order to avoid this situation, it is recommended to enable IGP-LDP synchronization on the LDP interface:

```
config>router>interface>ldp-sync-timer seconds
```

Reducing the Scope of the LFA Calculation by SPF

The user can instruct IGP to not include all interfaces participating in a specific IS-IS level or OSPF area in the SPF LFA computation. This provides a way of reducing the LFA SPF calculation where it is not needed.

```
config>router>isis>level>loopfree-alternate-exclude
config>router>ospf>area>loopfree-alternate-exclude
```

Note that if IGP shortcut are also enabled in LFA SPF, as explained in [Section 5.3.2](#), LSPs with destination address in that IS-IS level or OSPF area are also not included in the LFA SPF calculation.

The user can also exclude a specific IP interface from being included in the LFA SPF computation by IS-IS or OSPF:

```
config>router>isis>interface> loopfree-alternate-exclude
config>router>ospf>area>interface> loopfree-alternate-exclude
```

Note that when an interface is excluded from the LFA SPF in IS-IS, it is excluded in both level 1 and level 2. When the user excludes an interface from the LFA SPF in OSPF, it is excluded in all areas. However, the above OSPF command can only be executed under the area in which the specified interface is primary and once enabled, the interface is excluded in that area and in all other areas where the interface is secondary. If the user attempts to apply it to an area where the interface is secondary, the command will fail.

Finally, the user can apply the same above commands for an OSPF instance within a VPRN service:

```
config>service>vprn>ospf>area>loopfree-alternate-exclude
config>service>vprn>ospf>area>interface>loopfree-alternate-exclude
```

LDP FRR Procedures

The LDP FEC resolution when LDP FRR is not enabled operates as follows. When LDP receives a *FEC*, *label* binding for a prefix, it will resolve it by checking if the exact prefix, or a longest match prefix when the **aggregate-prefix-match option** is enabled in LDP, exists in the routing table and is resolved against a next-hop which is an address belonging to the LDP peer which

advertized the binding, as identified by its LSR-id. When the next-hop is no longer available, LDP de-activates the FEC and de-programs the NHLFE in the data path. LDP will also immediately withdraw the labels it advertised for this FEC and deletes the ILM in the data path unless the user configured the **label-withdrawal-delay** option to delay this operation. Traffic that is received while the ILM is still in the data path is dropped. When routing computes and populates the routing table with a new next-hop for the prefix, LDP resolves again the FEC and programs the data path accordingly.

When LDP FRR is enabled and an LFA backup next-hop exists for the FEC prefix in RTM, or for the longest prefix the FEC prefix matches to when **aggregate-prefix-match** option is enabled in LDP, LDP will resolve the FEC as above but will program the data path with both a primary NHLFE and a backup NHLFE for each next-hop of the FEC.

In order perform a switchover to the backup NHLFE in the fast path, LDP follows the uniform FRR failover procedures which are also supported with RSVP FRR.

When any of the following events occurs, LDP instructs in the fast path the IOM to enable the backup NHLFE for each FEC next-hop impacted by this event. The IOM do that by simply flipping a single state bit associated with the failed interface or neighbor/next-hop:

1. An LDP interface goes operationally down, or is admin shutdown. In this case, LDP sends a neighbor/next-hop down message to the IOM for each LDP peer it has adjacency with over this interface.
2. An LDP session to a peer went down as the result of the Hello or Keep-Alive timer expiring over a specific interface. In this case, LDP sends a neighbor/next-hop down message to the IOM for this LDP peer only.
3. The TCP connection used by a link LDP session to a peer went down, due say to next-hop tracking of the LDP transport address in RTM, which brings down the LDP session. In this case, LDP sends a neighbor/next-hop down message to the IOM for this LDP peer only.
4. A BFD session, enabled on a T-LDP session to a peer, times-out and as a result the link LDP session to the same peer and which uses the same TCP connection as the T-LDP session goes also down. In this case, LDP sends a neighbor/next-hop down message to the IOM for this LDP peer only.
5. A BFD session enabled on the LDP interface to a directly connected peer, times-out and brings down the link LDP session to this peer. In this case, LDP sends a neighbor/next-hop down message to the IOM for this LDP peer only. BFD support on LDP interfaces is a new feature introduced for faster tracking of link LDP peers. See Section 1.2.1 for more details.

The tunnel-down-dump-time option or the label-withdrawal-delay option, when enabled, does not cause the corresponding timer to be activated for a FEC as long as a backup NHLFE is still available.

Link LDP Hello Adjacency Tracking with BFD

LDP can only track an LDP peer with which it established a link LDP session with using the Hello and Keep-Alive timers. If an IGP protocol registered with BFD on an IP interface to track a neighbor, and the BFD session times out, the next-hop for prefixes advertised by the neighbor are no longer resolved. This however does not bring down the link LDP session to the peer since the LDP peer is not directly tracked by BFD. More importantly the LSR-id of the LDP peer may not coincide with the neighbor's router-id IGP is tracking by way of BFD.

In order to properly track the link LDP peer, LDP needs to track the Hello adjacency to its peer by registering with BFD. This way, the peer next-hop is tracked.

The user enables Hello adjacency tracking with BFD by enabling BFD on an LDP interface:

```
config>router>ldp>interface-parameters>interface>enable-bfd
```

The parameters used for the BFD session, i.e., transmit-interval, receive-interval, and multiplier, are those configured under the IP interface in existing implementation:

```
config>router>interface>bfd
```

When multiple links exist to the same LDP peer, a Hello adjacency is established over each link but only a single LDP session will exist to the peer and will use a TCP connection over one of the link interfaces. Also, a separate BFD session should be enabled on each LDP interface. If a BFD session times out on a specific link, LDP will immediately bring down the Hello adjacency on that link. In addition, if there are FECs which have their primary NHLFE over this link, LDP triggers the LDP FRR procedures by sending to IOM the neighbor/next-hop down message. This will result in moving the traffic of the impacted FECs to an LFA next-hop on a different link to the same LDP peer or to an LFA backup next-hop on a different LDP peer depending on the lowest backup cost path selected by the IGP SPF.

As soon as the last Hello adjacency goes down due to BFD timing out, the LDP session goes down and the LDP FRR procedures will be triggered. This will result in moving the traffic to an LFA backup next-hop on a different LDP peer.

ECMP Considerations

Whenever the SPF computation determined there is more than one primary next-hop for a prefix, it will not program any LFA next-hop in RTM. Thus, the LDP FEC will resolve to the multiple primary next-hops in this case which provides the required protection.

Also note that when the system ECMP value is set to **ecmp=1** or to **no ecmp**, which translates to the same and is the default value, SPF will be able to use the overflow ECMP links as LFA next-hops in these two cases.

LDP FRR and LDP Shortcut

When LDP FRR is enabled in LDP and the `ldp-shortcut` option is enabled in the router level, in transit IPv4 packets and specific CPM generated IPv4 control plane packets with a prefix resolving to the LDP shortcut are protected by the backup LDP NHLFE.

LDP FRR and LDP-over-RSVP

When LDP-over-RSVP is enabled, the RSVP LSP is modeled as an endpoint, i.e., the destination node of the LSP, and not as a link in the IGP SPF. Thus, it is not possible for IGP to compute a primary or alternate next-hop for a prefix which FEC path is tunneled over the RSVP LSP. Only LDP is aware of the FEC tunneling but it cannot determine on its own a loop-free backup path when it resolves the FEC to an RSVP LSP.

As a result, LDP does not activate the LFA next-hop it learned from RTM for a FEC prefix when the FEC is resolved to an RSVP LSP. LDP will activate the LFA next-hop as soon as the FEC is resolved to direct primary next-hop.

LDP FEC tunneled over an RSVP LSP due to enabling the LDP-over-RSVP feature will thus not support the LDP FRR procedures and will follow the slow path procedure of prior implementation.

Note that when the user enables the **lfa-only** option for an RSVP LSP, as described in [Loop-Free Alternate Calculation in the Presence of IGP shortcuts on page 465](#), such an LSP will not be used by LDP to tunnel an LDP FEC even when IGP shortcut is disabled but LDP-over-RSVP is enabled in IGP.

LDP FRR and RSVP Shortcut (IGP Shortcut)

When an RSVP LSP is used as a shortcut by IGP, it is included by SPF as a P2P link and can also be optionally advertised into the rest of the network by IGP. Thus the SPF is able of using a tunneled next-hop as the primary next-hop for a given prefix. LDP is also able of resolving a FEC to a tunneled next-hop when the IGP shortcut feature is enabled.

When both IGP shortcut and LFA are enabled in IS-IS or OSPF, and LDP FRR is also enabled, then the following additional LDP FRR capabilities are supported:

1. A FEC which is resolved to a direct primary next-hop can be backed up by a LFA tunneled next-hop.
2. A FEC which is resolved to a tunneled primary next-hop will not have an LFA next-hop. It will rely on RSVP FRR for protection.

The LFA SPF is extended to use IGP shortcuts as LFA next-hops as explained in [Loop-Free Alternate Calculation in the Presence of IGP shortcuts on page 465](#).

IS-IS and OSPF Support for Loop-Free Alternate Calculation

SPF computation in IS-IS and OSPF is enhanced to compute LFA alternate routes for each learned prefix and populate it in RTM.

Figure 40 illustrates a simple network topology with point-to-point (P2P) interfaces and highlights three routes to reach router R5 from router R1.

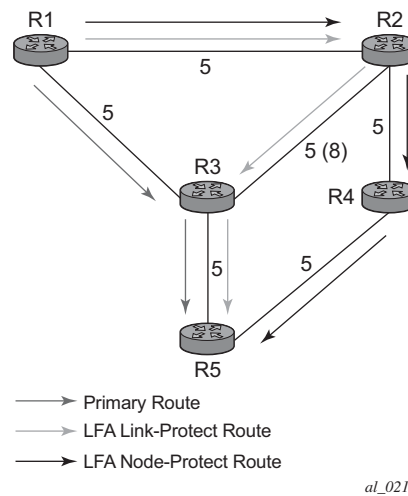



Figure 40: Topology with Primary and LFA Routes

The primary route is by way of R3. The LFA route by way of R2 has two equal cost paths to reach R5. The path by way of R3 protects against failure of link R1-R3. This route is computed by R1 by checking that the cost for R2 to reach R5 by way of R3 is lower than the cost by way of routes R1 and R3. This condition is referred to as the *loop-free criterion*. R2 must be loop-free with respect to source node R1.

The path by way of R2 and R4 can be used to protect against the failure of router R3. However, with the link R2-R3 metric set to 5, R2 sees the same cost to forward a packet to R5 by way of R3 and R4. Thus R1 cannot guarantee that enabling the LFA next-hop R2 will protect against R3 node failure. This means that the LFA next-hop R2 provides link-protection only for prefix R5. If the metric of link R2-R3 is changed to 8, then the LFA next-hop R2 provides node protection since a packet to R5 will always go over R4. In other words it is required that R2 becomes loop-free with respect to both the source node R1 and the protected node R3.

Consider the case where the primary next-hop uses a broadcast interface as illustrated in [Figure 41](#)



al_0216

Figure 41: Example Topology with Broadcast Interfaces

In order for next-hop R2 to be a link-protect LFA for route R5 from R1, it must be loop-free with respect to the R1-R3 link's Pseudo-Node (PN). However, since R2 has also a link to that PN, its cost to reach R5 by way of the PN or router R4 are the same. Thus R1 cannot guarantee that enabling the LFA next-hop R2 will protect against a failure impacting link R1-PN since this may cause the entire subnet represented by the PN to go down. If the metric of link R2-PN is changed to 8, then R2 next-hop will be an LFA providing link protection.

The following are the detailed rules for this criterion as provided in RFC 5286:

- **Rule 1: Link-protect LFA backup next-hop (primary next-hop R1-R3 is a P2P interface):**
 $Distance_opt(R2, R5) < Distance_opt(R2, R1) + Distance_opt(R1, R5)$
 and,
 $Distance_opt(R2, R5) \geq Distance_opt(R2, R3) + Distance_opt(R3, R5)$
- **Rule 2: Node-protect LFA backup next-hop (primary next-hop R1-R3 is a P2P interface):**
 $Distance_opt(R2, R5) < Distance_opt(R2, R1) + Distance_opt(R1, R5)$
 and,
 $Distance_opt(R2, R5) < Distance_opt(R2, R3) + Distance_opt(R3, R5)$
- **Rule 3: Link-protect LFA backup next-hop (primary next-hop R1-R3 is a broadcast interface):**
 $Distance_opt(R2, R5) < Distance_opt(R2, R1) + Distance_opt(R1, R5)$
 and,
 $Distance_opt(R2, R5) < Distance_opt(R2, PN) + Distance_opt(PN, R5)$
 where; PN stands for the R1-R3 link Pseudo-Node.

For the case of P2P interface, if SPF finds multiple LFA next-hops for a given primary next-hop, it follows the following selection algorithm:

- A) It will pick the node-protect type in favor of the link-protect type.

- B) If there is more than one LFA next-hop within the selected type, then it will pick one based on the least cost.
- C) If more than one LFA next-hop with the same cost results from Step B, then SPF will select the first one. This is not a deterministic selection and will vary following each SPF calculation.

For the case of a broadcast interface, a node-protect LFA is not necessarily a link protect LFA if the path to the LFA next-hop goes over the same PN as the primary next-hop. Similarly, a link protect LFA may not guarantee link protection if it goes over the same PN as the primary next-hop.

The selection algorithm when SPF finds multiple LFA next-hops for a given primary next-hop is modified as follows:

- A) The algorithm splits the LFA next-hops into two sets:
 - The first set consists of LFA next-hops which *do not* go over the PN used by primary next-hop.
 - The second set consists of LFA next-hops which *do* go over the PN used by the primary next-hop.
- B) If there is more than one LFA next-hop in the first set, it will pick the node-protect type in favor of the link-protect type.
- C) If there is more than one LFA next-hop within the selected type, then it will pick one based on the least cost.
- D) If more than one LFA next-hop with equal cost results from Step C, SPF will select the first one from the remaining set. This is not a deterministic selection and will vary following each SPF calculation.
- E) If no LFA next-hop results from Step D, SPF will rerun Steps B-D using the second set.

Note this algorithm is more flexible than strictly applying Rule 3 above; the link protect rule in the presence of a PN and specified in RFC 5286. A node-protect LFA which does not avoid the PN; does not guarantee link protection, can still be selected as a last resort. The same thing, a link-protect LFA which does not avoid the PN may still be selected as a last resort. Both the computed primary next-hop and LFA next-hop for a given prefix are programmed into RTM.

Loop-Free Alternate Calculation in the Presence of IGP shortcuts

In order to expand the coverage of the LFA backup protection in a network, RSVP LSP based IGP shortcuts can be placed selectively in parts of the network and be used as an LFA backup next-hop.

When IGP shortcut is enabled in IS-IS or OSPF on a given node, all RSVP LSP originating on this node and with a destination address matching the router-id of any other node in the network are included in the main SPF by default.

In order to limit the time it takes to compute the LFA SPF, the user must explicitly enable the use of an IGP shortcut as LFA backup next-hop using one of a couple of new optional argument for the existing LSP level IGP shortcut command:

```
config>router>mpls>lsp>igp-shortcut [lfa-protect | lfa-only]
```

The **lfa-protect** option allows an LSP to be included in both the main SPF and the LFA SPFs. For a given prefix, the LSP can be used either as a primary next-hop or as an LFA next-hop but not both. If the main SPF computation selected a tunneled primary next-hop for a prefix, the LFA SPF will not select an LFA next-hop for this prefix and the protection of this prefix will rely on the RSVP LSP FRR protection. If the main SPF computation selected a direct primary next-hop, then the LFA SPF will select an LFA next-hop for this prefix but will prefer a direct LFA next-hop over a tunneled LFA next-hop.

The **lfa-only** option allows an LSP to be included in the LFA SPFs only such that the introduction of IGP shortcuts does not impact the main SPF decision. For a given prefix, the main SPF always selects a direct primary next-hop. The LFA SPF will select an LFA next-hop for this prefix but will prefer a direct LFA next-hop over a tunneled LFA next-hop.

Thus the selection algorithm in Section 1.3 when SPF finds multiple LFA next-hops for a given primary next-hop is modified as follows:

- A) The algorithm splits the LFA next-hops into two sets:
 - the first set consists of direct LFA next-hops
 - the second set consists of tunneled LFA next-hops. after excluding the LSPs which use the same outgoing interface as the primary next-hop.
- B) The algorithms continues with first set if not empty, otherwise it continues with second set.
- C) If the second set is used, the algorithm selects the tunneled LFA next-hop which endpoint corresponds to the node advertising the prefix.
 - If more than one tunneled next-hop exists, it selects the one with the lowest LSP metric.
 - If still more than one tunneled next-hop exists, it selects the one with the lowest tunnel-id.
 - If none is available, it continues with rest of the tunneled LFAs in second set.

- D) Within the selected set, the algorithm splits the LFA next-hops into two sets:
 - The first set consists of LFA next-hops which do not go over the PN used by primary next-hop.
 - The second set consists of LFA next-hops which go over the PN used by the primary next-hop.
 - E) If there is more than one LFA next-hop in the selected set, it will pick the node-protect type in favor of the link-protect type.
 - F) If there is more than one LFA next-hop within the selected type, then it will pick one based on the least total cost for the prefix. For a tunneled next-hop, it means the LSP metric plus the cost of the LSP endpoint to the destination of the prefix.
 - G) If there is more than one LFA next-hop within the selected type (ecmp-case) in the first set, it will select the first direct next-hop from the remaining set. This is not a deterministic selection and will vary following each SPF calculation.
 - H) If there is more than one LFA next-hop within the selected type (ecmp-case) in the second set, it will pick the tunneled next-hop with the lowest cost from the endpoint of the LSP to the destination prefix. If there remains more than one, it will pick the tunneled next-hop with the lowest tunnel-id.
-

Loop-Free Alternate Shortest Path First (LFA SPF) Policies

An LFA SPF policy allows the user to apply specific criteria, such as admin group and SRLG constraints, to the selection of a LFA backup next-hop for a subset of prefixes that resolve to a specific primary next-hop. See more details in the section titled “*Loop-Free Alternate Shortest Path First (LFA SPF) Policies*” in the *Routing Protocols Guide*.

Loop-Free Alternate Calculation for Inter-Area/inter-Level Prefixes

When SPF resolves OSPF inter-area prefixes or IS-IS inter-level prefixes, it will compute an LFA backup next-hop to the same exit area/border router as used by the primary next-hop.

mLDP Fast Upstream Switchover

mLDP Fast Upstream Switchover allows a downstream LSR of an multicast LDP (mLDP) FEC to perform a fast switchover and source the traffic from another upstream LSR while IGP is converging due to a failure of the primary next-hop of the P2MP FEC. In a sense, it provides an upstream Fast-Reroute (FRR) capability for the mLDP packets. It does it at the expense of traffic duplication from two different upstream nodes into the node that performs the fast upstream switchover.

When this feature is enabled and LDP is resolving an mLDP FEC received from a downstream LSR, it checks if an Equal-Cost Multi-Path (ECMP) next-hop or a Loop-Free Alternate (LFA) next-hop exist to the root LSR node. If LDP finds one, it programs a primary ILM on the interface corresponding to the primary next-hop and a backup ILM on the interface corresponding to the ECMP or LFA next-hop. LDP then sends the corresponding labels to the upstream LSR nodes. In normal operation, the primary ILM accepts packets while the backup ILM drops them. If the node detects that the interface or the upstream LSR of the primary ILM is down, the backup ILM will then start accepting packets.

In order to make use of the ECMP next-hop, the user must configure the ECMP value in the system to at least two (2). In order to make use of the LFA next-hop, the user must enable LFA and IP FRR options under the IGP instance.

LDP FEC to BGP Label Route Stitching

The stitching of an LDP FEC to a BGP labeled route allows LDP capable PE devices to offer services to PE routers in other areas or domains without the need to support BGP labeled routes.

This feature is used in a large network to provide services across multiple areas or autonomous systems. [Figure 42](#) shows a network with a core area and regional areas.




Figure 42: Application of LDP to BGP FEC Stitching

Specific /32 routes in a regional area are not redistributed into the core area. Therefore, only nodes within a regional area and the ABR nodes in the same area exchange LDP FECs. A PE router, for example, PE21, in a regional area learns the reachability of PE routers in other regional areas by way of RFC 3107 BGP labeled routes redistributed by the remote ABR nodes by way of the core area. The remote ABR then sets the next-hop self on the labeled routes before re-distributing them into the core area. The local ABR for PE2, for example, ABR3 may or may not set next-hop self when it re-distributes these labeled BGP routes from the core area to the local regional area.

When forwarding a service packet to the remote PE, PE21 inserts a VC label, the BGP route label to reach the remote PE, and an LDP label to reach either ABR3, if ABR3 sets next-hop self, or ABR1.

In the same network, an MPLS capable DSLAM also act as PE router for VLL services and will need to establish a PW to a PE in a different regional area by way of router PE21, acting now as an LSR. To achieve that, PE21 is required to perform the following operations:

- Translate the LDP FEC it learned from the DSLAM into a BGP labeled route and re-distribute it by way of iBGP within its area. This is in addition to redistributing the FEC to its LDP neighbors in the same area.
 - Translate the BGP labeled routes it learns through iBGP into an LDP FEC and re-distribute it to its LDP neighbors in the same area. In the application in [Figure 42](#), the DSLAM requests the LDP FEC of the remote PE router using LDP Downstream on Demand (DoD).
 - When a packet is received from the DSLAM, PE21 swaps the LDP label into a BGP label and pushes the LDP label to reach ABR3 or ABR1. When a packet is received from ABR3, the top label is removed and the BGP label is swapped for the LDP label corresponding to the DSLAM FEC.
-

Configuration

The user enables the stitching of routes between LDP and BGP by configuring separately tunnel table route export policies in both protocols and enabling the advertising of RFC 3107 formatted labeled routes for prefixes learned from LDP FECs.

The route export policy in BGP instructs BGP to listen to LDP route entries in the CPM tunnel table. If a /32 LDP FEC prefix matches an entry in the export policy, BGP originates a BGP labeled route, stitches it to the LDP FEC, and re-distributes the BGP labeled route to its iBGP neighbors.

The user adds LDP FEC prefixes with the statement ‘from protocol ldp’ in the configuration of the existing BGP export policy at the global level, the peer-group level, or at the peer level using the commands:

- **configure>router>bgp>export** *policy-name*
- **configure>router>bgp>group>export** *policy-name*
- **configure>router>bgp>group>neighbour>export** *policy-name*

To indicate to BGP to evaluate the entries with the ‘from protocol ldp’ statement in the export policy when applied to a specific BGP neighbor, a new argument is added to the existing advertise-label command:

```
configure>router>bgp>group>neighbour>advertise-label ipv4 include-ldp-prefix
```

Without the new **include-ldp-prefix** argument, only core IPv4 routes learned from RTM are advertised as BGP labeled routes to this neighbor. And the stitching of LDP FEC to the BGP labeled route is not performed for this neighbor even if the same prefix was learned from LDP.

The tunnel table route export policy in LDP instructs LDP to listen to BGP route entries in the CPM Tunnel Table. If a /32 BGP labeled route matches a prefix entry in the export policy, LDP originates an LDP FEC for the prefix, stitches it to the BGP labeled route, and re-distributes the LDP FEC its iBGP neighbors.

The user adds BGP labeled route prefixes with the statement ‘from protocol bgp’ in the configuration of a new LDP tunnel table export policy using the command:

```
configure>router>ldp>export-tunnel-table policy-name.
```

Note that the ‘from protocol’ statement has an effect only when the protocol value is ldp. Policy entries with protocol values of rsvp, bgp, or any value other than ldp are ignored at the time the policy is applied to LDP.

Detailed LDP FEC Resolution

When a 7x50 LSR receives a FEC-label binding from an LDP neighbor for a given specific FEC1 element, the following procedures are performed.

1. LDP installs the FEC if:
 - It was able to perform a successful exact match or a longest match, if aggregate-prefix-match option is enabled in LDP, of the FEC /32 prefix with a prefix entry in the routing table.
 - The advertising LDP neighbor is the next-hop to reach the FEC prefix.
2. When such a FEC-label binding has been installed in the LDP FIB, LDP will perform the following:
 - Program a push and a swap NHLFE entries in the egress data path to forward packets to FEC1.
 - Program the CPM tunnel table with a tunnel entry for the NHLFE.
 - Advertise a new FEC-label binding for FEC1 to all its LDP neighbors according to the global and per-peer LDP prefix export policies.
 - Install the ILM entry pointing to the swap NHLFE.
3. When BGP learns the LDP FEC by way of the CPM tunnel table and the FEC prefix exists in the BGP route export policy, it will perform the following:
 - Originate a labeled BGP route for the same prefix with this node as the next-hop and advertise it by way of iBGP to its BGP neighbors, for example, the local ABR/ASBR nodes, which have the advertise-label for LDP FEC prefixes is enabled.
 - Install the ILM entry pointing to the swap NHLFE programmed by LDP.

Detailed BGP Labeled Route Resolution

When a 7x50 LSR receives a BGP labeled route by way of iBGP for a given specific /32 prefix, the following procedures are performed.

1. BGP resolves and installs the route in BGP if:
 - There exists an LDP LSP to the BGP neighbor, for example, the ABR or ASBR, which advertised it and which is the next-hop of the BGP labeled route.
 2. Once the BGP route is installed, BGP programs the following:
 - Push NHLFE in the egress data path to forward packets to this BGP labeled route.
 - The CPM tunnel table with a tunnel entry for the NHLFE.
 3. When LDP learns the BGP labeled route by way of the CPM tunnel table and the prefix exists in the new LDP tunnel table route export policy, it performs the following:
 - Advertise a new LDP FEC-label binding for the same prefix to its LDP neighbors according the global and per-peer LDP export prefix policies. If LDP already advertised a FEC for the same /32 prefix after receiving it from an LDP neighbor then no action is required. For LDP neighbors that negotiated LDP Downstream on Demand (DoD), the FEC is advertised only when this node receives a Label Request message for this FEC from its neighbor.
 - Install the ILM entry pointing the BGP NHLFE if a new LDP FEC-label binding is advertised. If an ILM entry exists and points to an LDP NHLFE for the same prefix then no update to ILM entry is performed. The LDP route has always preference over the BGP labeled route.
-

Data Plane Forwarding

When a packet is received from an LDP neighbor, the 7x50 LSR swaps the LDP label into a BGP label and pushes the LDP label to reach the BGP neighbor, for example, ABR/ASBR, which advertised the BGP labeled route with itself as the next-hop.

When a packet is received from a BGP neighbor such as an ABR/ASBR, the top label is removed and the BGP label is swapped for the LDP label to reach the next-hop for the prefix.

Automatic Creation of a Targeted Hello Adjacency and LDP Session

This feature enables the automatic creation of a targeted Hello adjacency and LDP session to a discovered peer.

Feature Configuration

The user first creates a targeted LDP session peer parameter template:

```
config>router>ldp>targeted-session>peer-template template-name
```

Inside the template the user configures the common T-LDP session parameters or options shared by all peers using this template. These are the following:

bfd-enable, hello, hello-reduction, keepalive, local-lsr-id, and tunneling.

Note that the tunneling option does not support adding explicit RSVP LSP names. Thus, LDP will select RSVP LSP for an endpoint in LDP-over-RSVP directly from the Tunnel Table Manager (TTM).

Then the user references the peer prefix list which is defined inside a policy statement defined in the global policy manager.

```
config>router>ldp>targeted-session>peer-template-map peer-template template-name policy peer-prefix-policy
```

Each application of a targeted session template to a given prefix in the prefix list will result in the establishment of a targeted Hello adjacency to an LDP peer using the template parameters as long as the prefix corresponds to a router-id for a node in the TE database. The targeted Hello adjacency will either trigger a new LDP session or will be associated with an existing LDP session to that peer. See section 7.1.2 for more details on the behavior of this feature when an already active targeted Hello adjacency and LDP session exist to the peer.

Up to five (5) peer prefix policies can be associated with a single peer template at all times. Also, the user can associate multiple templates with the same or different peer prefix policies. Thus multiple templates can match with a given peer prefix. In all cases, the targeted session parameters applied to a given peer prefix are taken from the first created template by the user. This provides a more deterministic behavior regardless of the order in which the templates are associated with the prefix policies.

Each time the user executes the above command, with the same or different prefix policy associations, or the user changes a prefix policy associated with a targeted peer template, the system re-evaluates the prefix policy. The outcome of the re-evaluation will tell LDP if an existing

targeted Hello adjacency needs to be torn down or if an existing targeted Hello adjacency needs to have its parameters updated on the fly.

If a /32 prefix is added to (removed from) or if a prefix range is expanded (shrunk) in a prefix list associated with a targeted peer template, the same prefix policy re-evaluation described above is performed.

The template comes up in the **no shutdown** state and as such it takes effect immediately. Once a template is in use, the user can change any of the parameters on the fly without shutting down the template. In this case, all targeted Hello adjacencies are.

There is no overall chassis mode restrictions enforced with the auto-created T-LDP session feature. If the chassis-mode, network chassis-mode or IOM type requirements for an LDP feature are not met, the configuration of the corresponding command will not be allowed as in existing implementation.

Feature Behavior

Whether the prefix list contains one or more specific /32 addresses or a range of addresses, an external trigger is required to indicate to LDP to instantiate a targeted Hello adjacency to a node which address matches an entry in the prefix list. The objective of the feature is to provide an automatic creation of a T-LDP session to the same destination as an auto-created RSVP LSP to achieve automatic tunneling of LDP-over-RSVP. The external trigger is when the router with the matching address appears in the Traffic Engineering database. In the latter case, an external module monitoring the TE database for the peer prefixes provides the trigger to LDP. As a result of this, the user must enable the **traffic-engineering** option in ISIS or OSPF.

Each mapping of a targeted session peer parameter template to a policy prefix which exists in the TE database will result in LDP establishing a targeted Hello adjacency to this peer address using the targeted session parameters configured in the template. This Hello adjacency will then either get associated with an LDP session to the peer if one exists or it will trigger the establishment of a new targeted LDP session to the peer.

The SR OS supports multiple ways of establishing a targeted Hello adjacency to a peer LSR:

- User configuration of the peer with the targeted session parameters inherited from the **config>router>ldp>targeted-session** in the top level context or explicitly configured for this peer in the **config>router>ldp>targeted-session>peer** context and which overrides the top level parameters shared by all targeted peers. Let us refer to the top level configuration context as the global context. Note that some parameters only exist in the global context and as such their value will always be inherited by all targeted peers regardless of which event triggered it.

- User configuration of an SDP of any type to a peer with the **signaling tldp** option enabled (default configuration). In this case the targeted session parameter values are taken from the global context.
- User configuration of a (FEC 129) PW template binding in a BGP-VPLS service. In this case the targeted session parameter values are taken from the global context.
- User configuration of a (FEC 129 type II) PW template binding in a VLL service (dynamic multi-segment PW). In this case the target session parameter values are taken from the global context
- This Release 11.0.R4 user configuration of a mapping of a targeted session peer parameter template to a prefix policy when the peer address exists in the TE database. In this case, the targeted session parameter values are taken from the template.
- Note that features using an LDP LSP, which itself is tunneled over an RSVP LSP (LDP-over-RSVP), as a shortcut do not trigger automatically the creation of the targeted Hello adjacency and LDP session to the destination of the RSVP LSP. The user must configure manually the peer parameters or configure a mapping of a targeted session peer parameter template to a prefix policy. These features are:
 - BGP shortcut (**igp-shortcut ldp** option in BGP),
 - IGP shortcut (**rsvp-shortcut** option in IGP),
 - LDP shortcut for IGP routes (**ldp-shortcut** option in router level),
 - static route LDP shortcut (**ldp** option in a static route),
 - VPRN service (**autobind ldp** option), and

Since the above triggering events can occur simultaneously or in any arbitrary order, the LDP code implements a priority handling mechanism in order to decide which event overrides the active targeted session parameters. The overriding trigger will become the owner of the targeted adjacency to a given peer and will be shown in 'show router ldp peer'.

The table below summarizes the triggering events and the associated priority.

Table 10: Triggering Events and the Associated Priority

Triggering Event	Automatic Creation of Targeted Hello Adjacency	Active Targeted Adjacency Parameter Override Priority
Manual configuration of peer parameters (creator=manual)	Yes	1
Mapping of targeted session template to prefix policy (creator=template)	Yes	2
Manual configuration of SDP with signaling tldp option enabled (creator=service manager)	Yes	3
PW template binding in BGP-AD VPLS (creator=service manager)	Yes	3
PW template binding in FEC 129 VLL (creator=service manager)	Yes	3
LDP-over-RSVP as a BGP/IGP/LDP/Static shortcut	No	N/A
LDP-over-RSVP in VPRN auto-bind	No	N/A
LDP-over-RSVP in BGP Label Route resolution	No	N/A

Triggering Event Automatic Creation of Targeted Hello Adjacency Active Targeted Adjacency Parameter Override Priority

Manual configuration of peer parameters (creator=manual) Yes 1

Mapping of targeted session template to prefix policy (creator=template) Yes 2

Manual configuration of SDP with signaling tldp option enabled (creator=service manager)Yes3

PW template binding in BGP-AD VPLS (creator=service manager)Yes3

PW template binding in FEC 129 VLL (creator=service manager)Yes3

LDP-over-RSVP as a BGP/IGP/ LDP/Static shortcutNoN/A

LDP-over-RSVP in VPRN auto-bindNoN/A

LDP-over-RSVP in BGP Label Route resolutionNoN/A

Table 5 1 Targeted LDP Adjacency Triggering Events and Priority

Note that any parameter value change to an active targeted Hello adjacency caused by any of the above triggering events is performed on the fly by having LDP immediately send a Hello message with the new parameters to the peer without waiting for the next scheduled time for the Hello message. This allows the peer to adjust its local state machine immediately and maintains both the Hello adjacency and the LDP session in UP state. The only exceptions are the following:

- The triggering event caused a change to the **local-lsr-id** parameter value. In this case, the Hello adjacency is brought down which will also cause the LDP session to be brought down if this is the last Hello adjacency associated with the session. A new Hello adjacency and LDP session will then get established to the peer using the new value of the local LSR ID.
- The triggering event caused the targeted peer **shutdown** option to be enabled. In this case, the Hello adjacency is brought down which will also cause the LDP session to be brought down if this is the last Hello adjacency associated with the session.

Finally, the value of any LDP parameter which is specific to the LDP/TCP session to a peer is inherited from the **config>router>ldp>peer-parameters>peer** context. This includes MD5 authentication, LDP prefix per-peer policies, label distribution mode (DU or DOD), etc.

Multicast P2MP LDP for GRT

P2MP LDP LSP setup is initiated by each leaf node of multicast tree. A leaf PE node learns to initiate a multicast tree setup from client application and sends a label map upstream towards the root node of the multicast tree. On propagation of label map, intermediate nodes that are common on path for multiple leaf nodes become branch nodes of the tree.

Figure 43 illustrates wholesale video distribution over P2MP LDP LSP. Static IGMP entries on edge are bound to P2MP LDP LSP tunnel-interface for multicast video traffic distribution.

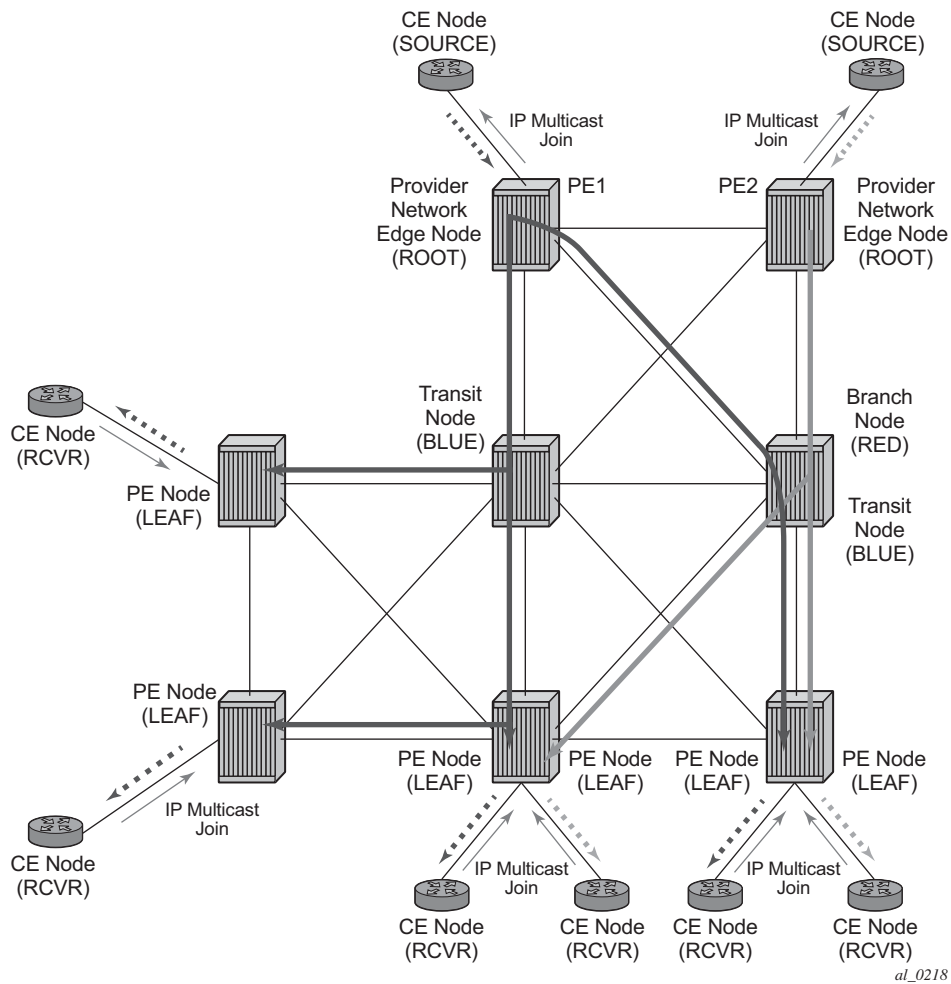


Figure 43: Video Distribution using P2MP LDP

LDP P2MP Support

LDP P2MP Configuration

A node running LDP also supports P2MP LSP setup using LDP. By default, it would advertise the capability to a peer node using P2MP capability TLV in LDP initialization message.

This configuration option per interface is provided to restrict/allow the use of interface in LDP multicast traffic forwarding towards a downstream node. Interface configuration option does not restrict/allow exchange of P2MP FEC by way of established session to the peer on an interface, but it would only restrict/allow use of next-hops over the interface.

LDP P2MP Protocol

Only a single generic identifier range is defined for signaling multipoint tree for all client applications. Implementation on 7x50 SR reserves the range (1..8292) of generic LSP P2MP-ID on root node for static P2MP LSP.

Make Before Break (MBB)

When a transit or leaf node detects that the upstream node towards the root node of multicast tree has changed, it follows graceful procedure that allows make-before-break transition to the new upstream node. Make-before-break support is optional. If the new upstream node does not support MBB procedures then the downstream node waits for the configured timer before switching over to the new upstream node.

ECMP Support

If multiple ECMP paths exist between two adjacent nodes then the upstream node of the multicast receiver programs all entries in forwarding plane. Only one entry is active based on ECMP hashing algorithm.

Multicast LDP Fast Upstream Switchover

This feature allows a downstream LSR of a multicast LDP (mLDP) FEC to perform a fast switchover and source the traffic from another upstream LSR while IGP and LDP are converging due to a failure of the upstream LSR which is the primary next-hop of the root LSR for the P2MP FEC. In essence it provides an upstream Fast-Reroute (FRR) node-protection capability for the mLDP FEC packets. It does it at the expense of traffic duplication from two different upstream nodes into the node which performs the fast upstream switchover.

The detailed procedures for this feature are described in *draft-pdutta-mpls-mldp-up-redundancy*.

Feature Configuration

The user enables the mLDP fast upstream switchover feature by configuring the following option in CLI:

```
configure>router>ldp>mcast-upstream-frr
```

When this command is enabled and LDP is resolving a mLDP FEC received from a downstream LSR, it checks if an ECMP next-hop or a LFA next-hop exist to the root LSR node. If LDP finds one, it programs a primary ILM on the interface corresponding to the primary next-hop and a backup ILM on the interface corresponding to the ECMP or LFA next-hop. LDP then sends the corresponding labels to both upstream LSR nodes. In normal operation, the primary ILM accepts packets while the backup ILM drops them. If the interface or the upstream LSR of the primary ILM goes down causing the LDP session to go down, the backup ILM will then start accepting packets.

In order to make use of the ECMP next-hop, the user must configure the **ecmp** value in the system to at least two (2) using the following command:

```
configure>router>ecmp
```

In order to make use of the LFA next-hop, the user must enable LFA using the following commands:

```
config>router>isis>loopfree-alternate
```

```
config>router>ospf>loopfree-alternate
```

Enabling IP FRR or LDP FRR using the following commands is not strictly required since LDP only needs to know where the alternate next-hop to the root LSR is to be able to send the Label Mapping message to program the backup ILM at the initial signaling of the tree. Thus enabling the LFA option is sufficient. If however, unicast IP and LDP prefixes need to be protected, then these features and the mLDP fast upstream switchover can be enabled concurrently:


```
config>router>ip-fast-reroute
```

```
config>router>ldp>fast-reroute
```

Note that mLdp FRR fast switchover relies on the fast detection of loss of **LDP session** to the upstream peer to which primary ILM label had been advertised. As a result, it is strongly recommended to perform the following:

1. Enable BFD on all LDP interfaces to upstream LSR nodes. When BFD detects the loss of the last adjacency to the upstream LSR, it will bring down immediately the LDP session which will cause the IOM to activate the backup ILM.
 2. If there is a concurrent TLDP adjacency to the same upstream LSR node, enable BFD on the T-LDP peer in addition to enabling it on the interface.
 3. Enable ldp-sync-timer option on all interfaces to the upstream LSR nodes. If an LDP session to the upstream LSR to which the primary ILM is resolved goes down for any other reason than a failure of the interface or of the upstream LSR, routing and LDP will go out of sync. This means the backup ILM will remain activated until the next time SPF is rerun by IGP. By enabling IGP-LDP synchronization feature, the advertised link metric will be changed to max value as soon as the LDP session goes down. This in turn will trigger an SPF and LDP will likely download a new set of primary and backup ILMs.
-

Feature Behavior

This feature allows a downstream LSR to send a label binding to a couple of upstream LSR nodes but only accept traffic from the ILM on the interface to the primary next-hop of the root LSR for the P2MP FEC in normal operation, and accept traffic from the ILM on the interface to the backup next-hop under failure. Obviously, a candidate upstream LSR node must either be an ECMP next-hop or a Loop-Free Alternate (LFA) next-hop. This allows the downstream LSR to perform a fast switchover and source the traffic from another upstream LSR while IGP is converging due to a failure of the LDP session of the upstream peer which is the primary next-hop of the root LSR for the P2MP FEC. In a sense it provides an upstream Fast-Reroute (FRR) node-protection capability for the mLDP FEC packets.

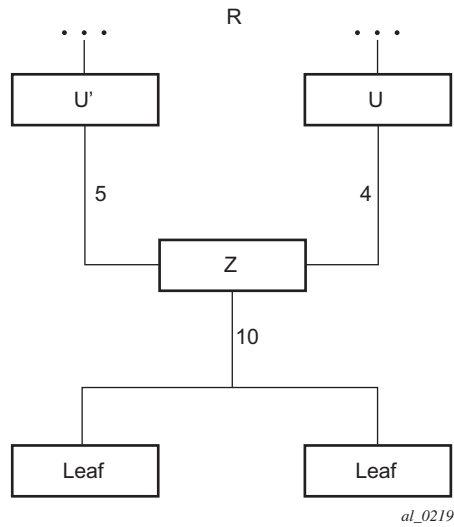


Figure 44: mLDP LSP with Backup Upstream LSR Nodes

Upstream LSR **U** in Figure 44 is the primary next-hop for the root LSR **R** of the P2MP FEC. This is also referred to as primary upstream LSR. Upstream LSR **U'** is an ECMP or LFA backup next-hop for the root LSR **R** of the same P2MP FEC. This is referred to as backup upstream LSR. Downstream LSR **Z** sends a label mapping message to both upstream LSR nodes and programs the primary ILM on the interface to LSR **U** and a backup ILM on the interface to LSR **U'**. The labels for the primary and backup ILMs must be different. LSR **Z** thus will attract traffic from both of them. However, LSR **Z** will block the ILM on the interface to LSR **U'** and will only accept traffic from the ILM on the interface to LSR **U**.

In case of a failure of the link to LSR **U** or of the LSR **U** itself causing the LDP session to LSR **U** to go down, LSR **Z** will detect it and reverse the ILM blocking state and will immediately start receiving traffic from LSR **U'** until IGP converges and provides a new primary next-hop, and ECMP or LFA backup next-hop, which may or may not be on the interface to LSR **U'**. At that point LSR **Z** will update the primary and backup ILMs in the data path.

Note that LDP will use the interface of either an ECMP next-hop or a LFA next-hop to the root LSR prefix, whichever is available, to program the backup ILM. ECMP next-hop and LFA next-hop are however mutually exclusive for a given prefix. IGP installs the ECMP next-hop in preference to an LFA next-hop for a prefix in the Routing Table Manager (RTM).

If one or more ECMP next-hops for the root LSR prefix exist, LDP picks the interface for the primary ILM based on the rules of mLDP FEC resolution specified in RFC 6388:

1. The candidate upstream LSRs are numbered from lower to higher IP address.

2. The following hash is performed: $H = (CRC32(\textit{Opaque Value})) \textit{ modulo } N$, where N is the number of upstream LSRs. The *Opaque Value* is the field identified in the P2MP FEC Element right after 'Opaque Length' field. The 'Opaque Length' indicates the size of the opaque value used in this calculation.
3. The selected upstream LSR U is the LSR that has the number H .

LDP then picks the interface for the backup ILM using the following new rules:

1. if $(H + I < NUM_ECMP)$ {
 - // If the hashed entry is not last in the next-hops then pick up the next as backup.
 - backup = $H + I$;
 - } else {
 - // Wrap around and pickup the first.
 - backup = 1;
 - }

In some topologies, it is possible that none of ECMP or LFA next-hop will be found. In this case, LDP programs the primary ILM only.

Uniform Failover from Primary to Backup ILM

When LDP programs the primary ILM record in the data path, it provides the IOM with the Protect-Group Identifier (PG-ID) associated with this ILM and which identifies which upstream LSR is protected.

In order for the system to perform a fast switchover to the backup ILM in the fast path, LDP applies to the primary ILM uniform FRR failover procedures similar in concept to the ones applied to an NHLFE in the existing implementation of LDP FRR for unicast FECs. There are however important differences to note. LDP associates a unique Protect Group ID (PG-ID) to all mLDP FECs which have their primary ILM on any LDP interface pointing **to the same upstream LSR**. This PG-ID is assigned per upstream LSR regardless of the number of LDP interfaces configured to this LSR. As such this PG-ID is different from the one associated with unicast FECs and which is assigned to each downstream LDP interface and next-hop. If however a failure caused an interface to go down and also caused the LDP session to upstream peer to go down, both PG-IDs have their state updated in the IOM and thus the uniform FRR procedures will be triggered for both the unicast LDP FECs forwarding packets towards the upstream LSR and the mLDP FECs receiving packets from the same upstream LSR.

When the mLDP FEC is programmed in the data path, the primary and backup ILM record thus contain the PG-ID the FEC is associated with. The IOM also maintains a list of PG-IDs and a state bit which indicates if it is UP or DOWN. When the PG-ID state is UP the primary ILM for each mLDP FEC is open and will accept mLDP packets while the backup ILM is blocked and drops mLDP packets. LDP sends a PG-ID DOWN notification to IOM when it detects that the LDP session to the peer is gone down. This notification will cause the backup ILMs associated with this PG-ID to open and accept mLDP packets immediately. When IGP re-converges, an updated pair of primary and backup ILMs is downloaded for each mLDP FEC by LDP into the IOM with the corresponding PG-IDs.

Note that if multiple LDP interfaces exist to the upstream LSR, a failure of one interface will bring down the link Hello adjacency on that interface but not the LDP session which is still associated with the remaining link Hello adjacencies. In this case, the upstream LSR updates in IOM the NHLFE for the mLDP FEC to use one of the remaining links. The switchover time in this case is not managed by the uniform failover procedures.

Multi-Area and Multi-Instance Extensions to LDP

In order to extend LDP across multiple areas of an IGP instance or across multiple IGP instances, the current standard LDP implementation based on RFC 3036 requires that all /32 prefixes of PEs be leaked between the areas or instances. This is because an exact match of the prefix in the routing table is required to install the prefix binding in the LDP Forwarding Information Base (FIB). Although a router will do this by default when configured as Area Border Router (ABR), this increases the convergence of IGP on routers when the number of PE nodes scales to thousands of nodes.

Multi-area and multi-instance extensions to LDP provide an optional behavior by which LDP installs a prefix binding in the LDP FIB by simply performing a longest prefix match with an aggregate prefix in the routing table (RIB). That way, the ABR will be configured to summarize the /32 prefixes of PE routers. This method is compliant to RFC 5283, *LDP Extension for Inter-Area Label Switched Paths (LSPs)*.

LDP Shortcut for BGP Next-Hop Resolution

LDP shortcut for BGP next-hop resolution shortcuts allow for the deployment of a 'route-less core' infrastructure. Many service providers either have or intend to remove the IBGP mesh from their network core, retaining only the mesh between routers connected to areas of the network that require routing to external routes.

Shortcuts are implemented by utilizing Layer 2 tunnels (i.e., MPLS LSPs) as next hops for prefixes that are associated with the far end termination of the tunnel. By tunneling through the network core, the core routers forwarding the tunnel have no need to obtain external routing information and are immune to attack from external sources.

The tunnel table contains all available tunnels indexed by remote destination IP address. LSPs derived from received LDP /32 route FECs will automatically be installed in the table associated with the advertising router-ID when IGP shortcuts are enabled.

Evaluating tunnel preference is based on the following order in descending priority:

1. LDP /32 route FEC shortcut
2. Actual IGP next-hop

If a higher priority shortcut is not available or is not configured, a lower priority shortcut is evaluated. When no shortcuts are configured or available, the IGP next-hop is always used. Shortcut and next-hop determination is event driven based on dynamic changes in the tunneling mechanisms and routing states.

Refer to the 7750 SR OS Routing Protocols Guide for details on the use of LDP FEC and RSVP LSP for BGP Next-Hop Resolution.

LDP Shortcut for IGP Routes

The LDP shortcut for IGP route resolution feature allows forwarding of packets to IGP learned routes using an LDP LSP. When LDP shortcut is enabled globally, IP packets forwarded over a network IP interface will be labeled with the label received from the next-hop for the route and corresponding to the FEC-prefix matching the destination address of the IP packet. In such a case, the routing table will have the shortcut next-hop as the best route. If such a LDP FEC does not exist, then the routing table will have the regular IP next-hop and regular IP forwarding will be performed on the packet.

An egress LER advertises and maintains a FEC, label binding for each IGP learned route. This is performed by the existing LDP fec-originate capability.

LDP Shortcut Configuration

The user enables the use of LDP shortcut for resolving IGP routes by entering the global command **config>router>ldp-shortcut**.

This command enables forwarding of user IP packets and specified control IP packets using LDP shortcuts over all network interfaces in the system which participate in the IS-IS and OSPF routing protocols. The default is to disable the LDP shortcut across all interfaces in the system.

IGP Route Resolution

When LDP shortcut is enabled, LDP populates the RTM with next-hop entries corresponding to all prefixes for which it activated an LDP FEC. For a given prefix, two route entries are populated in RTM. One corresponds to the LDP shortcut next-hop and has an owner of LDP. The other one is the regular IP next-hop. The LDP shortcut next-hop always has preference over the regular IP next-hop for forwarding user packets and specified control packets over a given outgoing interface to the route next-hop.

The prior activation of the FEC by LDP is done by performing an exact match with an IGP route prefix in RTM. It can also be done by performing a longest prefix-match with an IGP route in RTM if the aggregate-prefix-match option is enabled globally in LDP.

This feature is not restricted to /32 FEC prefixes. However only /32 FEC prefixes will be populated in the CPM Tunnel Table for use as a tunnel by services.

All user packets and specified control packets for which the longest prefix match in RTM yields the FEC prefix will be forwarded over the LDP LSP. Currently, the control packets that could be forwarded over the LDP LSP are ICMP ping and UDP-traceroute. The following is an example of the resolution process.

Assume the egress LER advertised a FEC for some /24 prefix using the fec-originate command. At the ingress LER, LDP resolves the FEC by checking in RTM that an exact match exists for this prefix. Once LDP activated the FEC, it programs the NHLFE in the egress data path and the LDP tunnel information in the ingress data path tunnel table.

Next, LDP provides the shortcut route to RTM which will associate it with the same /24 prefix. There will be two entries for this /24 prefix, the LDP shortcut next-hop and the regular IP next-hop. The latter was used by LDP to validate and activate the FEC. RTM then resolves all user prefixes which succeed a longest prefix match against the /24 route entry to use the LDP LSP.

Assume now the aggregate-prefix-match was enabled and that LDP found a /16 prefix in RTM to activate the FEC for the /24 FEC prefix. In this case, RTM adds a new more specific route entry of /24 and has the next-hop as the LDP LSP but it will still not have a specific /24 IP route entry. RTM then resolves all user prefixes which succeed a longest prefix match against the /24 route entry to use the LDP LSP while all other prefixes which succeed a longest prefix-match against the /16 route entry will use the IP next-hop.

LDP Shortcut Forwarding Plane

Once LDP activated a FEC for a given prefix and programmed RTM, it also programs the ingress Tunnel Table in forwarding engine with the LDP tunnel information.

When an IPv4 packet is received on an ingress network interfacea subscriber IES interface,, or a regular IES interface, the lookup of the packet by the ingress forwarding engine will result in the packet being sent labeled with the label stack corresponding to the NHLFE of the LDP LSP when the preferred RTM entry corresponds to an LDP shortcut.

If the preferred RTM entry corresponds to an IP next-hop, the IPv4 packet is forwarded unlabeled.

ECMP Considerations

When ECMP is enabled and multiple equal-cost next-hops exit for the IGP route, the ingress forwarding engine sprays the packets for this route based on hashing routine currently supported for IPv4 packets.

When the preferred RTM entry corresponds to an LDP shortcut route, spraying will be performed across the multiple next-hops for the LDP FEC. The FEC next-hops can either be direct link LDP neighbors or T-LDP neighbors reachable over RSVP LSPs in the case of LDP-over-RSVP but not both. This is as per ECMP for LDP in existing implementation.

When the preferred RTM entry corresponds to a regular IP route, spraying will be performed across regular IP next-hops for the prefix.

Disabling TTL Propagation in an LSP Shortcut

This feature provides the option for disabling TTL propagation from a transit or a locally generated IP packet header into the LSP label stack when an LDP LSP is used as a shortcut for BGP next-hop resolution, a static-route next-hop resolution, or for an IGP route resolution.

A transit packet is a packet received from an IP interface and forwarded over the LSP shortcut at ingress LER.

A locally-generated IP packet is any control plane packet generated from the CPM and forwarded over the LSP shortcut at ingress LER.

TTL handling can be configured for all LDP LSP shortcuts originating on an ingress LER using the following global commands:

```
config>router>ldp>[no] shortcut-transit-ttl-propagate  
config>router>ldp>[no] shortcut-local-ttl-propagate
```

These commands apply to all LDP LSPs which are used to resolve static routes, BGP routes, and IGP routes.

When the **no** form of the above command is enabled for local packets, TTL propagation is disabled on all locally generated IP packets, including ICMP Ping, traceroute, and OAM packets that are destined to a route that is resolved to the LSP shortcut. In this case, a TTL of 255 is programmed onto the pushed label stack. This is referred to as pipe mode.

Similarly, when the **no** form is enabled for transit packets, TTL propagation is disabled on all IP packets received on any IES interface and destined to a route that is resolved to the LSP shortcut. In this case, a TTL of 255 is programmed onto the pushed label stack.

LDP Graceful Handling of Resource Exhaustion

This feature enhances the behavior of LDP when a data path or a CPM resource required for the resolution of a FEC is exhausted. In prior releases, the LDP module shuts down. The user is required to fix the issue causing the FEC scaling to be exceeded and to restart the LDP module by executing the **unshut** command.

LDP Base Graceful Handling of Resources

This feature implements a base graceful handling capability by which the LDP interface to the peer, or the targeted peer in the case of Targeted LDP (T-LDP) session, is shutdown. If LDP tries to resolve a FEC over a link or a targeted LDP session and it runs out of data path or CPM resources, it will bring down that interface or targeted peer which will bring down the Hello adjacency over that interface to all link LDP peers or to the targeted peer. The interface is brought down in LDP context only and is still available to other applications such as IP forwarding and RSVP LSP forwarding.

Depending of what type of resource was exhausted, the scope of the action taken by LDP will be different. Some resource such as NHLFE have interface local impact, meaning that only the interface to the downstream LSR which advertised the label is shutdown. Some resources such as ILM have global impact, meaning that they will impact every downstream peer or targeted peer which advertised the FEC to the node. The following are examples to illustrate this.

- For NHLFE exhaustion, one or more interfaces or targeted peers, if the FEC is ECMP, will be shut down. ILM is maintained as long as there is at least one downstream for the FEC for which the NHLFE has been successfully programmed.
- For an exhaustion of an ILM for a unicast LDP FEC, all interfaces to peers or all target peers which sent the FEC will be shutdown. No deprogramming of data path is required since FEC is not programmed.
- An exhaustion of ILM for an mLDP FEC can happen during primary ILM programming, MBB ILM programming, or multicast upstream FRR backup ILM programming. In all cases, the P2MP index for the mLDP tree is deprogrammed and the interfaces to each downstream peer which sent a Label Mapping message associated with this ILM are shutdown.

After the user has taken action to free resources up, he/she will require manually unshut the interface or the targeted peer to bring it back into operation. This then re-establishes the Hello adjacency and resumes the resolution of FECs over the interface or to the targeted peer.

Detailed guidelines for using the feature and for troubleshooting a system which activated this feature are provided in the following sections.

This new behavior will become the new default behavior in Release 11.0.R4 and will interoperate with SROS based LDP implementation and any other third party LDP implementation.

The following are the data path resources which can trigger this mechanism:

- NHLFE, ILM, Label-to-NHLFE (LTN), Tunnel Index, P2MP Index.

The following are the CPM resources which can trigger this mechanism:

- Label allocation.

LDP Enhanced Graceful Handling of Resources

This feature is an enhanced graceful handling capability which is supported only among SROS based implementations. If LDP tries to resolve a FEC over a link or a targeted session and it runs out of data path or CPM resources, it will put the LDP/T-LDP session into overload state. As a result, it will release to its LDP peer the labels of the FECs which it could not resolve and will also send an LDP notification message to all LDP peers with the new status load of overload for the FEC type which caused the overload. The notification of overload is per FEC type, i.e., unicast IPv4, P2MP mLDP etc., and not per individual FEC. The peer which caused the overload and all other peers will stop sending any new FECs of that type until this node updates the notification stating that it is no longer in overload state for that FEC type. FECs of this type previously resolved and other FEC types to this peer and all other peers will continue to forward traffic normally.

After the user has taken action to free resources up, he/she will require manually clear the overload state of the LDP/T-LDP sessions towards its peers. Detailed guidelines for using the feature and for troubleshooting a system which activated this feature are provided in Section 7.3.

The enhanced mechanism will be enabled instead of the base mechanism only if both LSR nodes advertize this new LDP capability at the time the LDP session is initialized. Otherwise, they will continue to use the base mechanism.

This feature will operate among SROS LSR nodes using a couple of private vendor LDP capabilities:

- The first one is the LSR Overload Status TLV to signal or clear the overload condition.
- The second one is the Overload Protection Capability Parameter which allows LDP peers to negotiate the use or not of the overload notification feature and hence the enhanced graceful handling mechanism.

When interoperating with an LDP peer which does not support the enhanced resource handling mechanism, the 7x50 reverts automatically to the default base resource handling mechanism.

The following are the details of the mechanism.

LSR Overload Notification

When an upstream LSR is overloaded for a FEC type, it notifies one or more downstream peer LSRs that it is overloaded for the FEC type.

When a downstream LSR receives overload status ON notification from an upstream LSR, it does not send further label mappings for the specified FEC type. When a downstream LSR receives overload OFF notification from an upstream LSR, it sends pending label mappings to the upstream LSR for the specified FEC type.

This feature introduces a new TLV referred to as *LSR Overload Status TLV*. This TLV is encoded using vendor proprietary TLV encoding as per RFC 5036. It uses a TLV type value of 0x3E02 and the Timetra OUI value of 0003FA.

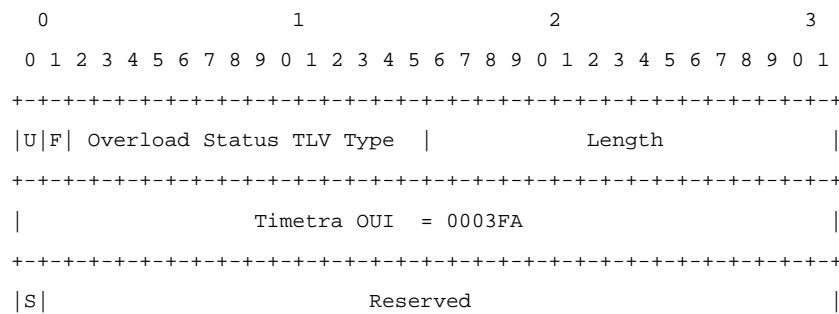


Figure 45: LSR Overload Status TLV (Type = 0x3E02)

where:

U-bit: Unknown TLV bit, as described in RFC 5036. The value MUST be 1 which means if unknown to receiver then receiver should ignore

F-bit: Forward unknown TLV bit, as described in RFC RFC5036. The value of this bit MUST be 1 since a LSR overload TLV is sent only between two immediate LDP peers, which are not forwarded.

S-bit: The State Bit. It indicates whether the sender is setting the LSR Overload Status ON or OFF. The State Bit value is used as follows:

- 1 - The TLV is indicating LSR overload status as ON.
- 0 - The TLV is indicating LSR overload status as OFF.

When a LSR that implements the procedures defined in this document generates LSR overload status, it MUST send LSR Overload Status TLV in a LDP Notification Message accompanied by a FEC TLV. The FEC TLV must contain one Typed Wildcard FEC TLV that specifies the FEC type to which the overload status notification applies.

The feature in this document re-uses the Typed Wildcard FEC Element which is defined in RFC 5918.

LSR Overload Protection Capability

To ensure backward compatibility with procedures in RFC 5036 an LSR supporting Overload Protection need means to determine whether a peering LSR supports overload protection or not.

An LDP speaker that supports the LSR Overload Protection procedures as defined in this document MUST inform its peers of the support by including a LSR Overload Protection Capability Parameter in its initialization message. The Capability parameter follows the guidelines and all Capability Negotiation Procedures as defined in RFC 5561. This TLV is encoded using vendor proprietary TLV encoding as per RFC 5036. It uses a TLV type value of 0x3E03 and the Timetra OUI value of 0003FA.




Figure 46: LSR Overload Protection Capability TLV (Type== 0x3E03)

Where:

U and F bits : MUST be 1 and 0 respectively as per section 3 of LDP Capabilities [RFC5561].

S-bit : MUST be 1 (indicates that capability is being advertised).

Procedures for LSR overload protection

The procedures defined in this document apply only to LSRs that support Downstream Unsolicited (DU) label advertisement mode and Liberal Label Retention Mode. An LSR that implements the LSR overload protection follows the following procedures:

1. A LSR MUST NOT use LSR Overload notification procedures with a peer LSR that has not specified LSR Overload Protection Capability in Initialization Message received from the peer LSR.
2. When an upstream LSR detects that it is overloaded with a FEC type then it MUST initiate a LDP Notification Message with S bit ON in LSR Overload Status TLV and a FEC TLV containing the Typed Wildcard FEC Element for the specified FEC type. The Message may be sent to one or more peers.
3. After it has notified overload status ON for a FEC type, the overloaded upstream LSR MAY send Label Release for a set of FEC elements to respective downstream LSRs to off load its LIB below certain watermark.
4. When an upstream LSR that was overloaded for a FEC type before, detects that it is no longer overloaded then it MUST send a LDP Notification Message with S bit OFF in LSR Overload Status TLV and FEC TLV containing the Typed Wildcard FEC Element for the specified FEC type.
5. When an upstream LSR has notified as overloaded for a FEC type, then a downstream LSR MUST NOT send new Label Mappings for the specified FEC type to the upstream LSR.
6. When a downstream LSR receives LSR Overload Notification from a peering LSR with status OFF for a FEC type then the receiving LSR MUST send any label mappings for the FEC type which were pending to the upstream LSR or which are eligible to be sent now.
7. When an upstream LSR is overloaded for a FEC type and it receives Label Mapping for that FEC type from a downstream LSR then it MAY send Label Release to the downstream for the received Label Mapping with LDP Status Code as *No_Label_Resources* as defined in RFC 5036.

User Guidelines and Troubleshooting Procedures

Common Procedures

When troubleshooting a LDP resource exhaustion situation on an LSR, the user must first determine which of the LSR and its peers supports the enhanced handling of resources. This is done by checking if the local LSR or its peers advertised the LSR Overload Protection Capability:

```

show router ldp status
=====
LDP Status for LSR ID 110.20.1.110
=====
Admin State      : Up                Oper State      : Up
Created at      : 07/17/13 21:27:41  Up Time        : 0d 01:00:41
Oper Down Reason : n/a                Oper Down Events : 1
Last Change     : 07/17/13 21:27:41  Tunn Down Damp Time : 20 sec
Label Withdraw Del*: 0 sec            Implicit Null Label : Enabled
Short. TTL Prop Lo*: Enabled          Short. TTL Prop Tran*: Enabled
Import Policies :                      Export Policies  :
    Import-LDP                          Import-LDP
    External                             External
Tunl Exp Policies :
    from-proto-bgp
Aggregate Prefix : False              Agg Prefix Policies : None
FRR              : Enabled            Mcast Upstream FRR  : Disabled
Dynamic Capability : False            P2MP Capability     : True
MP MBB Capability : True              MP MBB Time         : 10
Overload Capability: True <---- //Local Overload Capability
Active Adjacencies : 0                Active Sessions     : 0
Active Interfaces  : 2                Inactive Interfaces : 4
Active Peers       : 62               Inactive Peers      : 10
Addr FECs Sent     : 0                Addr FECs Recv      : 0
Serv FECs Sent     : 0                Serv FECs Recv      : 0
P2MP FECs Sent     : 0                P2MP FECs Recv     : 0
Attempted Sessions : 458
No Hello Err       : 0                Param Adv Err       : 0
Max PDU Err        : 0                Label Range Err     : 0
Bad LDP Id Err     : 0                Bad PDU Len Err    : 0
Bad Mesg Len Err   : 0                Bad TLV Len Err     : 0
Unknown TLV Err    : 0
Malformed TLV Err  : 0                Keepalive Expired Err: 4
Shutdown Notif Sent: 12              Shutdown Notif Recv : 5
=====

show router ldp session detail
=====
LDP Sessions (Detail)
=====
-----
Session with Peer 10.8.100.15:0, Local 110.20.1.110:0
-----
Adjacency Type      : Targeted          State              : Nonexistent
Up Time             : 0d 00:00:00
Max PDU Length      : 4096              KA/Hold Time Remaining : 0

```

```

Link Adjacencies      : 0                Targeted Adjacencies : 1
Local Address         : 110.20.1.110     Peer Address         : 10.8.100.15
Local TCP Port        : 0                Peer TCP Port        : 0
Local KA Timeout      : 40               Peer KA Timeout      : 40
Mesg Sent             : 0                Mesg Recv           : 1
FECs Sent            : 0                FECs Recv           : 0
Addrs Sent           : 0                Addrs Recv          : 0
GR State              : Capable          Label Distribution   : DU
Nbr Liveness Time    : 0                Max Recovery Time   : 0
Number of Restart     : 0                Last Restart Time   : Never
P2MP                  : Not Capable      MP MBB              : Not Capable
Dynamic Capability    : Not Capable      LSR Overload        : Not Capable <---- /
/Peer OverLoad Capab.
Advertise             : Address/Servi*
Addr FEC OverLoad Sent : No              Addr FEC OverLoad Recv : No
Mcast FEC Overload Sent : No             Mcast FEC Overload Recv : No
Serv FEC Overload Sent : No              Serv FEC Overload Recv : No
-----

```

Base Resource Handling Procedures

Step 1

If the peer OR the local LSR does not support the Overload Protection Capability it means that the associated adjacency [interface/peer] will be brought down as part of the base resource handling mechanism.

The user can determine which interface or targeted peer was shut down, by applying the following commands:

- [show router ldp interface resource-failures]
- [show router ldp peer resource-failures]

```

show router ldp interface resource-failures
=====
LDP Interface Resource Failures
=====
srl                               srr
sru4                              sr4-1-5-1
=====

show router ldp peer resource-failures
=====
LDP Peers Resource Failures
=====
10.20.1.22                        110.20.1.3
=====

```

A trap is also generated for each interface or targeted peer:

```
16 2013/07/17 14:21:38.06 PST MINOR: LDP #2003 Base LDP Interface Admin State
```



```
"Interface instance state changed - vRtrID: 1, Interface sr4-1-5-1, administrative state: inService, operational state: outOfService"
```

```
13 2013/07/17 14:15:24.64 PST MINOR: LDP #2003 Base LDP Interface Admin State
"Interface instance state changed - vRtrID: 1, Peer 10.20.1.22, administrative state: inService, operational state: outOfService"
```

The user can then check that the base resource handling mechanism has been applied to a specific interface or peer by running the following show commands:

- [show router ldp interface detail]

- [show router ldp peer detail]

```
show router ldp interface detail
=====
LDP Interfaces (Detail)
=====
-----
Interface "sr4-1-5-1"
-----
Admin State      : Up                Oper State      : Down
Oper Down Reason : noResources <----- //link LDP resource exhaustion handled
Hold Time       : 45                Hello Factor    : 3
Oper Hold Time  : 45
Hello Reduction : Disabled          Hello Reduction *: 3
Keepalive Timeout : 30            Keepalive Factor : 3
Transport Addr  : System          Last Modified   : 07/17/13 14:21:38
Active Adjacencies : 0
Tunneling       : Disabled
Lsp Name        : None
Local LSR Type  : System
Local LSR       : None
BFD Status      : Disabled
Multicast Traffic : Enabled
-----
```

```
show router ldp discovery interface "sr4-1-5-1" detail
=====
LDP Hello Adjacencies (Detail)
=====
-----
Interface "sr4-1-5-1"
-----
Local Address    : 223.0.2.110      Peer Address    : 224.0.0.2
Adjacency Type  : Link           State           : Down
=====
```

```
show router ldp peer detail
=====
LDP Peers (Detail)
=====
-----
Peer 10.20.1.22
-----
```

```

Admin State      : Up                Oper State       : Down
Oper Down Reason : noResources      <----- // T-LDP resource exhaustion handled
Hold Time       : 45                Hello Factor     : 3
Oper Hold Time  : 45
Hello Reduction : Disabled          Hello Reduction Fact*: 3
Keepalive Timeout : 40              Keepalive Factor : 4
Passive Mode    : Disabled          Last Modified    : 07/17/13 14:15:24
Active Adjacencies : 0              Auto Created     : No
Tunneling      : Enabled
Lsp Name       : None
Local LSR      : None
BFD Status     : Disabled
Multicast Traffic : Disabled

```

```

-----
show router ldp discovery peer 10.20.1.22 detail
=====
LDP Hello Adjacencies (Detail)
=====
Peer 10.20.1.22
-----
Local Address      : 110.20.1.110      Peer Address      : 10.20.1.22
Adjacency Type    : Targeted          State             : Down <----- //T-LDP
resource exhaustion handled
=====

```

Step 2

Besides interfaces and targeted peer, locally originated FECs may also be put into overload. These are the following:

- unicast fec-originate pop
- multicast local static p2mp-fec type=1 [on leaf LSR]
- multicast local Dynamic p2mp-fec type=3 [on leaf LSR]

The user can check if only remote and/or local FECs have been set in overload by the resource base resource exhaustion mechanism using the following command:

- [tools dump router ldp instance]

The relevant part of the output is described below:

```

{..... snip.....}
Num OLoad Interfaces: 4 <----- // #LDP interfaces resource in exhaustion
Num Targ Sessions: 72 Num Active Targ Sess: 62
Num OLoad Targ Sessions: 7 <----- // #T-LDP peers in resource exhaustion
Num Addr FECs Rcvd: 0 Num Addr FECs Sent: 0
Num Addr Fecs OLoad: 1 <----- // # of local/remote unicast FECs in Overload
Num Svc FECs Rcvd: 0 Num Svc FECs Sent: 0
Num Svc FECs OLoad: 0 <----- // # of local/remote service Fecs in Overload
Num mcast FECs Rcvd: 0 Num Mcast FECs Sent: 0
Num mcast FECs OLoad: 0 <----- // # of local/remote multicast Fecs in Over-

```

```
load
{..... snip.....}
```

When at least one local FEC has been set in overload the following trap will occur:

```
23 2013/07/17 15:35:47.84 PST MINOR: LDP #2002 Base LDP Resources Exhausted "Instance
state changed - vRtrID: 1, administrative state: inService, operationa l state: inService"
```

Step 3

After the user has detected that at least, one link LDP or T-LDP adjacency has been brought down by the resource exhaustion mechanism, he/she must protect the router by applying one or more of the following to free resources up:

- Identify the source for the [unicast/multicast/service] FEC flooding.
- Configure the appropriate [import/export] policies and/or delete the excess [unicast/multicast/service] FECs not currently handled.

Step 4

Next, the user has to manually attempt to clear the overload (no resource) state and allow the router to attempt to restore the link and targeted sessions to its peer.

Please note that due to the dynamic nature of FEC distribution and resolution by LSR nodes, one cannot predict exactly which FECs and which interfaces or targeted peers will be restored after performing the following commands if the LSR activates resource exhaustion again.

One of the following commands can be used:

- [clear router ldp resource-failures]

- Clears the overload state and attempt to restore adjacency and session for LDP interfaces and peers.
- Clear the overload state for the local FECs.

- [clear router ldp interface ifName]

- [clear router ldp peer peerAddress]

- Clears the overload state and attempt to restore adjacency and session for LDP interfaces and peers.
- These 2 commands ***DO NOT*** Clear the overload state for the local FECs.

Enhanced Resource Handling Procedures

Step 1

If the peer AND the local LSR do support the Overload Protection Capability it means that the LSR will signal the overload state for the FEC type which caused the resource exhaustion as part of the enhanced resource handling mechanism.

In order to verify if the local router has received or sent the overload status TLV, perform the following:

```
- [show router ldp session detail]
  show router ldp session 110.20.1.1 detail
-----
Session with Peer 110.20.1.1:0, Local 110.20.1.110:0
-----
Adjacency Type      : Both          State          : Established
Up Time             : 0d 00:05:48
Max PDU Length      : 4096          KA/Hold Time Remaining : 24
Link Adjacencies    : 1            Targeted Adjacencies  : 1
Local Address       : 110.20.1.110 Peer Address       : 110.20.1.1
Local TCP Port      : 51063        Peer TCP Port       : 646
Local KA Timeout    : 30          Peer KA Timeout     : 45
Msg Sent            : 442         Msg Recv           : 2984
FECs Sent           : 16         FECs Recv          : 2559
Addrs Sent          : 17         Addrs Recv         : 1054
GR State            : Capable      Label Distribution   : DU
Nbr Liveness Time   : 0          Max Recovery Time   : 0
Number of Restart   : 0          Last Restart Time   : Never
P2MP                : Capable     MP MBB             : Capable
Dynamic Capability  : Not Capable LSR Overload       : Capable
Advertise           : Address/Servi* BFD Operational Status : inService
Addr FEC OverLoad Sent : Yes      Addr FEC OverLoad Recv : No <---- // this
LSR sent overLoad for unicast FEC type to peer
Mcast FEC Overload Sent: No      Mcast FEC Overload Recv: No
Serv FEC Overload Sent : No      Serv FEC Overload Recv : No
-----

  show router ldp session 110.20.1.110 detail
-----
Session with Peer 110.20.1.110:0, Local 110.20.1.1:0
-----
Adjacency Type      : Both          State          : Established
Up Time             : 0d 00:08:23
Max PDU Length      : 4096          KA/Hold Time Remaining : 21
Link Adjacencies    : 1            Targeted Adjacencies  : 1
Local Address       : 110.20.1.1 Peer Address       : 110.20.1.110
Local TCP Port      : 646         Peer TCP Port       : 51063
Local KA Timeout    : 45          Peer KA Timeout     : 30
Msg Sent            : 3020        Msg Recv           : 480
FECs Sent           : 2867        FECs Recv          : 16
Addrs Sent          : 1054        Addrs Recv         : 17
GR State            : Capable      Label Distribution   : DU
Nbr Liveness Time   : 0          Max Recovery Time   : 0
Number of Restart   : 0          Last Restart Time   : Never
P2MP                : Capable     MP MBB             : Capable
Dynamic Capability  : Not Capable LSR Overload       : Capable
Advertise           : Address/Servi* BFD Operational Status : inService
Addr FEC OverLoad Sent : No      Addr FEC OverLoad Recv : Yes <---- // this
LSR received overLoad for unicast FEC type from peer
Mcast FEC Overload Sent: No      Mcast FEC Overload Recv: No
```

Serv FEC Overload Sent : No Serv FEC Overload Recv : No
=====

A trap is also generated:

70002 2013/07/17 16:06:59.46 PST MINOR: LDP #2008 Base LDP Session State Change "Session state is operational. Overload Notification message is sent to/from peer 110.20.1.1:0 with overload state true for fec type prefixes"

Step 2

Besides interfaces and targeted peer, locally originated FECs may also be put into overload. These are the following:

- unicast fec-originate pop
- multicast local static p2mp-fec type=1 [on leaf LSR]
- multicast local Dynamic p2mp-fec type=3 [on leaf LSR]

The user can check if only remote and/or local FECs have been set in overload by the resource enhanced resource exhaustion mechanism using the following command:

- [tools dump router ldp instance]

The relevant part of the output is described below:

Num Entities OLoad (FEC: Address Prefix): Sent: 7 Rcvd: 0 <----- // #
of session in OvLd for fec-type=unicast
Num Entities OLoad (FEC: PWE3): Sent: 0 Rcvd: 0 <----- // #
of session in OvLd for fec-type=service
Num Entities OLoad (FEC: GENPWE3): Sent: 0 Rcvd: 0 <----- // #
of session in OvLd for fec-type=service
Num Entities OLoad (FEC: P2MP): Sent: 0 Rcvd: 0 <----- // #
of session in OvLd for fec-type=MulticastP2mp
Num Entities OLoad (FEC: MP2MP UP): Sent: 0 Rcvd: 0 <----- // #
of session in OvLd for fec-type=MulticastMP2mp
Num Entities OLoad (FEC: MP2MP DOWN): Sent: 0 Rcvd: 0 <----- // #
of session in OvLd for fec-type=MulticastMP2mp
Num Active Adjacencies: 9
Num Interfaces: 6 Num Active Interfaces: 6
Num OLoad Interfaces: 0 <----- // link LDP interfaces in resource exhaus-
tion should be zero when Overload Protection Capability is supported
Num Targ Sessions: 72 Num Active Targ Sess: 67
Num OLoad Targ Sessions: 0 <----- // T-LDP peers in resource exhaustion should
be zero if Overload Protection Capability is supported
Num Addr FECs Rcvd: 8667 Num Addr FECs Sent: 91
Num Addr Fecs OLoad: 1 <----- // # of local/
remote unicast Fecs in Overload
Num Svc FECs Rcvd: 3111 Num Svc FECs Sent: 0
Num Svc FECs OLoad: 0 <----- // # of local/
remote service Fecs in Overload
Num mcast FECs Rcvd: 0 Num Mcast FECs Sent: 0
Num mcast FECs OLoad: 0 <----- // # of local/
remote multicast Fecs in Overload

Num MAC Flush Rcvd: 0 Num MAC Flush Sent: 0

When at least one local FEC has been set in overload the following trap will occur:

```
69999 2013/07/17 16:06:59.21 PST MINOR: LDP #2002 Base LDP Resources Exhausted "Instance
state changed - vRtrID: 1, administrative state: inService, operational state: inService"
```

Step 3

After the user has detected that at least one overload status TLV has been sent or received by the LSR, he/she must protect the router by applying one or more of the following to free resources up:

- Identify the source for the [unicast/multicast/service] FEC flooding. This is most likely the LSRs which session received the overload status TLV.
- Configure the appropriate [import/export] policies and/or delete the excess [unicast/multicast/service] FECs from the FEC type in overload.

Step 4

Next, the user has to manually attempt to clear the overload state on the affected sessions and for the affected FEC types and allow the router to clear the overload status TLV to its peers.

Please note that due to the dynamic nature of FEC distribution and resolution by LSR nodes, one cannot predict exactly which sessions and which FECs will be cleared after performing the following commands if the LSR activates overload again.

One of the following commands can be used depending if the user wants to clear all sessions or at once or one session at a time:

- [clear router ldp resource-failures]

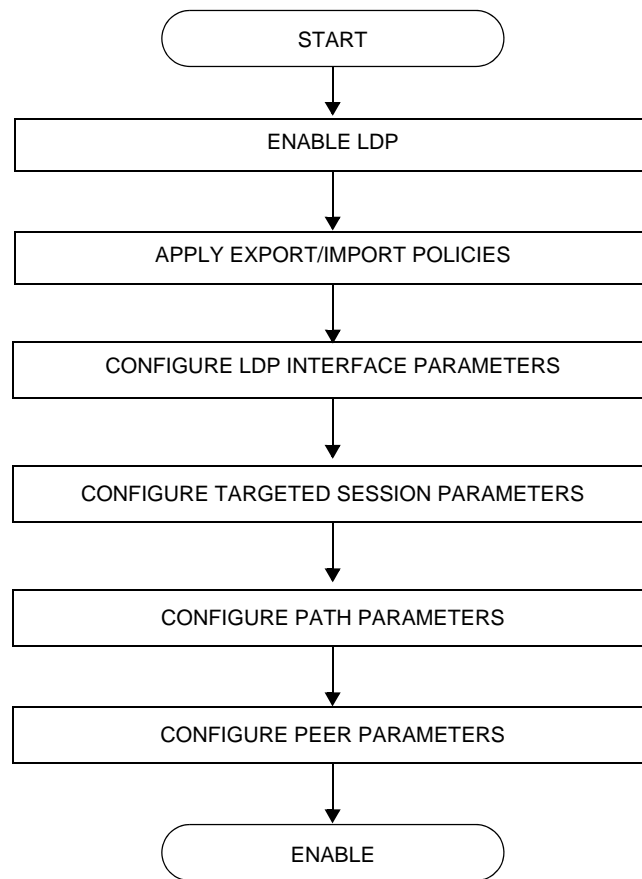
- Clears the overload state for the affected sessions and FEC types.
- Clear the overload state for the local FECs.

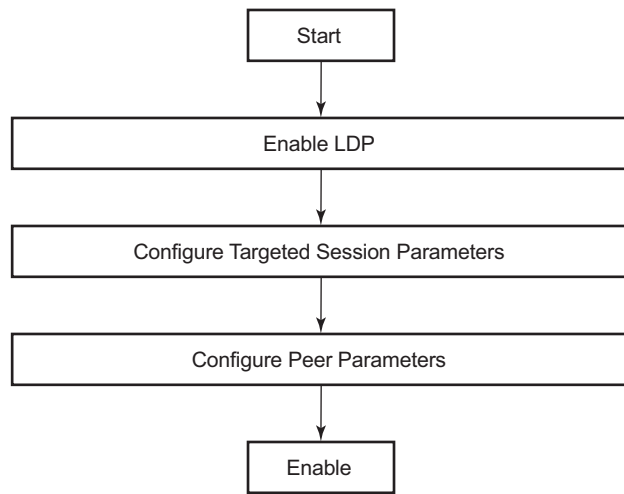
- [clear router ldp session a.b.c.d overload fec-type {services|prefixes|multicast}]

- Clears the overload state for the specified session and FEC type.
- Clears the overload state for the local FECs.

LDP Process Overview

Figure 47 displays the process to provision basic LDP parameters.





al_0220

Figure 47: LDP Configuration and Implementation

Configuring LDP with CLI

This section provides information to configure LDP using the command line interface.

Topics in this section include:

- [LDP Configuration Overview on page 506](#)
- [Basic LDP Configuration on page 507](#)
- [Common Configuration Tasks on page 508](#)
- [LDP Configuration Management Tasks on page 518](#)

LDP Configuration Overview

When the implementation of LDP is instantiated, the protocol is in the `no shutdown` state. In addition, targeted sessions are then enabled. The default parameters for LDP are set to the documented values for targeted sessions in *draft-ietf-mpls-ldp-mib-09.txt*.

LDP must be enabled in order for signaling to be used to obtain the ingress and egress labels in frames transmitted and received on the service distribution path (SDP). When signaling is *off*, labels must be manually configured when the SDP is bound to a service.

Basic LDP Configuration

This chapter provides information to configure LDP and remove configuration examples of common configuration tasks.

The LDP protocol instance is created in the `no shutdown` (enabled) state.

The following displays the default LDP configuration.

```
A:ALA-1>config>router>ldp# info
-----
      interface-parameters
      exit
      targeted-session
      exit
-----
A:ALA-1>config>router>ldp#
```

Common Configuration Tasks

This section provides information to configure:

- [Enabling LDP on page 508](#)
- [Configuring FEC Originate Parameters on page 509](#)
- [Configuring Graceful-Restart Helper Parameters on page 510](#)
- [Applying Export and Import Policies on page 511](#)
- [Targeted Session Parameters on page 512](#)
- [Interface Parameters on page 513](#)
- [Peer Parameters on page 514](#)
- [Interface Parameters on page 513](#)

Enabling LDP

LDP must be enabled in order for the protocol to be active. MPLS does not need to be enabled on the router except if the network interface uses the Packet over Sonet (POS) encapsulation (Sonet path encapsulation type set to ppp-auto). In this case, MPLS must be enabled and the interface name added into MPLS to allow for the MPLSCP to come up on the PPP link between the two peers and for MPLS to be used on the interface. MPLS is enabled in the `config>router>mpls` context.

Use the following syntax to enable LDP on a router:

CLI Syntax: `ldp`

Example: `config>router# ldp`

The following displays the enabled LDP configuration.

```
A:ALA-1>config>router# info
-----
...
#-----
echo "LDP Configuration"
#-----
      ldp
        interface-parameters
        exit
        targeted-session
        exit
      exit
-----
...
A:ALA-1>config>router#
```

Configuring FEC Originate Parameters

A FEC can be added to the LDP IP prefix database with a specific label operation on the node. Permitted operations are pop or swap. For a swap operation, an incoming label can be swapped with a label in the range of 16 to 1048575. If a swap-label is not configured then the default value is 3.

A route table entry is required for a FEC with a pop operation to be advertised. For a FEC with a swap operation, a route-table entry must exist and user configured next-hop for swap operation must match one of the next-hops in route-table entry.

Use the following syntax to configure FEC originate parameters:

CLI Syntax: `config>router>ldp`
`fec-originate ip-prefix/mask [advertised-label in-label]`
`next-hop ip-address [swap-label out-label]`
`fec-originate ip-prefix/mask [advertised-label in-label] pop`

The following displays a FEC originate configuration example.

```
A:ALA-5>config>router# info
-----
fec-originate 100.1.1.1/32 pop
fec-originate 100.2.1.1/32 advertised-label 1000 next-hop 10.10.1.2
fec-originate 100.3.1.1/32 advertised-label 1001 next-hop 10.10.2.3
    swap-label 131071
interface-parameters
exit
targeted-session
exit
exit
-----
A:ALA-5>config>router>ldp#
```

Configuring Graceful-Restart Helper Parameters

Graceful-restart helper advertises to its LDP neighbors by carrying the fault tolerant (FT) session TLV in the LDP initialization message, assisting the LDP in preserving its IP forwarding state across the restart. TiMetra Systems's recovery is self-contained and relies on information stored internally to self-heal. This feature is only used to help third-party routers without a self-healing capability to recover.

Maximum recovery time is the time (in seconds) the sender of the TLV would like the receiver to wait, after detecting the failure of LDP communication with the sender.

Neighbor liveness time is the time (in seconds) the LSR is willing to retain its MPLS forwarding state. The time should be long enough to allow the neighboring LSRs to re-sync all the LSPs in a graceful manner, without creating congestion in the LDP control plane.

Use the following syntax to configure graceful-restart parameters:

CLI Syntax: `config>router>ldp`
`[no] graceful-restart`

Applying Export and Import Policies

Both inbound and outbound label binding filtering are supported. Inbound filtering allows a route policy to control the label bindings an LSR accepts from its peers. An import policy can accept or reject label bindings received from LDP peers.

Label bindings can be filtered based on:

- Neighbor — Match on bindings received from the specified peer.
- Interface — Match on bindings received from a neighbor or neighbors adjacent over the specified interface.
- Prefix-list — Match on bindings with the specified prefix/prefixes.

Outbound filtering allows a route policy to control the set of LDP label bindings advertised by the LSR. An export policy can control the set of LDP label bindings advertised by the router. By default, label bindings for only the system address are advertised and propagate all FECs that are received. Matches can be based on:

- Loopback — loopback interfaces.
- All — all local subnets.
- Match — match on bindings with the specified prefix/prefixes.

Use the following syntax to apply import and export policies:

CLI Syntax: `config>router>ldp`
`export policy-name [policy-name...(upto 32 max)]`
`import policy-name [policy-name...(upto 32 max)]`

The following displays export and import policy configuration examples.

```
A:ALA-1>config>router# info
-----
export "LDP-export"
fec-originate 100.1.1.1/32 pop
fec-originate 100.2.1.1/32 advertised-label 1000 next-hop 10.10.1.2
import "LDP-import"
interface-parameters
exit
targeted-session
exit
-----
A:ALA-1>config>router#
```

Targeted Session Parameters

Use the following syntax to specify **targeted-session** parameters:

```
CLI Syntax: config>router# ldp
                targeted-session
                  disable-targeted-session
                  export-prefixes policy-name [policy-name...(up to 5 max)]
                  hello timeout factor
                  import-prefixes policy-name [policy-name...(up to 5 max)]
                  keepalive timeout factor
                  peer ip-address
                      hello timeout factor
                      keepalive timeout factor
                      no shutdown
                      tunneling
                      lsp lsp-name
```

The following example displays an LDP configuration example:

```
A:ALA-1>config>router>ldp# info
-----
...
    targeted-session
      hello 5000 255
      keepalive 5000 255
      peer 10.10.10.104
        hello 2500 104
        keepalive 15 3
      exit
    exit
-----
A:ALA-1>config>router>ldp#
```


Interface Parameters

Use the following syntax to configure interface parameters:

```
CLI Syntax: config>router# ldp
                interface-parameters
                  hello timeout factor
                  keepalive timeout factor
                  transport-address {system|interface}
                  interface ip-int-name
                    hello timeout factor
                    keepalive timeout factor
                    transport-address {system|interface}
                    no shutdown
```

The following example displays an interface parameter configuration example:

```
A:ALA-1>config>router>ldp# info
-----
...
    targeted-session
      no disable-targeted-session
      hello 5000 255
      keepalive 5000 255
      peer 10.10.10.104
        hello 2500 104
        keepalive 15 3
        no shutdown
      exit
    exit
    no shutdown
-----
A:ALA-1>config>router>ldp#
```

Peer Parameters

Use the following syntax to specify interface parameters:

```
CLI Syntax: config>router# ldp
                peer-parameters
                  peer ip-address
                      auth-keychain name
                      authentication-key [authentication-key|hash-key]
                      [hash|hash2]
                      ttl-security min-ttl-value [log log-id]
```

The following example displays an LDP configuration example:

```
A:ALA-1>config>router>ldp# info
-----
export "LDP-export"
import "LDP-import"
peer-parameters
  peer 10.10.10.104
      authentication-key "3WErEDozxyQ" hash
  exit
exit
interface-parameters
  interface "test"
  exit
  interface "to-104"
  hello 15 3
  exit
exit
targeted-session
  hello 5000 255
  keepalive 5000 255
  peer 10.10.10.104
  hello 2500 100
  keepalive 15 3
  exit
exit
-----
A:ALA-1>config>router>ldp#
```

LDP Signaling and Services

When LDP is enabled, targeted sessions can be established to create remote adjacencies with nodes that are not directly connected. When service distribution paths (SDPs) are configured, extended discovery mechanisms enable LDP to send periodic targeted hello messages to the SDP far-end point. The exchange of LDP hellos trigger session establishment. The SDP signaling default enables **tldp**. The service SDP uses the targeted-session parameters configured in the **config>router>ldp>targeted-session** context.

The SDP LDP and LSP commands are mutually exclusive; either one LSP can be specified or LDP can be enabled. If LDP is already enabled on an MPLS SDP, then an LSP cannot be specified on the SDP. If an LSP is specified on an MPLS SDP, then LDP cannot be enabled on the SDP.

To enable LDP on the SDP when an LSP is already specified, the LSP must be removed from the configuration using the **no lsp lsp-name** command. For further information about configuring SDPs, refer to the 7750 SR OS Services Guide.

The following example displays the command syntax usage to configure enable LDP on an MPLS SDP:

```
CLI Syntax: config>service>sdp#
                ldp
                signaling {off|tldp}
```

The following displays an example of an SDP configuration showing the signaling default **tldp** enabled.

```
A:ALA-1>config>service>sdp# info detail
-----
description "MPLS: to-99"
far-end 10.10.10.99
ldp
signaling tldp
path-mtu 4462
keep-alive
    hello-time 10
    hold-down-time 10
    max-drop-count 3
    timeout 5
    no message-length
    no shutdown
exit
no shutdown
-----
A:ALA-1>config>service>sdp#
```

The following shows a working configuration of LDP over RSVP-TE (1) where tunnels look like the second example (2):

```
1. *A:ALA-1>config>router>ldp# info
-----
prefer-tunnel-in-tunnel
interface-parameters
  interface "port-1/1/3"
  exit
  interface "port-lag-1"
  exit
exit
targeted-session
  peer 10.51.0.1
  shutdown
  tunneling
    lsp "to_P_1"
  exit
  peer 10.51.0.17
  shutdown
  tunneling
    lsp "to_P_6"
  exit
exit
exit
-----
*A:ALA-1>config>router>ldp#
```

```
2. *A:ALA-1>config>router>mpls# info
-----
resignal-timer 30
admin-group "lower" 2
admin-group "upper" 1
interface "system"
exit
interface "port-1/1/3"
exit
interface "port-lag-1"
exit
path "dyn"
  no shutdown
exit
lsp "to_P_1"
  to 10.51.0.1
  cspf
  fast-reroute facility
  exit
  primary "dyn"
  exit
  no shutdown
exit
lsp "to_P_6"
  to 10.51.0.17
  cspf
  fast-reroute facility
```

```
        exit
        primary "dyn"
        exit
        no shutdown
    exit
    no shutdown
-----
*A:ALA-1>config>router>mpls#
```

LDP Configuration Management Tasks

This section discusses the following LDP configuration management tasks:

- [Disabling LDP on page 518](#)
 - [Modifying Targeted Session Parameters on page 519](#)
 - [Modifying Interface Parameters on page 520](#)
 - [Modifying Interface Parameters on page 520](#)
-

Disabling LDP

The **no ldp** command disables the LDP protocol on the router. All parameters revert to the default settings. LDP must be shut down before it can be disabled.

Use the following command syntax to disable LDP:

CLI Syntax: `no ldp
shutdown`

Modifying Targeted Session Parameters

The modification of LDP targeted session parameters does not take effect until the next time the session goes down and is re-establishes. Individual parameters cannot be deleted. The **no** form of a **targeted-session** parameter command reverts modified values back to the default.

The following example displays the command syntax usage to revert targeted session parameters back to the default values:

```
Example :  config>router# ldp
            config>router>ldp# targeted-session
            config>router>ldp>targeted# no authentication-key
            config>router>ldp>targeted# no disable-targeted-session
            config>router>ldp>targeted# no hello
            config>router>ldp>targeted# no keepalive
            config>router>ldp>targeted# no peer 10.10.10.99
```

The following output displays the default values:

```
A:ALA-1>config>router>ldp>targeted# info detail
-----
                no disable-targeted-session
                hello 45 3
                keepalive 40 4
-----
A:ALA-1>config>router>ldp>targeted#
```

Modifying Interface Parameters

Individual parameters cannot be deleted. The **no** form of a **interface-parameter** command reverts modified values back to the defaults.

The following output displays the default values:

```
A:ALA-1>config>router>ldp>targeted# info detail
-----
                hello 15 3
                keepalive 30 3
                no transport-address
-----
A:ALA-1>config>router>ldp>targeted#
```

LDP Command Reference

Command Hierarchies

- [LDP Commands on page 521](#)
- [Show Commands on page 524](#)
- [Clear Commands on page 525](#)
- [Debug Commands on page 525](#)
- [Tools Commands on page 525](#)

LDP Commands

```
config
  — router
    — [no] ldp
      — [no] aggregate-prefix-match
        — prefix-exclude policy-name [policy-name...(up to 5 max)]
        — no prefix-exclude
        — [no] shutdown
      — egress-statistics
        — [no] fec-prefix ip-prefix/mask
          — accounting-policy policy-id
          — no accounting-policy
          — [no] collect-stats
          — [no] shutdown
      — export policy-name [policy-name...(up to 5 max)]
      — no export
      — [no] fast-reroute
      — [no] export-tunnel-table policy-name
      — fec-originate ip-prefix/mask [advertised-label in-label] [swap-label out-label]
      — interface interface-name
      — fec-originate ip-prefix/mask [advertised-label i in-label] next-hop ip-address
      — [swap-label out-label]
      — fec-originate ip-prefix/mask [advertised-label in-label] next-hop ip-address
      — [swap-label out-label] interface interface-name
      — fec-originate ip-prefix/mask [advertised-label in-label] pop
      — no fec-originate ip-prefix/mask interface interface-name
      — no fec-originate ip-prefix/mask next-hop ip-address
      — no fec-originate ip-prefix/mask next-hop ip-address interface interface-name
      — no fec-originate ip-prefix/mask pop
      — [no] graceful-restart
        — maximum-recovery-time interval
        — no maximum-recovery-time
        — neighbor-liveness-time interval
        — no neighbor-liveness-time
      — [no] implicit-null-label
      — import policy-name [policy-name...(up to 5 max)]
      — no import
```

- **interface-parameters**
 - **hello** *timeout factor*
 - **no hello**
 - **[no] interface** *ip-int-name*
 - **[no] bfd-enable**
 - **hello** *timeout factor*
 - **no hello**
 - **keepalive** *timeout factor*
 - **no keepalive**
 - **[no] multicast-traffic** {enable|disable}
 - **local-lsr-id** {system | interface | interface-name *interface-name*}
 - **no local-lsr-id**
 - **[no] shutdown**
 - **transport-address** {system | interface}
 - **no transport-address**
 - **keepalive** *timeout factor*
 - **no keepalive**
 - **transport-address** {system | interface}
 - **no transport-address**
- **label-withdrawal-delay** *seconds*
- **[no] mcast-upstream-frr**
- **mp-mbb-time** *interval*
- **no mp-mbb-time**
- **peer-parameters**
 - **peer** *ip-address*
 - **no peer** [*ip-address*]
 - **[no] aggregate-prefix-match**
 - **auth-keychain** *name*
 - **authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]
 - **no authentication-key**
 - **[no] dod-label-distribution**
 - **[no] export-prefixes** *policy-name*
 - **[no] fec129-cisco-interop**
 - **[no] import-prefixes** *policy-name*
 - **[no] path-mtu-discovery**
 - **[no] pe-id-mac-flush-interop**
 - **ttl-security** *min-ttl-value*
 - **no ttl-security**
- **[no] mp-mbb-time**
- **[no] prefer-tunnel-in-tunnel**
- **[no] shortcut-transit-ttl-propagate**
- **[no] shortcut-local-ttl-propagate**
- **[no] shutdown**
- **targeted-session**
 - **[no] disable-targeted-session**
 - **export-prefixes** *policy-name* [*policy-name...*(up to 5 max)]
 - **no export-prefixes**
 - **hello** *timeout factor*
 - **no hello**
 - **hello-reduction** {enable *factor* | **disable**}
 - **no hello-reduction**
 - **import-prefixes** *policy-name* [*policy-name...*(up to 5 max)]
 - **no import-prefixes**

- **keepalive** *timeout factor*
- **no keepalive**
- **peer** *ip-address*
- **no peer** *ip-address*
 - **[no] bfd-enable**
 - **hello** *timeout factor*
 - **no hello**
 - **hello-reduction** {**enable** *factor* / **disable**}
 - **no hello-reduction**
 - **keepalive** *timeout factor*
 - **no keepalive**
 - **local-lsr-id** *interface-name*
 - **no local-lsr-id**
 - **[no] shutdown**
 - **[no] tunneling**
 - **[no] lsp**
- **[no] peer-template** *template-name*
 - **[no] bfd-enable**
 - **hello** *timeout factor*
 - **no hello**
 - **hello-reduction** {**enable** *factor* / **disable**}
 - **no hello-reduction**
 - **keepalive** *timeout factor*
 - **no keepalive**
 - **local-lsr-id** *interface-name*
 - **no local-lsr-id**
 - **[no] shutdown**
 - **[no] tunneling**
- **peer-template-map** *template-name* **policy** *peer-prefix-policy1* [*peer-prefix-policy2... up to 5*]
- **no peer-template-map** *template-name*
- **tunnel-down-damp-time** *seconds*
- **no tunnel-down-damp-time**
- **[no] ldp-shortcut**

Show Commands

```
show
  — router
    — ldp
      — adv-adj-addr-only
      — auth-keychain [keychain]
      — bindings active [prefix ip-prefix/mask] [summary | egress-nh ip-prefix/mask |
        egress-if port-id | egress-lsp tunnel-id]
      — bindings active [fec-type prefixes] [prefix <ip-prefix/mask>] [egress-nh <ip-pre-
        fix/mask> | egress-if <port-id> | egress-lsp <tunnel-id>] [summary]
      — bindings active [fec-type p2mp] [p2mp-id <identifier> root <ip-address>]
        [egress-nh <ip-prefix/mask> | egress-if <port-id> | egress-lsp <tunnel-id>] [sum-
        mary]
      — bindings[fec-type fec-type [detail | summary]] [session ip-addr[:label-space]]
      — bindings [label-type] [start-label [end-label]]
      — bindings {prefix ip-prefix/mask [detail]}[session ip-addr[:label-space]]
      — bindings active [prefix ip-prefix/mask] [summary | egress-nh ip-prefix/mask |
        egress-if port-id | egress-lsp tunnel-id]
      — bindings service-id service-id [detail]
      — bindings vc-type vc-type [{vc-id vc-id| agi agi} [session ip-addr[:lab el-space]]]
      — discovery [{peer [ip-address]} | {interface [ip-int-name]}] [state state] [detail]
        [adjacency-type type]
      — fec-egress-stats [ip-prefix/mask]
      — fec-egress-stats active
      — fec-originate ip-prefix/mask [operation-type]
      — interface [ip-int-name | ip-address] [detail]
      — parameters
      — peer [ip-address] [detail]
      — peer-parameters peer-ip-address
      — peer-template
      — peer-template-map [tldp-peers]
      — session [ip-addr[:label-space]] [detail | statistics [packet-type]] [session-type]
      — status
      — statistics-summary
```

Note: See 7750 SR *OS OAM and Diagnostics Guide* for tools command descriptions, syntax, and usage information.

Clear Commands

```
clear
  — router
    — ldp
      — fec-egress-statistics [ip-prefix/mask]
      — instance
      — interface [ip-int-name]
      — peer [ip-address] [statistics]
      — session [ip-addr[:label-space]] [statistics]
      — statistics
```

Debug Commands

```
[no] debug
  — router
    — [no] ldp
      — [no] interface interface-name
        — [no] event
          — [no] messages
        — [no] packet [detail]
          — hello [detail]
          — no hello
      — peer ip-address
        — [no] event
          — [no] bindings
          — [no] messages
        — [no] packet
          — hello [detail]
          — no hello
          — init [detail]
          — no init
          — [no] keepalive
          — label [detail]
          — no label
```

Tools Commands

See 7750 SR OS OAM and Diagnostics Guide for CLI description and syntax.

```
tools
  — dump
    — ldp-treetrace {prefix ip-prefix/mask| manual-prefix ip-prefix/mask}[path-destination ip-address] [trace-tree]
    — router
      — ldp
        — peer ip-address
```

LDP Configuration Commands

Generic Commands

ldp

Syntax	<code>[no] ldp</code>
Context	config>router
Default	<p>This command creates the context to configure an LDP parameters. LDP is not enabled by default and must be explicitly enabled (no shutdown).</p> <p>To suspend the LDP protocol, use the shutdown command. Configuration parameters are not affected.</p> <p>The no form of the command deletes the LDP protocol instance, removing all associated configuration parameters. The LDP instance must first be disabled with the shutdown command before being deleted.</p>
Default	none (LDP must be explicitly enabled)

ldp-shortcut

Syntax	<code>[no] ldp-shortcut</code>
Context	config>router
Description	<p>This command enables the resolution of IGP routes using LDP LSP across all network interfaces participating in the IS-IS and OSPF routing protocol in the system.</p> <p>When LDP shortcut is enabled, LDP populates the routing table with next-hop entries corresponding to all prefixes for which it activated an LDP FEC. For a given prefix, two route entries are populated in the system routing table. One corresponds to the LDP shortcut next-hop and has an owner of LDP. The other one is the regular IP next-hop. The LDP shortcut next-hop always has preference over the regular IP next-hop for forwarding user packets and specified control packets over a given outgoing interface to the route next-hop.</p> <p>All user and specified control packets for which the longest prefix match in RTM yields the FEC prefix will be forwarded over the LDP LSP.</p> <p>When an IPv4 packet is received on an ingress network interface, a subscriber IES interface, or a regular IES interface, the lookup of the packet by the ingress forwarding engine will result in the packet being sent labeled with the label stack corresponding to the NHLFE of the LDP LSP when the preferred RTM entry corresponds to an LDP shortcut.</p> <p>If the preferred RTM entry corresponds to an IP next-hop, the IPv4 packet is forwarded unlabelled.</p> <p>When ECMP is enabled and multiple equal-cost next-hops exit for the IGP route, the ingress forwarding engine will spray the packets for this route based on hashing routine currently supported</p>

for IPv4 packets. When the preferred RTM entry corresponds to an LDP shortcut route, spraying will be performed across the multiple next-hops for the LDP FEC. The FEC next-hops can either be direct link LDP neighbors or T-LDP neighbors reachable over RSVP LSPs in the case of LDP-over-RSVP but not both.

When the preferred RTM entry corresponds to a regular IP route, spraying will be performed across regular IP next-hops for the prefix..

The **no** form of this command disables the resolution of IGP routes using LDP shortcuts.

Default no ldp-shortcut

shutdown

Syntax [no] shutdown

Context config>router>ldp
config>router>ldp>if-params>if
config>router>ldp>targ-session>peer
config>router>ldp>egr-stats>fec-prefix
config>router>ldp>aggregate-prefix-match

Description This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.

The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

The **no** form of this command administratively enables an entity.

Unlike other commands and parameters where the default state is not indicated in the configuration file, the **shutdown** and **no shutdown** states are always indicated in system generated configuration files.

The **no** form of the command places an entity in an administratively enabled state.

Default no shutdown

adv-adj-addr-only

Syntax [no] adv-adj-addr-only

Context config>router>ldp>peer-parameters>peer

Description This command provides a means for an LDP router to advertise only the local interfaces it uses to establish hello adjacencies with an LDP peer. By default, when a router establishes an LDP session with a peer, it advertises in an LDP Address message the addresses of all local interfaces to allow the peer to resolve LDP FECs distributed by this router. Similarly, a router sends a Withdraw Address message to of all its peers to withdraw a local address if the corresponding interface went down or was deleted.

This new option reduces CPU processing when a large number of LDP neighbors come up or go down. The new CLI option is strongly recommended in mobile backhaul networks where the number of LDP peers can be very large.

The **no** version of this command reverts LDP to the default behaviour of advertising all local interfaces.

aggregate-prefix-match

Syntax	[no] aggregate-prefix-match
Context	config>router>ldp
Description	<p>The command enables the use by LDP of the aggregate prefix match procedures.</p> <p>When this option is enabled, LDP performs the following procedures for all prefixes. When an LSR receives a FEC-label binding from an LDP neighbor for a given specific FEC1 element, it will install the binding in the LDP FIB if:</p> <ul style="list-style-type: none">• It is able to perform a successful longest IP match of the FEC prefix with an entry in the routing table, and• The advertising LDP neighbor is the next-hop to reach the FEC prefix. <p>When such a FEC-label binding has been installed in the LDP FIB, then LDP programs an NHLFE entry in the egress data path to forward packets to FEC1. It also advertises a new FEC-label binding for FEC1 to all its LDP neighbors.</p> <p>When a new prefix appears in the routing table, LDP inspects the LDP FIB to determine if this prefix is a better match (a more specific match) for any of the installed FEC elements. For any FEC for which this is true, LDP may have to update the NHLFE entry for this FEC.</p> <p>When a prefix is removed from the routing table, LDP inspects the LDP FIB for all FEC elements which matched this prefix to determine if another match exists in the routing table. If so, it updates the NHLFE entry accordingly. If not, it sends a label withdraw message to its LDP neighbors to remove the binding.</p> <p>When the next hop for a routing prefix changes, LDP updates the LDP FIB entry for the FEC elements which matched this prefix. It also updates the NHLFE entry for these FEC elements accordingly.</p> <p>The no form of this command disables the use by LDP of the aggregate prefix procedures and deletes the configuration. LDP resumes performing exact prefix match for FEC elements.</p>
Default	no aggregate-prefix-match

prefix-exclude

Syntax	prefix-exclude <i>policy-name</i> [<i>policy-name...</i> (up to 5 max)] no prefix-exclude
Context	config>router>ldp>aggregate-prefix-match
Description	This command specifies the policy name containing the prefixes to be excluded from the aggregate prefix match procedures. In this case, LDP will perform an exact match of a specific FEC element prefix as opposed to a longest match of one or more LDP FEC element prefixes, against this prefix when it receives a FEC-label binding or when a change to this prefix occurs in the routing table. The no form of this command removes all policies from the configuration.
Default	no prefix-exclude.

egress-statistics

Syntax	egress-statistics
Context	config>router>ldp
Description	This command provides the context for the user to enter the LDP FEC prefix for the purpose of enabling egress data path statistics at the ingress LER for this FEC.
Default	none

fec-prefix

Syntax	[no] fec-prefix <i>ip-prefix</i> [/ <i>mask</i>]
Context	config>router>ldp>egr-stats
Description	This command configures statistics in the egress data path at the ingress LER or LSR for an LDP FEC. The user must execute the no shutdown command for this command to effectively enable statistics. The egress data path counters will be updated for both originating and transit packets. Originating packets may be service packets or IP user and control packets forwarded over the LDP LSP when used as an IGP shortcut. Transit packets of the FEC which are label switched on this node. When ECMP is enabled and multiple paths exist for a FEC, the same set of counters are updated for each packet forwarded over any of the NHLFEs associated with this FEC and for as long as this FEC is active. The statistics can be enabled on prefix FECs imported from both LDP neighbors and T-LDP neighbors (LDP over RSVP). Only /32 FEC prefixes are accepted. Service FECs, i.e., FEC 128 and FEC 129 are not valid. LDP FEC egress statistics are not collected at the Penultimate-Popping Hop (PHP) node for a LDP FEC using an implicit null label. The no form of this command disables the statistics in the egress data path and removes the accounting policy association from the LDP FEC.
Default	none

accounting-policy

Syntax	accounting-policy <i>acct-policy-id</i> no accounting-policy
Context	config>router>ldp>egr-stats
Description	This command associates an accounting policy to the MPLS instance. An accounting policy must be defined before it can be associated else an error message is generated. The no form of this command removes the accounting policy association.
Default	none
Parameters	<i>acct-policy-id</i> — Enter the accounting <i>policy-id</i> as configured in the config>log>accounting-policy context. Values 1 — 99

collect-stats

Syntax	[no] collect-stats
Context	config>router>ldp>egr-stats
Description	This command enables accounting and statistical data collection. When applying accounting policies the data, by default, is collected in the appropriate records and written to the designated billing file. When the no collect-stats command is issued the statistics are still accumulated by the forwarding engine. However, the CPU will not obtain the results and write them to the billing file. If a subsequent collect-stats command is issued then the counters written to the billing file include all the traffic while the no collect-stats command was in effect.
Default	collect-stats

export

Syntax	export <i>policy-name</i> [<i>policy-name</i> ...upto 5 max] no export
Context	config>router>ldp
Description	This command specifies the export route policies used to determine which routes are exported to LDP. Policies are configured in the config>router>policy-options context. If no export policy is specified, non-LDP routes will not be exported from the routing table manager to LDP. LDP-learned routes will be exported to LDP neighbors. Present implementation of export policy (outbound filtering) can be used “only” to add FECs for label propagation. The export policy does not control propagation of FECs that an LSR receives from its neighbors.

If multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied. If multiple export commands are issued, the last command entered will override the previous command. A maximum of 5 policy names can be specified.

The **no** form of the command removes all policies from the configuration.

Default **no export** — No export route policies specified.

Parameters *policy-name* — The export route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

The specified name(s) must already be defined.

fast-reroute

Syntax **[no] fast-reroute**

Context config>router>ldp

Description This command enables LDP Fast-Reroute (FRR) procedures. When enabled, LDP uses both the primary next-hop and LFA next-hop, when available, for resolving the next-hop of an LDP FEC against the corresponding prefix in the routing table. This will result in LDP programming a primary NHLFE and a backup NHLFE into the forwarding engine for each next-hop of a FEC prefix for the purpose of forwarding packets over the LDP FEC.

When any of the following events occurs, LDP instructs in the fast path the forwarding engines to enable the backup NHLFE for each FEC next-hop impacted by this event:

- An LDP interface goes operationally down, or is admin shutdown.
- An LDP session to a peer went down as the result of the Hello or Keep-Alive timer expiring.
- The TCP connection used by a link LDP session to a peer went down, due say to next-hop tracking of the LDP transport address in RTM, which brings down the LDP session.
- A BFD session, enabled on a T-LDP session to a peer, times-out and as a result the link LDP session to the same peer and which uses the same TCP connection as the T-LDP session goes also down.
- A BFD session enabled on the LDP interface to a directly connected peer, times out and brings down the link LDP session to this peer.

The **tunnel-down-dump-time** option or the **label-withdrawal-delay** option, when enabled, does not cause the corresponding timer to be activated for a FEC as long as a backup NHLFE is still available.

Note that because LDP can detect the loss of a neighbor/next-hop independently, it is possible that it switches to the LFA next-hop while IGP is still using the primary next-hop. Also, when the interface for the previous primary next-hop is restored, IGP may re-converge before LDP completed the FEC exchange with it neighbor over that interface. This may cause LDP to de-program the LFA next-hop from the FEC and blackhole traffic. In order to avoid this situation, it is recommended to enable IGP-LDP synchronization on the LDP interface.

When the SPF computation determines there is more than one primary next-hop for a prefix, it will not program any LFA next-hop in RTM. Thus, the LDP FEC will resolve to the multiple primary next-hops that provide the required protection.

The **no** form of this command disables LDP FRR.

Default no fast-reroute

export-tunnel-table

Syntax **[no] export-tunnel-table** *policy-name*

Context config>router>ldp

Description This command applies a tunnel table export policy to LDP for the purpose of learning BGP labeled routes from the CPM tunnel table and stitching them to LDP FEC for the same prefix.

The user enables the stitching of routes between LDP and BGP by configuring separately tunnel table route export policies in both protocols and enabling the advertising of RFC 3107, *Carrying Label Information in BGP-4*, formatted labeled routes for prefixes learned from LDP FECs.

The route export policy in BGP instructs BGP to listen to LDP route entries in the CPM Tunnel Table. If a /32 LDP FEC prefix matches an entry in the export policy, BGP originates a BGP labeled route, stitches it to the LDP FEC, and re-distributes the BGP labeled route to its iBGP neighbors.

The user adds LDP FEC prefixes with the statement '**from protocol ldp**' in the configuration of the existing BGP export policy at the global level, the peer-group level, or at the peer level using the commands:

- **configure>router>bgp>export** *policy-name*
- **configure>router>bgp>group>export** *policy-name*
- **configure>router>bgp>group>neighbour>export** *policy-name*

To indicate to BGP to evaluate the entries with the '**from protocol ldp**' statement in the export policy when applied to a specific BGP neighbor, a new argument is added to the existing advertise-label command:

```
configure>router>bgp>group>neighbour>advertise-label ipv4 include-ldp-prefix
```

Without the new **include-ldp-prefix** argument, only core IPv4 routes learned from RTM are advertised as BGP labeled routes to the neighbor. No stitching of LDP FEC to the BGP labeled route will be performed for this neighbor even if the same prefix was learned from LDP.

The tunnel table route export policy in LDP instructs LDP to listen to BGP route entries in the CPM Tunnel Table. If a /32 BGP labeled route matches a prefix entry in the export policy, LDP originates an LDP FEC for the prefix, stitches it to the BGP labeled route, and re-distributes the LDP FEC to its iBGP neighbors.

The user can add BGP labeled route prefixes with the statement '**from protocol bgp**' in the configuration of the LDP tunnel table export policy. Note that the '**from protocol**' statement has an effect only when the protocol value is ldp. Policy entries with protocol values of rsvp, bgp, or any value other than ldp are ignored at the time the policy is applied to LDP.

The **no** form of the command removes the policy from the configuration.

Default **no export-tunnel-table** — no tunnel table export route policy is specified.

Parameters *policy-name* — The export-tunnel-table route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the

string contains spaces, use double quotes to delimit the start and end of the string. The specified name(s) must already be defined.

fec-originate

Syntax **fec-originate** *ip-prefix/mask* [**advertised-label** *in-label*] [**swap-label** *out-label*] **interface** *interface-name*
fec-originate *ip-prefix/mask* [**advertised-label** *in-label*] **next-hop** *ip-address* [**swap-label** *out-label*]
fec-originate *ip-prefix/mask* [**advertised-label** *in-label*] **next-hop** *ip-address* [**swap-label** *out-label*] **interface** *interface-name*
fec-originate *ip-prefix/mask* [**advertised-label** *in-label*] **pop**
no fec-originate *ip-prefix/mask* **interface** *interface-name*
no fec-originate *ip-prefix/mask* **next-hop** *ip-address*
no fec-originate *ip-prefix/mask* **next-hop** *ip-address* **interface** *interface-name*
no fec-originate *ip-prefix/mask* **pop**

Context config>router>ldp

Description This command defines a way to originate a FEC (with a swap action) for which the LSR is not egress, or to originate a FEC (with a pop action) for which the LSR is egress.

Parameters *ip-prefix/mask* — Specify information for the specified IP prefix and mask length.

next-hop — Specify the IP address of the next hop of the prefix.

advertised-label — Specify the label advertised to the upstream peer. If not configured, then the label advertised should be from the label pool. If the configured static label is not available then the IP prefix is not advertised.

out-label — Specify the LSR to swap the label. If configured, then the LSR should swap the label with the configured swap-label. If not configured, then the default action is pop if the next-hop parameter is not defined.

NOTE: The next-hop, advertised-label, swap-label parameters are all optional. If next-hop is configured but no swap label specified, then it will be a swap with label 3, such as, pop and forward to the next-hop. If the next-hop and swap-label are configured, then it is a regular swap. If no parameters are specified, then a pop and route is performed.

Values 16 — 1048575

in-label — Specifies the number of labels to send to the peer associated with this FEC.

Values 32 — 1023

pop — Specifies to pop the label and transmit without the label.

interface *interface-name* — Specifies the name of the interface the label for the originated FEC is swapped to. For an unnumbered interface, this parameter is mandatory since there is no address for the next-hop. For a numbered interface, it is optional.

graceful-restart

Syntax	[no] graceful-restart
Context	config>router>ldp
Description	This command enables graceful restart helper. The no form of the command disables graceful restart. Note that graceful restart helper configuration changes, enable/disable or change of a parameter, will cause the LDP session to bounce.
Default	no graceful-restart (disabled) — Graceful-restart must be explicitly enabled.

implicit-null-label

Syntax	[no] implicit-null-label
Context	config>router>ldp
Description	This command enables the use of the implicit null label. Use this command to signal the IMPLICIT NULL option for all LDP FECs for which this node is the egress LER. The no form of this command disables the signaling of the implicit null label.
Default	no implicit-null-label

maximum-recovery-time

Syntax	maximum-recovery-time <i>interval</i> no maximum-recovery-time
Context	config>router>ldp
Description	This command configures the local maximum recovery time. The no form of the command returns the default value.
Default	120
Parameters	<i>interval</i> — Specifies the length of time in seconds. Values 15 — 1800

neighbor-liveness-time

Syntax	neighbor-liveness-time <i>interval</i> no neighbor-liveness-time
Context	config>router>ldp
Description	This command configures the neighbor liveness time.

The **no** form of the command returns the default value.

Default 120

Parameters *interval* — Specifies the length of time in seconds.

Values 5 — 300

import

Syntax **import** *policy-name* [*policy-name* ...upto 5 max]
no import

Context config>router>ldp

Description This command configures import route policies to determine which label bindings (FECs) are accepted from LDP neighbors. Policies are configured in the **config>router>policy-options** context.

If no import policy is specified, LDP accepts all label bindings from configured LDP neighbors. Import policies can be used to limit or modify the routes accepted and their corresponding parameters and metrics.

If multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied. If multiple import commands are issued, the last command entered will override the previous command. A maximum of five policy names can be specified.

The **no** form of the command removes all policies from the configuration.

Default **no import** — No import route policies specified.

Parameters *policy-name* — The import route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

The specified name(s) must already be defined.

label-withdrawal-delay

Syntax **label-withdrawal-delay** *seconds*

Context config>router>ldp

Description This command specifies configures the time interval, in seconds, LDP will delay for the withdrawal of FEC-label binding it distributed to its neighbors when FEC is de-activated. When the timer expires, LDP then sends a label withdrawal for the FEC to all its neighbours. This is applicable only to LDP IPv4 prefix FECs and is not applicable to pseudowires (service FECs).

When there is an upper layer (user of LDP) which depends of LDP control plane for failover detection then label withdrawal delay and tunnel-down-damp-time options must be set to 0.

An example is PW redundancy where the primary PW doesn't have its own fast failover detection mechanism and the node depends on LDP tunnel down event to activate the standby PW.

Default no label-withdrawal-delay

Parameters *seconds* — Specifies the time that LDP delays the withdrawal of FEC-label binding it distributed to its neighbors when FEC is de-activated.

Values 3 — 120

mcast-upstream-frr

Syntax `[no] mcast-upstream-frr`

Context `config>router>ldp`

Description When LDP programs the primary ILM record in the data path, it provides the IOM with the This command enables the mLDP fast upstream switchover feature.

When this command is enabled and LDP is resolving a mLDP FEC received from a downstream LSR, it checks if an ECMP next-hop or a LFA next-hop exist to the root LSR node. If LDP finds one, it programs a primary ILM on the interface corresponding to the primary next-hop and a backup ILM on the interface corresponding to the ECMP or LFA next-hop. LDP then sends the corresponding labels to both upstream LSR nodes. In normal operation, the primary ILM accepts packets while the backup ILM drops them. If the interface or the upstream LSR of the primary ILM goes down causing the LDP session to go down, the backup ILM will then start accepting packets.

In order to make use of the ECMP next-hop, the user must configure the **ecmp** value in the system to at least 2 using the following command:

```
configure>router>ecmp
```

In order to make use of the LFA next-hop, the user must enable LFA using the following commands:

```
config>router>isis>loopfree-alternate
```

```
config>router>ospf>loopfree-alternate
```

Enabling IP FRR or LDP FRR features is not strictly required since LDP only needs to know where the alternate next-hop to the root LSR is to be able to send the Label Mapping message to program the backup ILM at the initial signaling of the tree. Thus enabling the LFA option is sufficient. If however, unicast IP and LDP prefixes need to be protected, then these features and the mLDP fast upstream switchover can be enabled concurrently.

Note that mLdp FRR fast switchover relies on the fast detection of loss of ****LDP session**** to the upstream peer to which primary ILM label had been advertised. As a result it is strongly recommended to perform the following:

- Enable BFD on all LDP interfaces to upstream LSR nodes. When BFD detects the loss of the last adjacency to the upstream LSR, it will bring down immediately the LDP session which will cause the IOM to activate the backup ILM.
- If there is a concurrent TLDP adjacency to the same upstream LSR node, enable BFD on the T-LDP peer in addition to enabling it on the interface.
- Enable the **ldp-sync-timer** option on all interfaces to the upstream LSR nodes. If an LDP session to the upstream LSR to which the primary ILM is resolved goes down for any other reason than a failure of the interface or of the upstream LSR, routing and LDP will go out of sync. This means the backup ILM will remain activated until the next time SPF is rerun by IGP. By enabling IGP-LDP synchronization feature, the advertised link metric will be changed to max value as soon as the LDP session goes down. This in turn will trigger an SPF and LDP will likely download a new set of primary and backup ILMs.

The **no** form of this command disables the fast upstream switchover for mLDP FECs.

Default no mcast-upstream-frr

mp-mbb-time

Syntax **mp-mbb-time** *interval*
no mp-mbb-time

Context config>router>ldp

Description This command configures the maximum time a Multi Point (MP) transit node must wait before switching over to a new path if the new node does not send Make Before Break (MBB) Tag Length Value (TLV) to inform of the availability of the data plane.

Parameters *interval* — Specifies the MP MBB time, in seconds.

Values 0 — 10

tunnel-down-damp-time

Syntax **tunnel-down-damp-time** *seconds*
no tunnel-down-damp-time

Context config>router>ldp

Description This command specifies the time interval, in seconds, that LDP waits before posting a tunnel down event to the Tunnel Table Manager (TTM).

When LDP can no longer resolve a FEC and de-activates it, it de-programs the NHLFE in the data path. It will however delay deleting the LDP tunnel entry in the TTM until the tunnel-down-damp-time timer expires. This means users of the LDP tunnel, such as SDPs (all services) and BGP (L3 VPN), will not be notified immediately. Traffic is still blackholed because the forwarding engine NHLFE has been de-programmed.

If the FEC gets resolved before the tunnel-down-damp-time timer expires, then LDP programs the forwarding engine with the new NHLFE and performs a tunnel modify event in TTM updating the dampened entry in TTM with the new NHLFE information. If the FEC does not get resolved and the tunnel-down-damp-time timer expires, LDP posts a tunnel down event to TTM which deletes the LDP tunnel.

When there is an upper layer (user of LDP) which depends of LDP control plane for failover detection then label withdrawal delay and tunnel-down-damp-time options must be set to 0.

An example is pseudowire redundancy where the primary PW doesn't have its own fast failover detection mechanism and the node depends on LDP tunnel down event to activate the standby PW.

The **no** form of this command then tunnel down events are not damped.

Parameters *seconds* — Specifies the time interval, in seconds, that LDP waits before posting a tunnel down event to the Tunnel Table Manager.

keepalive

Syntax **keepalive** *timeout factor*
no keepalive

Context config>router>ldp>if-params
config>router>ldp>if-params>if
config>router>ldp>targ-session
config>router>ldp>targ-session>peer

Description This command configures the time interval, in seconds, that LDP waits before tearing down the session. The **factor** parameter derives the keepalive interval.

If no LDP messages are exchanged for the configured time interval, the LDP session is torn down. Keepalive timeout is usually three times the keepalive interval. To maintain the session permanently, regardless of the activity, set the value to zero.

When LDP session is being set up, the keepalive timeout is negotiated to the lower of the two peers. Once a operational value is agreed upon, the keepalive factor is used to derive the value of the keepalive interval.

The **no** form of the command at the interface-parameters and targeted-session levels sets the **keepalive timeout** and the **keepalive factor** to the default value.

The **no** form of the command, at the interface level, sets the **keepalive timeout** and the **keepalive factor** to the value defined under the **interface-parameters** level.

The **no** form of the command, at the peer level, will set the **keepalive timeout** and the **keepalive factor** to the value defined under the **targeted-session** level.

Note that the session needs to be flapped for the new args to operate.

Default

Context	timeout	factor
config>router>ldp>if-params	30	3
config>router>ldp>targ-session	40	4
config>router>ldp>if-params>if	Inherits values from interface-parameters context.	
config>router>ldp>targ-session>peer	Inherits values from targeted-session context.	

Parameters *timeout* — Configures the time interval, expressed in seconds, that LDP waits before tearing down the session.

Values 1 — 65535

factor — Specifies the number of keepalive messages, expressed as a decimal integer, that should be sent on an idle LDP session in the keepalive timeout interval.

Values 1 — 255

local-lsr-id

Syntax	local-lsr-id { system interface interface-name <i>interface-name</i> } no local-lsr-id
Context	config>router>ldp>interface-parameters>interface
Description	<p>This command enables the use of the address of the local LDP interface, or any other network interface configured on the system, as the LSR-ID to establish link LDP Hello adjacency and LDP session with directly connected LDP peers. The network interface can be a loopback or not.</p> <p>Link LDP sessions to all peers discovered over a given LDP interface share the same local LSR-ID. However, LDP sessions on different LDP interfaces can use different network interface addresses as their local LSR-ID.</p> <p>By default, the LDP session to a peer uses the system interface address as the LSR-ID unless explicitly configured using the above command. Note, however, that the system interface must always be configured on the router, or the LDP protocol will not come up on the node. There is no requirement to include it in any routing protocol.</p> <p>At initial configuration, the LDP session to a peer will remain down while the network interface used as LSR-ID is down. LDP will not try to bring it up using the system interface.</p> <p>At any time the network IP interface used as LSR-ID goes down, the LDP sessions to all discovered peers using this LSR-ID go down.</p> <p>If the user changes the LSR-ID value on the fly between system, interface, and <i>interface-name</i> while the LDP session is up, LDP will immediately tear down all sessions using this LSR-ID and will attempt to re-establish them using the new LSR-ID.</p> <p>Note that when an interface other than system is used as the LSR-ID, the transport connection (TCP) for the link LDP session will also use the address of that interface as the transport address. If system or interface value is configured in the configure>router>ldp>interface-parameters>interface>transport-address context, it will be overridden.</p> <p>The no form of the command returns to the default behavior in which case the system interface address is used as the LSR-ID.</p>
Default	no local-lsr-id
Parameters	<i>interface-name</i> — Specifies the name, up to 32 character in length, of the network IP interface. AN interface name cannot be in the form of an IP address. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

tunneling

Syntax	[no] tunneling
Context	config>router>ldp>targ-session>peer
Description	<p>This command enables LDP over tunnels.</p> <p>The no form of the command disables tunneling.</p>
Default	no tunneling

lsp

Syntax	[no] lsp <i>lsp-name</i>
Context	config>router>ldp>targ-session>tunneling
Description	<p>This command configures a specific LSP destined to this peer and to be used for tunneling of LDP FEC over RSVP. A maximum of 4 RSVP LSPs can be explicitly used for tunneling LDP FECs to the T-LDP peer.</p> <p>It is not necessary to specify any RSVP LSP in this context unless there is a need to restrict the tunneling to selected LSPs. All RSVP LSPs with a to address matching that of the T-LDP peer are eligible by default. The user can also exclude specific LSP names by using the ldp-over-rsvp exclude command in the configure->router->mpls->lsp <i>lsp-name</i> context.</p>

Interface Parameters Commands

interface-parameters

Syntax	interface-parameters
Context	config>router>ldp
Description	This command enables the context to configure LDP interfaces and parameters applied to LDP interfaces.

bfd-enable

Syntax	[no] bfd-enable
Context	config>router>ldp>if-params>if
Description	<p>This command enables tracking of the Hello adjacency to an LDP peer using BFD.</p> <p>When this command is enabled on an LDP interface, LDP registers with BFD and starts tracking the LSR-id of all peers it formed Hello adjacencies with over that LDP interface. The LDP hello mechanism is used to determine the remote address to be used for the BFD session. The parameters used for the BFD session, that is, transmit-interval, receive-interval, and multiplier are those configured under the IP interface in existing implementation: config>router>interface>bfd</p> <p>If a BFD session fails then the associated LDP adjacency is also declared down and LDP will immediately begin its re-convergence.</p> <p>When multiple links exist to the same LDP peer, a Hello adjacency is established over each link and a separate BFD session is enabled on each LDP interface. If a BFD session times out on a specific link, LDP will immediately associate the LDP session with one of the remaining Hello adjacencies and trigger the LDP FRR procedures. As soon as the last Hello adjacency goes down due to BFD timing out, the LDP session goes down and the LDP FRR procedures will be triggered.</p> <p>The no form of this command disables BFD on the LDP interface.</p>
Default	no bfd-enable

hello

Syntax **hello** *timeout factor*
no hello

Context config>router>ldp>if-params
config>router>ldp>if-params>if
config>router>ldp>targ-session
config>router>ldp>targ-session>peer

Description This command configures the time interval to wait before declaring a neighbor down. The **factor** parameter derives the hello interval.

Hold time is local to the system and sent in the hello messages to the neighbor. Hold time cannot be less than three times the hello interval. The hold time can be configured globally (applies to all LDP interfaces) or per interface. The most specific value is used.

When LDP session is being set up, the holddown time is negotiated to the lower of the two peers. Once a operational value is agreed upon, the hello factor is used to derive the value of the hello interval.

The **no** form of the command at the interface-parameters and targeted-session level sets the **hello timeout** and the **hello factor** to the default values.

The **no** form of the command, at the interface level, will set the **hello timeout** and the **hello factor** to the value defined under the interface-parameters level.

The **no** form of the command, at the peer level, will set the **hello timeout** and the **hello factor** to the value defined under the targeted-session level.

Note that the session needs to be flapped for the new args to operate.

Default

Context	Timeout	Factor
config>router>ldp>if-params	15	3
config>router>ldp>targ-session	45	3
config>router>ldp>if-params>if	Inherits values from interface-parameters context.	
config>router>ldp>targ-session>peer	Inherits values from targeted-session context.	

Parameters *timeout* — Configures the time interval, in seconds, that LDP waits before a neighbor down.
Values 1 — 65535

factor — Specifies the number of keepalive messages that should be sent on an idle LDP session in the hello timeout interval.
Values 1 — 255

hello-reduction

Syntax	hello-reduction { enable <i>factor</i> / disable } no hello-reduction
Context	config>router>ldp>targeted-session config>router>ldp>targeted-session>peer
Description	<p>This command enables the suppression of periodic targeted Hello messages between LDP peers once the targeted LDP session is brought up.</p> <p>When this feature is enabled, the target Hello adjacency is brought up by advertising the Hold-Time value the user configured in the “hello timeout” parameter for the targeted session. The LSR node will then start advertising an exponentially increasing Hold-Time value in the Hello message as soon as the targeted LDP session to the peer is up. Each new incremented Hold-Time value is sent in a number of Hello messages equal to the value of the argument <i>factor</i>, which represents the dampening factor, before the next exponential value is advertised. This provides time for the two peers to settle on the new value. When the Hold-Time reaches the maximum value of 0xffff (binary 65535), the two peers will send Hello messages at a frequency of every $[(65535-1)/\text{local helloFactor}]$ seconds for the lifetime of the targeted-LDP session (for example, if the local Hello Factor is three (3), then Hello messages will be sent every 21844 seconds.</p> <p>The LSR node continues to compute the frequency of sending the Hello messages based on the minimum of its local Hold-time value and the one advertised by its peer as in RFC 5036. Thus for the targeted LDP session to suppress the periodic Hello messages, both peers must bring their advertised Hold-Time to the maximum value. If one of the LDP peers does not, the frequency of the Hello messages sent by both peers will continue to be governed by the smaller of the two Hold-Time values.</p> <p>When the user enables the hello reduction option on the LSR node while the targeted LDP session to the peer is operationally up, the change will take effect immediately. In other words, the LSR node will start advertising an exponentially increasing Hold-Time value in the Hello message, starting with the current configured Hold-Time value.</p> <p>When the user disables the hello reduction option while the targeted LDP session to the peer is operationally up, the change in the Hold-Time from 0xffff (binary 65535) to the user configured value for this peer will take effect immediately. The local LSR will immediately advertise the value of the user configured Hold-Time value and will not wait until the next scheduled time to send a Hello to make sure the peer adjusts its local hold timeout value immediately.</p> <p>In general, any configuration change to the parameters of the T-LDP Hello adjacency (i.e., modifying the hello adjacency Hello Timeout or factor, enabling/disabling hello reduction, or modifying hello reduction factor) will cause the LSR node to trigger immediately an updated Hello message with the updated Hold Time value without waiting for the next scheduled time to send a Hello.</p> <p>The no form of this command disables the hello reduction feature.</p>
Default	no hello-reduction
Parameters	<i>factor</i> — Specifies the integer that specifies the Hello reduction dampening factor.
Values	3 —20

interface

Syntax	[no] interface <i>ip-int-name</i>
Context	config>router>ldp>if-params
Description	<p>This command enables LDP on the specified IP interface.</p> <p>The no form of the command deletes the LDP interface and all configuration information associated with the LDP interface.</p> <p>The LDP interface must be disabled using the shutdown command before it can be deleted.</p>
Parameters	<i>ip-int-name</i> — The name of an existing interface. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

transport-address

Syntax	transport-address { interface system } no transport-address
Context	config>router>ldp>if-params config>router>ldp>if-params>if
Description	<p>This command configures the transport address to be used when setting up the LDP TCP sessions. The transport address can be configured as interface or system. The transport address can be configured globally (applies to all LDP interfaces) or per interface. The most specific value is used.</p> <p>With the transport-address command, you can set up the LDP interface to the connection which can be set to the interface address or the system address. However, there can be an issue of which address to use when there are parallel adjacencies. This situation can not only happen with parallel links, it could be a link and a targeted adjacency since targeted adjacencies request the session to be set up only to the system IP address.</p> <p>Note that the transport-address value should not be interface if multiple interfaces exist between two LDP neighbors. Depending on the first adjacency to be formed, the TCP endpoint is chosen. In other words, if one LDP interface is set up as transport-address interface and another for transport-address system, then, depending on which adjacency was set up first, the TCP endpoint addresses are determined. After that, because the hello contains the LSR ID, the LDP session can be checked to verify that it is set up and then match the adjacency to the session.</p> <p>Note that for any given ILDP interface, as the local-lsr-id parameters is changed to interface, the transport-address configuration loses effectiveness. Since it will be ignored and the ILDP session will <i>always</i> use the relevant interface IP address as transport-address even though system is chosen.</p>

The **no** form of the command, at the global level, sets the transport address to the default value. The **no** form of the command, at the interface level, sets the transport address to the value defined under the global level.

Default **system** — The system IP address is used.

Parameters **interface** — The IP interface address is used to set up the LDP session between neighbors. The transport address interface cannot be used if multiple interfaces exist between two neighbors, since only one LDP session is set up between two neighbors.

system — The system IP address is used to set up the LDP session between neighbors.

multicast-traffic

Syntax **[no] multicast-traffic**

Context config>router>ldp>interface-parameters>interface

Description This command enables P2MP multicast traffic forwarding on the interface.

The **no** form of command disables P2MP LDP multicast traffic on the interface. P2MP tree branching out on the interface would not withdraw label map from the peer session on interface shutdown or multicast traffic is disabled. Session may exist on multiple parallel interfaces. Only forwarding entry is changed when interface is shutdown or multicast traffic support is disabled.

Note that LDP may choose to egress the mLDP tree over this interface, but if this it is enabled (multicast-traffic disabled), the dataplane will not forward traffic on this branch.

Default multicast-traffic enable

mp-mbb-time

Syntax **[no] mp-mbb-time**

Context config>router>ldp

Description This command configures the maximum time a P2MP transit/bud node must wait before switching over to the new path if the new node does not send MBB TLV to inform of the availability of data plane.

The **no** form of command should configure the default timer of 3 seconds.

Default 3 seconds

Parameters *interval* — seconds.

Values 1-10 seconds

Peer Parameters Commands

peer-parameters

Syntax	peer-parameters
Context	config>router>ldp
Description	This command enables the context to configure peer specific parameters.

peer

Syntax	[no] peer <i>ip-address</i>
Context	config>router>ldp>peer-parameters
Description	This command configures parameters for an LDP peer.
Default	none
Parameters	<i>ip-addr</i> — The IP address of the LDP peer in dotted decimal notation.

auth-keychain

Syntax	auth-keychain <i>name</i>
Context	config>router>ldp>peer-parameters>peer
Description	This command configures TCP authentication keychain to use for the session.
Parameters	<i>name</i> — Specifies the name of the keychain to use for the specified TCP session or sessions. This keychain allows the rollover of authentication keys during the lifetime of a session up to 32 characters in length. Peer address has to be the TCP session transport address.

authentication-key

Syntax	authentication-key [<i>authentication-key</i> <i>hash-key</i>] [hash hash2] no authentication-key
Context	config>router>ldp>peer-parameters>peer
Description	This command specifies the authentication key to be used between LDP peers before establishing sessions. Authentication uses the MD-5 message-based digest. Peer address has to be the TCP session transport address. The no form of this command disables authentication.

Default	none
Parameters	<p><i>authentication-key</i> — The authentication key. The key can be any combination of ASCII characters up to 16 characters in length (unencrypted). If spaces are used in the string, enclose the entire string in quotation marks (“ ”).</p> <p><i>hash-key</i> — The hash key. The key can be any combination of up to 33 alphanumeric characters. If spaces are used in the string, enclose the entire string in quotation marks (“ ”).</p> <p>This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.</p> <p>hash — Specifies the key is entered in an encrypted form. If the hash keyword is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the hash parameter specified.</p> <p>hash2 — Specifies the key is entered in a more complex encrypted form. If the hash2 parameter is not used, the less encrypted hash form is assumed.</p>

dod-label-distribution

Syntax	[no] dod-label-distribution
Context	config>router>ldp>peer-parameters>peer
Description	<p>This command enables the use of the LDP Downstream-on-Demand (DoD) label distribution procedures.</p> <p>When this option is enabled, LDP will set the A-bit in the Label Initialization message when the LDP session to the peer is established. When both peers set the A-bit, they will both use the DoD label distribution method over the LDP session [rfc5036].</p> <p>This feature can only be enabled on a link-level LDP session and therefore will apply to prefix labels only, not service labels.</p> <p>As soon as the link LDP session comes up, the 7x50 will send a label request to its DoD peer for the FEC prefix corresponding to the peer’s LSR-id. The DoD peer LSR-id is found in the basic Hello discovery messages the peer used to establish the Hello adjacency with the 7x50.</p> <p>Similarly if the 7x50 and the directly attached DoD peer entered into extended discovery and established a targeted LDP session, the 7x50 will immediately send a label request for the FEC prefix corresponding to the peer’s LSR-id found in the extended discovery messages.</p> <p>However, the 7x50 node will not advertise any <FEC, label> bindings, including the FEC of its own LSR-id, unless the DoD peer requested it using a Label Request Message.</p> <p>When the DoD peer sends a label request for any FEC prefix, the 7x50 will reply with a <FEC, label> binding for that prefix if the FEC was already activated on the 7x50. If not, the 7x50 replies with a notification message containing the status code of “no route.” The 7x50 will not attempt in the latter case to send a label request to the next-hop for the FEC prefix when the LDP session to this next-hop uses the DoD label distribution mode. Hence the reference to single-hop LDP DoD procedures.</p> <p>As soon as the link LDP session comes up, the 7x50 will send a label request to its DoD peer for the FEC prefix corresponding to the peer’s LSR-id. The DoD peer LSR-id is found in the basic Hello discovery messages the peer used to establish the Hello adjacency with the 7x50.</p>

Similarly if the 7x50 and the directly attached DoD peer entered into extended discovery and established a targeted LDP session, the 7x50 will immediately send a label request for the FEC prefix corresponding to the peer's LSR-id found in the extended discovery messages. Peer address has to be the peer LSR-ID address.

The **no** form of this command disables the DoD label distribution with an LDP neighbor.

Default no dod-label-distribution

export-prefixes

Syntax [no] export-prefixes *policy-name*

Context config>router>ldp>peer-parameters>peer

Description This command specifies the export route policy used to determine which prefixes received from other LDP and T-LDP peers are re-distributed to this LDP peer via the LDP/T-LDP session to this peer. A prefix that is filtered out (deny) will not be exported. A prefix that is filtered in (accept) will be exported.

If no export policy is specified, all FEC prefixes learned will be exported to this LDP peer. This policy is applied in addition to the global LDP policy and targeted session policy.

Policies are configured in the **config>router>policy-options** context. A maximum of five policy names can be specified. Peer address has to be the peer LSR-ID address.

The **no** form of the command removes the policy from the configuration.

Default no export-prefixes - no export route policy is specified

Parameters *policy-name* — The export-prefix route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains spaces, use double quotes to delimit the start and end of the string. The specified name(s) must already be defined.

fec129-cisco-interop

Syntax [no] fec129-cisco-interop

Context config>router>ldp>peer-parameters>peer

Description This command specifies whether LDP will provide translation between non-compliant FEC 129 formats of Cisco. Peer LDP sessions must be manually configured towards the non-compliant Cisco PEs.

When enabled, Cisco non-compliant format will be used to send and interpret received label release messages i.e. the FEC129 SAII and TAI fields will be reversed.

When the disabled, Cisco non-compliant format will not be used or supported. Peer address has to be the peer LSR-ID address.

The **no** form of the command returns the default .

Default no fec129-cisco-interop

import-prefixes

Syntax	[no] import-prefixes <i>policy-name</i>
Context	config>router>ldp>peer-parameters>peer
Description	<p>This command configures the import FEC prefix policy to determine which prefixes received from this LDP peer are imported and installed by LDP on this node. If resolved these FEC prefixes are then re-distributed to other LDP and T-LDP peers. A FEC prefix that is filtered out (deny) will not be imported. A FEC prefix that is filtered in (accept) will be imported.</p> <p>If no import policy is specified, the node will import all prefixes received from this LDP/T-LDP peer. This policy is applied in addition to the global LDP policy and targeted session policy.</p> <p>Policies are configured in the config>router>policy-options context. A maximum of five policy names can be specified. Peer address has to be the peer LSR-ID address.</p> <p>The no form of the command removes the policy from the configuration.</p>
Default	no import-prefixes - no import route policy is specified
Parameters	<i>policy-name</i> — The import-prefix route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains spaces, use double quotes to delimit the start and end of the string. The specified name(s) must already be defined

path-mtu-discovery

Syntax	path-mtu-discovery no path-mtu-discovery
Context	config>router>ldp>peer-parameters>peer
Description	<p>This command enables Path MTU discovery for the associated TCP connections. When enabled, the MTU for the associated TCP session is initially set to the egress interface MTU. The DF bit is also set so that if a router along the path of the TCP connection cannot handle a packet of a particular size without fragmenting, it sends back an ICMP message to set the path MTU for the given session to a lower value that can be forwarded without fragmenting.</p>
Default	no path-mtu-discovery

pe-id-mac-flush-interop

Syntax	[no] pe-id-mac-flush-interop
Context	config>router>ldp>peer-parameters>peer
Description	<p>This command enables the addition of the PE-ID TLV in the LDP MAC withdrawal (mac-flush) message, under certain conditions, and modifies the mac-flush behavior for interoperability with other vendors that do not support the flush-all-from-me vendor-specific TLV. This flag can be enabled on a per LDP peer basis and allows the flush-all-from-me interoperability with other vendors. When the pe-id-mac-flush-interop flag is enabled for a given peer, the current mac-flush behavior is modified in terms of mac-flush generation, mac-flush propagation and behavior upon receiving a mac-flush.</p>

The mac-flush generation will be changed depending on the type of event and according to the following rules:

- Any all-from-me mac-flush event will trigger a mac-flush all-but-mine message (RFC 4762 compliant format) with the addition of a PE-ID TLV. The PE-ID TLV contains the IP address of the sending PE.
- Any all-but-mine mac-flush event will trigger a mac-flush all-but-mine message WITHOUT the addition of the PE-ID TLV, as long as the source spoke-sdp is not part of an end-point.
- Any all-but-mine mac-flush event will trigger a mac-flush all-but-mine message WITH the addition of the PE-ID TLV, if the source spoke-sdp is part of an end-point and the spoke-sdp goes from down/standby state to active state. In this case, the PE-ID TLV will contain the IP address of the PE to which the previous active spoke-sdp was connected to.

Any other case will follow the existing mac-flush procedures.

When the pe-id-mac-flush-interop flag is enabled for a given LDP peer, the mac-flush ingress processing is modified according to the following rules:

- Any received all-from-me mac-flush will follow the existing mac-flush all-from-me rules regardless of the existence of the PE-ID.
- Any received all-but-mine mac-flush will take into account the received PE-ID, i.e. all the mac addresses associated to the PE-ID will be flushed. If the PE-ID is not included, the mac addresses associated to the sending PE will be flushed.
- Any other case will follow the existing mac-flush procedures.

When a mac-flush message has to be propagated (for an ingress sdp-binding to an egress sdp-binding) and the pe-id-mac-flush-interop flag is enabled for the ingress and egress TLDP peers, the following behavior is observed:

- If the ingress and egress bindings are spoke-sdp, the PE will propagate the mac-flush message with its own PE-ID.
- If the ingress binding is an spoke-sdp and the egress binding a mesh-sdp, the PE will propagate the mac-flush message without modifying the PE-ID included in the PE-ID TLV.
- If the ingress binding is a mesh-sdp and the egress binding an spoke-sdp, the PE will propagate the mac-flush message with its own PE-ID.
- When ingress and egress bindings are mesh-sdp, the mac-flush message is never propagated. This is the behavior regardless of the pe-id-mac-flush-interop flag configuration.

Note that the PE-ID TLV is never added when generating a mac-flush message on a B-VPLS if the send-bvpls-flush command is enabled in the I-VPLS. In the same way, no PE-ID is added when propagating mac-flush from a B-VPLS to a I-VPLS when the propagate-mac-flush-from-bvpls command is enabled. Mac-flush messages for peers within the same I-VPLS or within the same B-VPLS domain follow the procedures described above.

Default no pe-id-mac-flush-interop

ttl-security

Syntax **ttl-security** *min-ttl-value*
no ttl-security

Context config>router>ldp>peer-parameters>peer

Description This command configures TTL security parameters for incoming packets. When the feature is enabled, BGP/LDP will accept incoming IP packets from a peer only if the TTL value in the packet is greater than or equal to the minimum TTL value configured for that peer. Peer address has to be the TCP session transport address.

The **no** form of the command disables TTL security.

Default no ttl-security

Parameters *min-ttl-value* — Specify the minimum TTL value for an incoming packet.

Values 1 — 255

prefer-tunnel-in-tunnel

Syntax [**no**] **prefer-tunnel-in-tunnel**

Context config>router>ldp

Description This command specifies to use tunnel-in-tunnel over a simple LDP tunnel. Specifically, the user packets for LDP FECs learned over this targeted LDP session can be sent inside an RSVP LSP which terminates on the same egress router as the destination of the targeted LDP session. The user can specify an explicit list of RSVP LSP tunnels under the Targeted LDP session or LDP will perform a lookup in the Tunnel Table Manager (TTM) for the best RSVP LSP. In the former case, only the specified LSPs will be considered to tunnel LDP user packets. In the latter case, all LSPs available to the TTM and which terminate on the same egress router as this targeted LDP session will be considered. In both cases, the metric specified under the LSP configuration is used to control this selection.

Note that the lookup in the TTM will prefer a LDP tunnel over an LDP-over-RSVP tunnel if both are available. Also note that the tunneling operates on the dataplane only. Control packets of this targeted LDP session are sent over the IGP path.

shortcut-transit-ttl-propagate

Syntax	[no] shortcut-transit-ttl-propagate
Context	config>router>ldp config>router>mpls
Description	<p>This command configures the TTL handling of transit packets for all LSP shortcuts originating on this ingress LER. It applies to all LDP or RSVP LSPs that are used to resolve static routes, BGP routes, and IGP routes.</p> <p>The user can enable or disable the propagation of the TTL from the header of an IP packet into the header of the resulting MPLS packet independently for local and transit packets forwarded over an LSP shortcut.</p> <p>By default, the feature propagates the TTL from the header of transit IP packets into the label stack of the resulting MPLS packets forwarded over the LSP shortcut. This is referred to as Uniform mode.</p> <p>When the no form of the command is enabled, TTL propagation is disabled on all transit IP packets received on any IES interface and destined to a route that is resolved to the LSP shortcut. In this case, a TTL of 255 is programmed onto the pushed label stack. This is referred to as Pipe mode.</p>
Default	shortcut-transit-ttl-propagate

shortcut-local-ttl-propagate

Syntax	[no] shortcut-local-ttl-propagate
Context	config>router>ldp config>router>mpls
Description	<p>This command configures the TTL handling of locally generated packets for all LSP shortcuts originating on this ingress LER. It applies to all LDP or RSVP LSPs that are used to resolve static routes, BGP routes, and IGP routes.</p> <p>The user can enable or disable the propagation of the TTL from the header of an IP packet into the header of the resulting MPLS packet independently for local and transit packets forwarded over an LSP shortcut.</p> <p>Local IP packets include ICMP Ping, traceroute, and OAM packets, that are destined to a route that is resolved to the LSP shortcut. Transit IP packets are all IP packets received on any IES interface and destined to a route that is resolved to the LSP shortcut</p> <p>By default, the feature propagates the TTL from the header of locally generated IP packets into the label stack of the resulting MPLS packets forwarded over the LSP shortcut. This is referred to as Uniform mode.</p> <p>When the no form of the above command is enabled, TTL propagation is disabled on all locally generated IP packets, including ICMP Ping, traceroute, and OAM packets, that are destined to a route that is resolved to the LSP shortcut. In this case, a TTL of 255 is programmed onto the pushed label stack. This is referred to as Pipe mode.</p>
Default	shortcut-local-ttl-propagate

Targeted Session Commands

targeted-session

Syntax	targeted-session
Context	config>router>ldp
Description	<p>This command configures targeted LDP sessions. Targeted sessions are LDP sessions between non-directly connected peers. Hello messages are sent directly to the peer platform instead of to all the routers on this subnet multicast address.</p> <p>The discovery messages for an indirect LDP session are addressed to the specified peer and not to the multicast address.</p>
Default	none

bfd-enable

Syntax	[no] bfd-enable
Context	config>router>ldp>targ-session>peer
Description	<p>This command enables the bidirectional forwarding detection (BFD) session for the selected TLDP session. By enabling BFD for a selected targeted session, the state of that session is tied to the state of the underneath BFD session between the two nodes.</p> <p>The parameters used for the BFD are set via the BFD command under the IP interface.</p> <p>The no form of this command removes the TLDP session operational state binding to the central BFD session one.</p>
Default	no bfd-enable

disable-targeted-session

Syntax	[no] disable-targeted-session
Context	config>router>ldp>targ-session
Description	This command disables support for SDP triggered automatic generated targeted sessions. Targeted sessions are LDP sessions between non-directly connected peers. The discovery messages for an indirect LDP session are addressed to the specified peer and not to the multicast address. The no form of the command enables the set up of any targeted sessions.
Default	no disable-targeted-session

peer

Syntax	[no] peer <i>ip-address</i>
Context	config>router>ldp>targeted-session
Description	This command configures parameters for an LDP peer.
Default	none
Parameters	<i>ip-address</i> — The IP address of the LDP peer in dotted decimal notation.

peer-template-map

Syntax	peer-template-map peer-template <i>template-name</i> policy <i>peer-prefix-policy1</i> [<i>peer-prefix-policy2..up to 5</i>] no peer-template-map peer-template <i>template-name</i>
Context	config>router>ldp>targeted-session
Description	This command enables the automatic creation of a targeted Hello adjacency and LDP session to a discovered peer. The user configures a targeted session peer parameter template and binds it to a peer prefix policy. Each application of a targeted session template to a given prefix in the prefix list will result in the establishment of a targeted Hello adjacency to an LDP peer using the template parameters as long as the prefix corresponds to a router-id for a node in the TE database. As a result of this, the user must enable the traffic-engineering option in ISIS or OSPF. The targeted Hello adjacency will either trigger a new LDP session or will be associated with an existing LDP session to that peer. Up to 5 peer prefix policies can be associated with a single peer template at all times. Also, the user can associate multiple templates with the same or different peer prefix policies. Thus multiple templates can match with a given peer prefix. In all cases, the targeted session parameters applied to a given peer prefix are taken from the first created template by the user. This provides a more deterministic behavior regardless of the order in which the templates are associated with the prefix policies.

Each time the user executes the above command, with the same or different prefix policy associations, or the user changes a prefix policy associated with a targeted peer template, the system re-evaluates the prefix policy. The outcome of the re-evaluation will tell LDP if an existing targeted Hello adjacency needs to be torn down or if an existing targeted Hello adjacency needs to have its parameters updated on the fly.

If a /32 prefix is added to (removed from) or if a prefix range is expanded (shrunk) in a prefix list associated with a targeted peer template, the same prefix policy re-evaluation described above is performed.

The template comes up in the **no shutdown** state and as such it takes effect immediately. Once a template is in use, the user can change any of the parameters on the fly without shutting down the template. In this case, all targeted Hello adjacencies are updated.

The SR OS supports multiple ways of establishing a targeted Hello adjacency to a peer LSR:

- User configuration of the peer with the targeted session parameters inherited from the **config>router>ldp>targeted-session** in the top level context or explicitly configured for this peer in the **config>router>ldp>targeted-session>peer** context and which overrides the top level parameters shared by all targeted peers. Let us refer to the top level configuration context as the global context. Note that some parameters only exist in the global context and as such their value will always be inherited by all targeted peers regardless of which event triggered it.
- User configuration of an SDP of any type to a peer with the signaling tldp option enabled (default configuration). In this case the targeted session parameter values are taken from the global context.
- User configuration of a (FEC 129) PW template binding in a BGP-VPLS service. In this case the targeted session parameter values are taken from the global context.
- User configuration of a (FEC 129 type II) PW template binding in a VLL service (dynamic multi-segment PW). In this case the target session parameter values are taken from the global context.
- User configuration of a mapping of a targeted session peer parameter template to a prefix policy when the peer address exists in the TE database (this feature). In this case, the targeted session parameter values are taken from the template.

Since the above triggering events can occur simultaneously or in any arbitrary order, the LDP code implements a priority handling mechanism in order to decide which event overrides the active targeted session parameters. The overriding trigger will become the owner of the targeted adjacency to a given peer. The following is the priority order:

- Priority 1: manual configuration of peer parameters
- Priority 2: mapping of targeted session template to prefix policy.
- Priority 3: manual configuration of SDP, PW template binding in BGP-AD VPLS and in FEC 129 VLL.

Note that any parameter value change to an active targeted Hello adjacency caused by any of the above triggering events is performed on the fly by having LDP immediately send a Hello message with the new parameters to the peer without waiting for the next scheduled time for the Hello message. This allows the peer to adjust its local state machine immediately and maintains both the Hello adjacency and the LDP session in UP state. The only exceptions are the following:

- The triggering event caused a change to the local-lsr-id parameter value. In this case, the Hello adjacency is brought down which will also cause the LDP session to be brought down if this is

the last Hello adjacency associated with the session. A new Hello adjacency and LDP session will then get established to the peer using the new value of the local LSR ID.

- The triggering event caused the targeted peer shutdown option to be enabled. In this case, the Hello adjacency is brought down which will also cause the LDP session to be brought down if this is the last Hello adjacency associated with the session.

Finally, the value of any LDP parameter which is specific to the LDP/TCP session to a peer is inherited from the **config>router>ldp>peer-parameters>peer** context. This includes MD5 authentication, LDP prefix per-peer policies, label distribution mode (DU or DOD), etc.

The no form of this command deletes the binding of the template to the peer prefix list and brings down all Hello adjacencies to the discovered LDP peers.

peer-template

Syntax	[no] peer-template template-name
Context	config>router>ldp>targeted-session
Description	<p>This command creates a targeted session peer parameter template that can be referenced in the automatic creation of targeted Hello adjacency and LDP session to a discovered peer.</p> <p>The no form of command deletes the peer template. A peer template cannot be deleted if it is bound to a peer prefix list.</p>
Parameters	<i>template-name</i> — Specifies the template name to identify targeted peer template. It must be 32 characters maximum.

export-prefixes

Syntax	export-prefixes policy-name [policy-name...(up to 5 max)] no export-prefixes
Context	config>router>ldp>targeted-session
Description	<p>This command specifies the export route policy used to determine which FEC prefix label bindings are exported from a targeted LDP session. A route that is filtered out (deny) will not be exported. A route that is filtered in (accept) will be exported.</p> <p>If no export policy is specified, all bindings learned through a targeted LDP session will be exported to all targeted LDP peers. This policy is applied in addition to the global LDP policy.</p> <p>Policies are configured in the config>router>policy-options context. A maximum of five policy names can be specified.</p> <p>The no form of the command removes the policy from the configuration.</p>
Parameters	<i>policy-name</i> — The export policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

import-prefixes

Syntax	import-prefixes <i>policy-name</i> [<i>policy-name...</i> (up to 5 max)] no import-prefixes
Context	config>router>ldp>targeted-session
Description	<p>This command configures the import route policy to determine which FEC prefix label bindings are accepted from targeted LDP neighbors into this node. A label binding that is filtered out (deny) will not be imported. A route that is filtered in (accept) will be imported.</p> <p>If no import policy is specified, this node session will accept all bindings from configured targeted LDP neighbors. This policy is applied in addition to the global LDP policy.</p> <p>Policies are configured in the config>router>policy-options context. A maximum of five policy names can be specified.</p> <p>The no form of the command removes the policy from the configuration.</p>
Parameters	<i>policy-name</i> — The import policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

Show LDP Commands

adv-adj-addr-only

Syntax	adv-adj-addr-only
Context	show>router>ldp
Description	This command displays the local interfaces used to establish hello adjacencies with an LDP peer.
Output	<pre> *A:SR4>config>router>ldp>peer-params# peer 110.20.1.1 adv-adj-addr-only *A:SR4>config>router>ldp>peer-params# *A:SR4>config>router>ldp>peer-params# show router ldp session - session [<ip-addr[:label-space]>] [session-type] [state <state>] [summary detail] - session [<ip-addr[:label-space]>] local-addresses [sent recv] [ip-addr <ip- address>] - session [<ip-addr[:label-space]>] statistics [packet-type] [session-type] <ip-addr[:label-sp*> : ip-addr - a.b.c.d label-space - [0..65535] <statistics> : keyword - display statistics <packet-type> : hello keepalive init label notification address - keywords <session-type> : link targeted both <state> : up - Established down - Initialized, OpenRecv, OpenSent, Nonexistent <summary detail> : summary detail <local-addresses> : keyword <sent recv> : keyword <ip-address> : a.b.c.d *A:SR4>config>router>ldp>peer-params# show router ldp session 110.20.1.1 local- addresses ===== LDP Session Local-Addresses ===== ----- Session with Peer 110.20.1.1:0, Local 110.20.1.3:0 ----- Sent Addresses: 6.2.6.3 6.2.7.3 6.2.49.3 6.2.57.3 110.20.1.3 200.0.0.3 Recv Addresses: ===== *A:SR4>config>router>ldp>peer-params# peer 110.20.1.1 no adv-adj-addr-only *A:SR4>config>router>ldp>peer-params# show router ldp session 110.20.1.1 local- addresses ===== LDP Session Local-Addresses ===== ----- Session with Peer 110.20.1.1:0, Local 110.20.1.3:0 ----- Sent Addresses: 1.2.11.3 1.2.12.3 1.2.111.3 1.2.121.3 6.2.6.3 6.2.7.3 6.2.49.3 6.2.50.3 6.2.57.3 6.2.58.3 6.2.61.3 6.2.71.3 10.100.40.3 34.34.34.34 35.35.35.35 60.60.60.60 71.17.71.1 88.88.87.1 88.88.88.1 88.88.89.1 104.104.0.3 104.104.1.3 104.104.2.3 104.104.3.3 </pre>

Show Commands

```
104.104.4.3 104.104.5.3 104.104.6.3 104.104.7.3
104.104.8.3 104.104.9.3 104.104.10.3 104.104.11.3
104.104.12.3 104.104.13.3 104.104.20.3 104.104.21.3
104.104.22.3 104.104.23.3 104.104.24.3 104.104.25.3
104.104.26.3 104.104.27.3 104.104.28.3 104.104.29.3
104.104.30.3 104.104.31.3 104.104.32.3 104.104.33.3
104.104.40.3 104.104.41.3 104.104.42.3 104.104.43.3
104.104.44.3 104.104.45.3 104.104.46.3 104.104.47.3
104.104.48.3 104.104.49.3 104.104.50.3 104.104.51.3
104.104.52.3 104.104.53.3 104.104.60.3 104.104.61.3
104.104.62.3 104.104.63.3 104.104.64.3 104.104.65.3
104.104.66.3 104.104.67.3 104.104.68.3 104.104.69.3
104.104.70.3 104.104.71.3 104.104.72.3 104.104.73.3
110.3.1.51 110.3.2.51 110.3.3.51 110.3.4.51
110.3.5.51 110.3.6.51 110.3.7.51 110.3.8.51
110.3.9.51 110.3.10.51 110.3.11.51 110.3.12.51
110.3.13.51 110.3.14.51 110.3.15.51 110.3.16.51
110.3.17.51 110.3.18.51 110.3.19.51 110.3.20.51
110.3.21.51 110.3.22.51 110.3.23.51 110.3.24.51
110.3.25.51 110.3.26.51 110.3.27.51 110.3.28.51
110.3.29.51 110.3.30.51 110.3.31.51 110.3.32.51
110.3.33.51 110.3.34.51 110.3.35.51 110.3.36.51
110.3.37.51 110.3.38.51 110.3.39.51 110.3.40.51
110.3.41.51 110.3.42.51 110.3.43.51 110.3.44.51
110.3.45.51 110.3.46.51 110.3.47.51 110.3.48.51
110.3.49.51 110.3.50.51 110.3.51.51 110.3.52.51
110.3.53.51 110.3.54.51 110.3.55.51 110.3.56.51
110.3.57.51 110.3.58.51 110.3.59.51 110.3.60.51
110.3.61.51 110.3.62.51 110.3.63.51 110.3.64.51
110.3.65.51 110.3.66.51 110.3.67.51 110.3.68.51
110.3.69.51 110.3.70.51 110.3.71.51 110.3.72.51
110.3.73.51 110.3.74.51 110.3.75.51 110.3.76.51
110.3.77.51 110.3.78.51 110.3.79.51 110.3.80.51
110.3.81.51 110.3.82.51 110.3.83.51 110.3.84.51
110.3.85.51 110.3.86.51 110.3.87.51 110.3.88.51
110.3.89.51 110.3.90.51 110.3.91.51 110.3.92.51
110.3.93.51 110.3.94.51 110.3.95.51 110.3.96.51
110.3.97.51 110.3.98.51 110.3.99.51 110.3.100.51
110.3.101.51 110.3.102.51 110.3.103.51 110.3.104.51
110.3.105.51 110.3.106.51 110.3.107.51 110.3.108.51
110.3.109.51 110.3.110.51 110.3.111.51 110.3.112.51
110.3.113.51 110.3.114.51 110.3.115.51 110.3.116.51
110.3.117.51 110.3.118.51 110.3.119.51 110.3.120.51
110.3.121.51 110.3.122.51 110.3.123.51 110.3.124.51
110.3.125.51 110.3.126.51 110.3.127.51 110.3.128.51
110.3.129.51 110.3.130.51 110.3.131.51 110.3.132.51
110.3.133.51 110.3.134.51 110.3.135.51 110.3.136.51
110.3.137.51 110.3.138.51 110.3.139.51 110.3.140.51
110.3.141.51 110.3.142.51 110.3.143.51 110.3.144.51
110.3.145.51 110.3.146.51 110.3.147.51 110.3.148.51
110.3.149.51 110.3.150.51 110.3.151.51 110.3.152.51
110.3.153.51 110.3.154.51 110.3.155.51 110.3.156.51
110.3.157.51 110.3.158.51 110.3.159.51 110.3.160.51
110.3.161.51 110.3.162.51 110.3.163.51 110.3.164.51
110.3.165.51 110.3.166.51 110.3.167.51 110.3.168.51
110.3.169.51 110.3.170.51 110.3.171.51 110.3.172.51
110.3.173.51 110.3.174.51 110.3.175.51 110.3.176.51
110.3.177.51 110.3.178.51 110.3.179.51 110.3.180.51
110.3.181.51 110.3.182.51 110.3.183.51 110.3.184.51
110.3.185.51 110.3.186.51 110.3.187.51 110.3.188.51
110.3.189.51 110.20.1.3 110.20.1.51 110.20.3.1
```



```

110.20.3.2 110.20.3.3 110.20.3.4 110.20.3.5
110.20.3.6 110.20.3.7 110.20.3.8 110.20.3.9
110.20.3.10 110.20.3.11 110.20.3.12 110.20.3.13
110.20.3.14 110.20.3.15 110.20.3.16 110.20.3.17
110.20.3.18 110.20.3.19 110.20.3.20 110.20.3.21
110.20.3.22 110.20.3.23 110.20.3.24 110.20.3.25
110.20.3.26 110.20.3.27 110.20.3.28 110.20.3.29
110.20.3.30 110.20.3.31 150.50.0.3 150.50.1.3
150.50.2.3 150.50.3.3 150.50.4.3 150.50.5.3
150.50.6.3 150.50.7.3 150.50.8.3 150.50.9.3
150.50.10.3 150.50.11.3 150.50.12.3 150.50.13.3
150.50.20.3 150.50.21.3 150.50.22.3 150.50.23.3
150.50.24.3 150.50.25.3 150.50.26.3 150.50.27.3
150.50.28.3 150.50.29.3 150.50.30.3 150.50.31.3
150.50.32.3 150.50.33.3 150.50.40.3 150.50.41.3
150.50.42.3 150.50.43.3 150.50.44.3 150.50.45.3
150.50.46.3 150.50.47.3 150.50.48.3 150.50.49.3
150.50.50.3 150.50.51.3 150.50.52.3 150.50.53.3
150.50.60.3 150.50.61.3 150.50.62.3 150.50.63.3
150.50.64.3 150.50.65.3 150.50.66.3 150.50.67.3
150.50.68.3 150.50.69.3 150.50.70.3 150.50.71.3
150.50.72.3 150.50.73.3 150.60.30.3 150.60.31.3
150.60.31.10 150.60.31.19 150.60.31.27 150.60.31.35
150.60.31.43 150.60.31.51 150.60.31.59 150.60.31.67
150.60.31.75 150.60.31.83 150.60.31.98 150.60.31.106
150.60.70.3 150.60.71.3 150.60.71.10 150.60.71.19
150.60.71.27 150.60.71.35 150.60.71.43 150.60.71.51
150.60.71.59 150.60.71.67 150.60.71.75 150.60.71.83
150.60.71.98 150.60.71.106 150.60.75.3 150.60.76.3
150.60.76.10 150.60.76.19 150.60.76.27 150.60.76.35
150.60.76.43 150.60.76.51 150.60.76.59 150.60.76.67
150.60.76.75 150.60.76.83 150.60.76.98 150.60.76.106
170.70.90.3 170.70.91.3 180.60.100.3 180.60.110.3
180.100.3.3 193.127.0.1 200.0.0.3 203.0.0.3
Recv Addresses: 40.40.1.1 40.40.2.1
=====
*A:SR4>config>router>ldp>peer-params#
*A:SR4>config>router>ldp>peer-params# peer 110.20.1.1 adv-adj-addr-only
*A:SR4>config>router>ldp>peer-params# show router ldp session 110.20.1.1 local-
addresses sent
=====
LDP Session Local-Addresses
=====
-----
Session with Peer 110.20.1.1:0, Local 110.20.1.3:0
-----
Sent Addresses: 6.2.6.3 6.2.7.3 6.2.49.3 6.2.57.3
110.20.1.3 200.0.0.3
=====
*A:SR4>config>router>ldp>peer-params# show router ldp session 110.20.1.1 local-
addresses recv
=====
LDP Session Local-Addresses
=====
-----
Session with Peer 110.20.1.1:0, Local 110.20.1.3:0
-----
Recv Addresses: 6.2.6.1 6.2.7.1 6.2.49.1 6.2.50.1
6.2.57.1 6.2.58.1 6.2.61.1 6.2.71.1
7.1.1.1 7.1.2.1 7.1.3.1 7.1.4.1
7.1.5.1 7.1.6.1 7.1.7.1 7.1.8.1

```

Show Commands

```
7.1.9.1 7.1.10.1 7.1.11.1 7.1.12.1
7.1.13.1 7.1.14.1 7.1.15.1 7.1.16.1
7.1.17.1 7.1.18.1 7.1.19.1 7.1.20.1
7.1.21.1 7.1.22.1 7.1.23.1 7.1.24.1
7.1.25.1 7.1.26.1 7.1.27.1 7.1.28.1
7.1.29.1 7.1.30.1 7.1.31.1 7.1.32.1
7.1.33.1 7.1.34.1 7.1.35.1 7.1.36.1
7.1.37.1 7.1.38.1 7.1.39.1 7.1.40.1
7.1.41.1 7.1.42.1 7.1.43.1 7.1.44.1
7.1.45.1 7.1.46.1 7.1.47.1 7.1.48.1
7.1.49.1 7.1.50.1 7.1.51.1 7.1.52.1
7.1.53.1 7.1.54.1 7.1.55.1 7.1.56.1
7.1.57.1 7.1.58.1 7.1.59.1 7.1.60.1
7.1.61.1 7.1.62.1 7.1.63.1 7.1.64.1
7.1.65.1 7.1.66.1 7.1.67.1 7.1.68.1
7.1.69.1 7.1.70.1 7.1.71.1 7.1.72.1
7.1.73.1 7.1.74.1 7.1.75.1 7.1.76.1
7.1.77.1 7.1.78.1 7.1.79.1 7.1.80.1
7.1.81.1 7.1.82.1 7.1.83.1 7.1.84.1
7.1.85.1 7.1.86.1 7.1.87.1 7.1.88.1
7.1.89.1 7.1.90.1 7.1.91.1 7.1.92.1
7.1.93.1 7.1.94.1 7.1.95.1 7.1.96.1
7.1.97.1 7.1.98.1 7.1.99.1 7.1.100.1
7.1.101.1 7.1.102.1 7.1.103.1 7.1.104.1
7.1.105.1 7.1.106.1 7.1.107.1 7.1.108.1
7.1.109.1 7.1.110.1 7.1.111.1 7.1.112.1
7.1.113.1 7.1.114.1 7.1.115.1 7.1.116.1
7.1.117.1 7.1.118.1 7.1.119.1 7.1.120.1
7.1.121.1 7.1.122.1 7.1.123.1 7.1.124.1
7.1.125.1 7.1.126.1 7.1.127.1 7.1.128.1
7.1.129.1 7.1.130.1 7.1.131.1 7.1.132.1
7.1.133.1 7.1.134.1 7.1.135.1 7.1.136.1
7.1.137.1 7.1.138.1 7.1.139.1 7.1.140.1
7.1.141.1 7.1.142.1 7.1.143.1 7.1.144.1
7.1.145.1 7.1.146.1 7.1.147.1 7.1.148.1
7.1.149.1 7.1.150.1 7.1.151.1 7.1.152.1
7.1.153.1 7.1.154.1 7.1.155.1 7.1.156.1
7.1.157.1 7.1.158.1 7.1.159.1 7.1.160.1
7.1.161.1 7.1.162.1 7.1.163.1 7.1.164.1
7.1.165.1 7.1.166.1 7.1.167.1 7.1.168.1
7.1.169.1 7.1.170.1 7.1.171.1 7.1.172.1
7.1.173.1 7.1.174.1 7.1.175.1 7.1.176.1
7.1.177.1 7.1.178.1 7.1.179.1 7.1.180.1
7.1.181.1 7.1.182.1 7.1.183.1 7.1.184.1
7.1.185.1 7.1.186.1 7.1.187.1 7.1.188.1
7.1.189.1 7.1.190.1 7.1.191.1 7.1.192.1
7.1.193.1 7.1.194.1 7.1.195.1 7.1.196.1
7.1.197.1 7.1.198.1 7.1.199.1 7.1.200.1
7.1.201.1 7.1.202.1 7.1.203.1 7.1.204.1
7.1.205.1 7.1.206.1 7.1.207.1 7.1.208.1
7.1.209.1 7.1.210.1 7.1.211.1 7.1.212.1
7.1.213.1 7.1.214.1 7.1.215.1 7.1.216.1
7.1.217.1 7.1.218.1 7.1.219.1 7.1.220.1
7.1.221.1 7.1.222.1 7.1.223.1 7.1.224.1
7.1.225.1 7.1.226.1 7.1.227.1 7.1.228.1
7.1.229.1 7.1.230.1 7.1.231.1 7.1.232.1
7.1.233.1 7.1.234.1 7.1.235.1 7.1.236.1
7.1.237.1 7.1.238.1 7.1.239.1 7.1.240.1
7.1.241.1 7.1.242.1 7.1.243.1 7.1.244.1
7.1.245.1 7.1.246.1 7.1.247.1 7.1.248.1
7.1.249.1 7.1.250.1 7.2.1.1 7.2.2.1
```

7.2.3.1 7.2.4.1 7.2.5.1 7.2.6.1
7.2.7.1 7.2.8.1 7.2.9.1 7.2.10.1
7.2.11.1 7.2.12.1 7.2.13.1 7.2.14.1
7.2.15.1 7.2.16.1 7.2.17.1 7.2.18.1
7.2.19.1 7.2.20.1 7.2.21.1 7.2.22.1
7.2.23.1 7.2.24.1 7.2.25.1 7.2.26.1
7.2.27.1 7.2.28.1 7.2.29.1 7.2.30.1
7.2.31.1 7.2.32.1 7.2.33.1 7.2.34.1
7.2.35.1 7.2.36.1 7.2.37.1 7.2.38.1
7.2.39.1 7.2.40.1 7.2.41.1 7.2.42.1
7.2.43.1 7.2.44.1 7.2.45.1 7.2.46.1
7.2.47.1 7.2.48.1 7.2.49.1 7.2.50.1
7.2.51.1 7.2.52.1 7.2.53.1 7.2.54.1
7.2.55.1 7.2.56.1 7.2.57.1 7.2.58.1
7.2.59.1 7.2.60.1 7.2.61.1 7.2.62.1
7.2.63.1 7.2.64.1 7.2.65.1 7.2.66.1
7.2.67.1 7.2.68.1 7.2.69.1 7.2.70.1
7.2.71.1 7.2.72.1 7.2.73.1 7.2.74.1
7.2.75.1 7.2.76.1 7.2.77.1 7.2.78.1
7.2.79.1 7.2.80.1 7.2.81.1 7.2.82.1
7.2.83.1 7.2.84.1 7.2.85.1 7.2.86.1
7.2.87.1 7.2.88.1 7.2.89.1 7.2.90.1
7.2.91.1 7.2.92.1 7.2.93.1 7.2.94.1
7.2.95.1 7.2.96.1 7.2.97.1 7.2.98.1
7.2.99.1 7.2.100.1 7.2.101.1 7.2.102.1
7.2.103.1 7.2.104.1 7.2.105.1 7.2.106.1
7.2.107.1 7.2.108.1 7.2.109.1 7.2.110.1
7.2.111.1 7.2.112.1 7.2.113.1 7.2.114.1
7.2.115.1 7.2.116.1 7.2.117.1 7.2.118.1
7.2.119.1 7.2.120.1 7.2.121.1 7.2.122.1
7.2.123.1 7.2.124.1 7.2.125.1 7.2.126.1
7.2.127.1 7.2.128.1 7.2.129.1 7.2.130.1
7.2.131.1 7.2.132.1 7.2.133.1 7.2.134.1
7.2.135.1 7.2.136.1 7.2.137.1 7.2.138.1
7.2.139.1 7.2.140.1 7.2.141.1 7.2.142.1
7.2.143.1 7.2.144.1 7.2.145.1 7.2.146.1
7.2.147.1 7.2.148.1 7.2.149.1 7.2.150.1
7.2.151.1 7.2.152.1 7.2.153.1 7.2.154.1
7.2.155.1 7.2.156.1 7.2.157.1 7.2.158.1
7.2.159.1 7.2.160.1 7.2.161.1 7.2.162.1
7.2.163.1 7.2.164.1 7.2.165.1 7.2.166.1
7.2.167.1 7.2.168.1 7.2.169.1 7.2.170.1
7.2.171.1 7.2.172.1 7.2.173.1 7.2.174.1
7.2.175.1 7.2.176.1 7.2.177.1 7.2.178.1
7.2.179.1 7.2.180.1 7.2.181.1 7.2.182.1
7.2.183.1 7.2.184.1 7.2.185.1 7.2.186.1
7.2.187.1 7.2.188.1 7.2.189.1 7.2.190.1
7.2.191.1 7.2.192.1 7.2.193.1 7.2.194.1
7.2.195.1 7.2.196.1 7.2.197.1 7.2.198.1
7.2.199.1 7.2.200.1 7.2.201.1 7.2.202.1
7.2.203.1 7.2.204.1 7.2.205.1 7.2.206.1
7.2.207.1 7.2.208.1 7.2.209.1 7.2.210.1
7.2.211.1 7.2.212.1 7.2.213.1 7.2.214.1
7.2.215.1 7.2.216.1 7.2.217.1 7.2.218.1
7.2.219.1 7.2.220.1 7.2.221.1 7.2.222.1
7.2.223.1 7.2.224.1 7.2.225.1 7.2.226.1
7.2.227.1 7.2.228.1 7.2.229.1 7.2.230.1
7.2.231.1 7.2.232.1 7.2.233.1 7.2.234.1
7.2.235.1 7.2.236.1 7.2.237.1 7.2.238.1
7.2.239.1 7.2.240.1 7.2.241.1 7.2.242.1
7.2.243.1 7.2.244.1 7.2.245.1 7.2.246.1

Show Commands

```
7.2.247.1 7.2.248.1 7.2.249.1 7.2.250.1
7.3.1.1 7.3.2.1 7.3.3.1 7.3.4.1
7.3.5.1 7.3.6.1 7.3.7.1 7.3.8.1
7.3.9.1 7.3.10.1 7.3.11.1 7.3.12.1
7.3.13.1 7.3.14.1 7.3.15.1 7.3.16.1
7.3.17.1 7.3.18.1 7.3.19.1 7.3.20.1
7.3.21.1 7.3.22.1 7.3.23.1 7.3.24.1
7.3.25.1 7.3.26.1 7.3.27.1 7.3.28.1
7.3.29.1 7.3.30.1 7.3.31.1 7.3.32.1
7.3.33.1 7.3.34.1 7.3.35.1 7.3.36.1
7.3.37.1 7.3.38.1 7.3.39.1 7.3.40.1
7.3.41.1 7.3.42.1 7.3.43.1 7.3.44.1
7.3.45.1 7.3.46.1 7.3.47.1 7.3.48.1
7.3.49.1 7.3.50.1 7.3.51.1 7.3.52.1
7.3.53.1 7.3.54.1 7.3.55.1 7.3.56.1
7.3.57.1 7.3.58.1 7.3.59.1 7.3.60.1
7.3.61.1 7.3.62.1 7.3.63.1 7.3.64.1
7.3.65.1 7.3.66.1 7.3.67.1 7.3.68.1
7.3.69.1 7.3.70.1 7.3.71.1 7.3.72.1
7.3.73.1 7.3.74.1 7.3.75.1 7.3.76.1
7.3.77.1 7.3.78.1 7.3.79.1 7.3.80.1
7.3.81.1 7.3.82.1 7.3.83.1 7.3.84.1
7.3.85.1 7.3.86.1 7.3.87.1 7.3.88.1
7.3.89.1 7.3.90.1 7.3.91.1 7.3.92.1
7.3.93.1 7.3.94.1 7.3.95.1 7.3.96.1
7.3.97.1 7.3.98.1 7.3.99.1 7.3.100.1
7.3.101.1 7.3.102.1 7.3.103.1 7.3.104.1
7.3.105.1 7.3.106.1 7.3.107.1 7.3.108.1
7.3.109.1 7.3.110.1 7.3.111.1 7.3.112.1
7.3.113.1 7.3.114.1 7.3.115.1 7.3.116.1
7.3.117.1 7.3.118.1 7.3.119.1 7.3.120.1
7.3.121.1 7.3.122.1 7.3.123.1 7.3.124.1
7.3.125.1 7.4.1.1 7.4.2.1 7.4.3.1
7.4.4.1 7.4.5.1 7.4.6.1 7.4.7.1
7.4.8.1 7.4.9.1 7.4.10.1 7.4.11.1
7.4.12.1 7.4.13.1 7.4.14.1 7.4.15.1
7.4.16.1 7.4.17.1 7.4.18.1 7.4.19.1
7.4.20.1 7.4.21.1 7.4.22.1 7.4.23.1
7.4.24.1 7.4.25.1 7.4.26.1 7.4.27.1
7.4.28.1 7.4.29.1 7.4.30.1 7.4.31.1
7.4.32.1 7.4.33.1 7.4.34.1 7.4.35.1
7.4.36.1 7.4.37.1 7.4.38.1 7.4.39.1
7.4.40.1 7.4.41.1 7.4.42.1 7.4.43.1
7.4.44.1 7.4.45.1 7.4.46.1 7.4.47.1
7.4.48.1 7.4.49.1 7.4.50.1 7.4.51.1
7.4.52.1 7.4.53.1 7.4.54.1 7.4.55.1
7.4.56.1 7.4.57.1 7.4.58.1 7.4.59.1
7.4.60.1 7.4.61.1 7.4.62.1 7.4.63.1
7.4.64.1 7.4.65.1 7.4.66.1 7.4.67.1
7.4.68.1 7.4.69.1 7.4.70.1 7.4.71.1
7.4.72.1 7.4.73.1 7.4.74.1 7.4.75.1
7.4.76.1 7.4.77.1 7.4.78.1 7.4.79.1
7.4.80.1 7.4.81.1 7.4.82.1 7.4.83.1
7.4.84.1 7.4.85.1 7.4.86.1 7.4.87.1
7.4.88.1 7.4.89.1 7.4.90.1 7.4.91.1
7.4.92.1 7.4.93.1 7.4.94.1 7.4.95.1
7.4.96.1 7.4.97.1 7.4.98.1 7.4.99.1
7.4.100.1 7.4.101.1 7.4.102.1 7.4.103.1
7.4.104.1 7.4.105.1 7.4.106.1 7.4.107.1
7.4.108.1 7.4.109.1 7.4.110.1 7.4.111.1
7.4.112.1 7.4.113.1 7.4.114.1 7.4.115.1
```

7.4.116.1 7.4.117.1 7.4.118.1 7.4.119.1
7.4.120.1 7.4.121.1 7.4.122.1 7.4.123.1
7.4.124.1 7.4.125.1 7.5.1.1 7.5.2.1
7.5.3.1 7.5.4.1 7.5.5.1 7.5.6.1
7.5.7.1 7.5.8.1 7.5.9.1 7.5.10.1
7.5.11.1 7.5.12.1 7.5.13.1 7.5.14.1
7.5.15.1 7.5.16.1 7.5.17.1 7.5.18.1
7.5.19.1 7.5.20.1 7.5.21.1 7.5.22.1
7.5.23.1 7.5.24.1 7.5.25.1 7.5.26.1
7.5.27.1 7.5.28.1 7.5.29.1 7.5.30.1
7.5.31.1 7.5.32.1 7.5.33.1 7.5.34.1
7.5.35.1 7.5.36.1 7.5.37.1 7.5.38.1
7.5.39.1 7.5.40.1 7.5.41.1 7.5.42.1
7.5.43.1 7.5.44.1 7.5.45.1 7.5.46.1
7.5.47.1 7.5.48.1 7.5.49.1 7.5.50.1
7.5.51.1 7.5.52.1 7.5.53.1 7.5.54.1
7.5.55.1 7.5.56.1 7.5.57.1 7.5.58.1
7.5.59.1 7.5.60.1 7.5.61.1 7.5.62.1
7.5.63.1 7.5.64.1 7.5.65.1 7.5.66.1
7.5.67.1 7.5.68.1 7.5.69.1 7.5.70.1
7.5.71.1 7.5.72.1 7.5.73.1 7.5.74.1
7.5.75.1 7.5.76.1 7.5.77.1 7.5.78.1
7.5.79.1 7.5.80.1 7.5.81.1 7.5.82.1
7.5.83.1 7.5.84.1 7.5.85.1 7.5.86.1
7.5.87.1 7.5.88.1 7.5.89.1 7.5.90.1
7.5.91.1 7.5.92.1 7.5.93.1 7.5.94.1
7.5.95.1 7.5.96.1 7.5.97.1 7.5.98.1
7.5.99.1 7.5.100.1 7.5.101.1 7.5.102.1
7.5.103.1 7.5.104.1 7.5.105.1 7.5.106.1
7.5.107.1 7.5.108.1 7.5.109.1 7.5.110.1
7.5.111.1 7.5.112.1 7.5.113.1 7.5.114.1
7.5.115.1 7.5.116.1 7.5.117.1 7.5.118.1
7.5.119.1 7.5.120.1 7.5.121.1 7.5.122.1
7.5.123.1 7.5.124.1 7.5.125.1 7.6.1.1
7.6.2.1 7.6.3.1 7.6.4.1 7.6.5.1
7.6.6.1 7.6.7.1 7.6.8.1 7.6.9.1
7.6.10.1 7.6.11.1 7.6.12.1 7.6.13.1
7.6.14.1 7.6.15.1 7.6.16.1 7.6.17.1
7.6.18.1 7.6.19.1 7.6.20.1 7.6.21.1
7.6.22.1 7.6.23.1 7.6.24.1 7.6.25.1
7.6.26.1 7.6.27.1 7.6.28.1 7.6.29.1
7.6.30.1 7.6.31.1 7.6.32.1 7.6.33.1
7.6.34.1 7.6.35.1 7.6.36.1 7.6.37.1
7.6.38.1 7.6.39.1 7.6.40.1 7.6.41.1
7.6.42.1 7.6.43.1 7.6.44.1 7.6.45.1
7.6.46.1 7.6.47.1 7.6.48.1 7.6.49.1
7.6.50.1 7.6.51.1 7.6.52.1 7.6.53.1
7.6.54.1 7.6.55.1 7.6.56.1 7.6.57.1
7.6.58.1 7.6.59.1 7.6.60.1 7.6.61.1
7.6.62.1 7.6.63.1 7.6.64.1 7.6.65.1
7.6.66.1 7.6.67.1 7.6.68.1 7.6.69.1
7.6.70.1 7.6.71.1 7.6.72.1 7.6.73.1
7.6.74.1 7.6.75.1 7.6.76.1 7.6.77.1
7.6.78.1 7.6.79.1 7.6.80.1 7.6.81.1
7.6.82.1 7.6.83.1 7.6.84.1 7.6.85.1
7.6.86.1 7.6.87.1 7.6.88.1 7.6.89.1
7.6.90.1 7.6.91.1 7.6.92.1 7.6.93.1
7.6.94.1 7.6.95.1 7.6.96.1 7.6.97.1
7.6.98.1 7.6.99.1 7.6.100.1 7.6.101.1
7.6.102.1 7.6.103.1 7.6.104.1 7.6.105.1
7.6.106.1 7.6.107.1 7.6.108.1 7.6.109.1

Show Commands

```
7.6.110.1 7.6.111.1 7.6.112.1 7.6.113.1
7.6.114.1 7.6.115.1 7.6.116.1 7.6.117.1
7.6.118.1 7.6.119.1 7.6.120.1 7.6.121.1
7.6.122.1 7.6.123.1 7.6.124.1 7.6.125.1
10.100.14.1 10.100.15.1 20.20.1.1 20.20.2.1
33.66.33.1 33.66.34.1 33.66.35.1 33.66.63.1
33.66.64.1 33.66.65.1 40.40.1.1 40.40.2.1
57.57.57.57 60.60.1.1 60.60.2.1 70.70.1.1
70.71.2.1 99.99.0.1 110.1.1.51 110.1.2.51
110.1.3.51 110.1.4.51 110.1.5.51 110.1.6.51
110.1.7.51 110.1.8.51 110.1.9.51 110.1.10.51
110.1.11.51 110.1.12.51 110.1.13.51 110.1.14.51
110.1.15.51 110.1.16.51 110.1.17.51 110.1.18.51
110.1.19.51 110.1.20.51 110.1.21.51 110.1.22.51
110.1.23.51 110.1.24.51 110.1.25.51 110.1.26.51
110.1.27.51 110.1.28.51 110.1.29.51 110.1.30.51
110.1.31.51 110.1.32.51 110.1.33.51 110.1.34.51
110.1.35.51 110.1.36.51 110.1.37.51 110.1.38.51
110.1.39.51 110.1.40.51 110.1.41.51 110.1.42.51
110.1.43.51 110.1.44.51 110.1.45.51 110.1.46.51
110.1.47.51 110.1.48.51 110.1.49.51 110.1.50.51
110.1.51.51 110.1.52.51 110.1.53.51 110.1.54.51
110.1.55.51 110.1.56.51 110.1.57.51 110.1.58.51
110.1.59.51 110.1.60.51 110.1.61.51 110.1.62.51
110.1.63.51 110.1.64.51 110.1.65.51 110.1.66.51
110.1.67.51 110.1.68.51 110.1.69.51 110.1.70.51
110.1.71.51 110.1.72.51 110.1.73.51 110.1.74.51
110.1.75.51 110.1.76.51 110.1.77.51 110.1.78.51
110.1.79.51 110.1.80.51 110.1.81.51 110.1.82.51
110.1.83.51 110.1.84.51 110.1.85.51 110.1.86.51
110.1.87.51 110.1.88.51 110.1.89.51 110.1.90.51
110.1.91.51 110.1.92.51 110.1.93.51 110.1.94.51
110.1.95.51 110.1.96.51 110.1.97.51 110.1.98.51
110.1.99.51 110.1.100.51 110.1.101.51 110.1.102.51
110.1.103.51 110.1.104.51 110.1.105.51 110.1.106.51
110.1.107.51 110.1.108.51 110.1.109.51 110.1.110.51
110.1.111.51 110.1.112.51 110.1.113.51 110.1.114.51
110.1.115.51 110.1.116.51 110.1.117.51 110.1.118.51
110.1.119.51 110.1.120.51 110.1.121.51 110.1.122.51
110.1.123.51 110.1.124.51 110.1.125.51 110.1.126.51
110.1.127.51 110.1.128.51 110.1.129.51 110.1.130.51
110.1.131.51 110.1.132.51 110.1.133.51 110.1.134.51
110.1.135.51 110.1.136.51 110.1.137.51 110.1.138.51
110.1.139.51 110.1.140.51 110.1.141.51 110.1.142.51
110.1.143.51 110.1.144.51 110.1.145.51 110.1.146.51
110.1.147.51 110.1.148.51 110.1.149.51 110.1.150.51
110.1.151.51 110.1.152.51 110.1.153.51 110.1.154.51
110.1.155.51 110.1.156.51 110.1.157.51 110.1.158.51
110.1.159.51 110.1.160.51 110.1.161.51 110.1.162.51
110.1.163.51 110.1.164.51 110.1.165.51 110.1.166.51
110.1.167.51 110.1.168.51 110.1.169.51 110.1.170.51
110.1.171.51 110.1.172.51 110.1.173.51 110.1.174.51
110.1.175.51 110.1.176.51 110.1.177.51 110.1.178.51
110.1.179.51 110.1.180.51 110.1.181.51 110.1.182.51
110.1.183.51 110.1.184.51 110.1.185.51 110.1.186.51
110.1.187.51 110.1.188.51 110.1.189.51 110.20.1.1
150.50.80.1 150.50.81.1 150.50.100.1 150.50.101.1
150.50.120.1 150.50.121.1 150.50.140.1 150.50.141.1
150.60.10.1 150.60.20.1 150.60.40.1 160.60.80.1
160.60.81.1 170.70.70.1 180.1.110.1 180.1.111.1
180.100.1.1 200.0.0.1 201.0.0.1
```

```

=====
*A:SR4>config>router>ldp>peer-params#
*A:SR4>config>router>ldp>peer-params# show router ldp session 110.20.1.1 local-
addresses recv ip-addr 201.0.0.1
=====
LDP Session Local-Addresses
=====
-----
Session with Peer 110.20.1.1:0, Local 110.20.1.3:0
-----
Recv Addresses: 201.0.0.1
=====
*A:SR4>config>router>ldp>peer-params#
*A:SR4>config>router>ldp>peer-params# show router ldp session 110.20.1.1 local-
addresses sent ip-addr 6.2.7.3
=====
LDP Session Local-Addresses
=====
-----
Session with Peer 110.20.1.1:0, Local 110.20.1.3:0
-----
Sent Addresses: 6.2.7.3
=====

```

auth-keychain

Syntax	auth-keychain [<i>keychain</i>]
Context	show>router>ldp
Description	This command displays LDP sessions using a particular authentication key-chain.
Parameters	<i>keychain</i> — Specifies an existing keychain name.

Sample Output

```

*A:ALA-48>config>router>ldp# show router ldp auth-keychain
=====
LDP Peers
=====
Peer          TTL Security Min-TTL-Value Authentication Auth key chain
-----
10.20.1.3     Disabled      n/a           Disabled      eta_keychain1
-----
No. of Peers: 1
=====
*A:ALA-48>config>router>ldp#

```

bindings

Syntax	<p>bindings [fec-type <i>fec-type</i> [detail summary]] [session <i>ip-addr[:label-space]</i>]</p> <p>bindings [fec-type p2mp] [p2mp <i>identifier root ip-address</i>] [detail summary] [session <i>ip-addr[:label-space]</i>]</p> <p>bindings [<i>label-type</i>] [<i>start-label</i> [<i>end-label</i>]]</p> <p>bindings {prefix <i>ip-prefix/mask</i> [detail]} [session <i>ip-addr[:label-space]</i>]</p> <p>bindings active [prefix <i>ip-prefix/mask</i>][summary egress-nh <i>ip-prefix/mask</i> egress-if <i>port-id</i> egress-lsp <i>tunnel-id</i>]</p> <p>bindings active [fec-type prefixes] [prefix <i>ip-prefix/mask</i>] [egress-nh <i>ip-prefix/mask</i> egress-if <i>port-id</i> egress-lsp <i>tunnel-id</i>] [summary]</p> <p>bindings active [fec-type p2mp] [p2mp-id <i>identifier root ip-address</i>] [egress-nh <i>ip-prefix/mask</i> egress-if <i>port-id</i> egress-lsp <i>tunnel-id</i>] [summary]</p> <p>bindings service-id <i>service-id</i> [detail]</p> <p>bindings vc-type <i>vc-type</i> [{vc-id <i>vc-id</i> agi <i>agi</i>] [session <i>ip-addr[:lab el-space]</i>]]</p> <p>bindings vc-type <i>vc-type</i> agi <i>agi</i> <i>sail-type2</i> <i>global-id:prefix:ac-id</i> <i>taii-type2</i> <i>global_id:prefix:ac_id</i> [detail]</p> <p>bindings p2mp-id <i>identifier root ip-address</i> [detail]</p>
Context	show>router>ldp
Description	This command displays the contents of the label information base.
Parameters	<p>fec-type <i>fec-type</i> — Specify the kind of FEC that the label mapping, withdraw, release and request messages are referring to.</p> <p>detail — Displays detailed information.</p> <p>summary — Displays information in a summarized format.</p> <p>session <i>ip-addr</i> — displays configuration information about LDP sessions.</p> <p><i>ip-prefix</i> — Specify information for the specified IP prefix and mask length. Host bits must be 0.</p> <p><i>mask</i> — Specifies the 32-bit address mask used to indicate the bits of an IP address that are being used for the subnet address.</p> <p>Values 0 — 32</p> <p><i>label-space</i> — Specifies the label space identifier that the router is advertising on the interface.</p> <p>Values 0 — 65535</p> <p><i>start-label</i> — Specifies a label value to begin the display.</p> <p>Values 16 — 1048575</p> <p><i>end-label</i> — Specifies a label value to end the display.</p> <p>Values 17 — 1048575</p> <p><i>vc-type</i> — Specifies the VC type to display.</p> <p>Values ethernet , vlan , mirror</p> <p><i>vc-id</i> — Specifies the VC ID to display.</p> <p>Values 1 — 4294967295</p> <p><i>service-id</i> — Specifies the service ID number to display.</p> <p>Values 1 — 2147483647</p>

egress-lsp tunnel-id —

Values 0 — 4294967295

egress-if port-id — Specifies the egress interface port ID.

Values slot[/mda[/port]] or slot/mda/port[.channel]

aps-id	aps-group-id[.channel]
aps	keyword
group-id	1— 64
ccag-id	slot/mda/path-id[cc-type]
path-id	a, b
cc-type	.sap-net, .net-sap

Output LDP Bindings Output — The following table describes the LDP bindings fields.

Label	Description
Legend	U: Label In Use A: Apipe service N: Label Not In Use F: Fpipe service W: Label Withdrawn I: IES service S: Status Signaled Up R: VPRN service D: Status Signaled Down P: Ipipe service E: Epipe service WP: Label Withdraw Pending V: VPLS service C: Cpipe service M: Mirror service TLV: (Type, Length: Value)
Type	The service type exchanging labels in the SDP. The possible types displayed are VPLS, Epipe, Spoke, and Unknown.
VCId	The value used by each end of an SDP tunnel to identify the VC.
SvcID	The unique service identification number identifying the service in the service domain.
SDPId	The SDP number identifying the SDP in the service domain.
Peer	The IP address of the peer.
EgrIntf/LspId	Displays the LSP Tunnel ID (not the LSP path ID).
IngLbl	The ingress LDP label. U — Label in use. R — Label released.
EgrLbl	The egress LDP label.
LMTU	The local MTU value.
RMTU	The remote MTU value.
No. of Service Bindings	The total number of LDP bindings on the router.

Sample Output

```
*A:Dut-C# show router ldp bindings active fec-type p2mp
=====
LDP Generic P2MP Bindings (Active)
=====
P2MP-Id      RootAddr
Interface    Op           IngLbl      EgrLbl      EgrIntf/   EgrNextHop
                LspId
-----
1            10.20.1.3
73728      Push        --          262142     1/1/2:0   Unnumbered

1            10.20.1.3
73728      Push        --          262137     2/1/2:0   Unnumbered

2            10.20.1.3
73729      Push        --          262141     1/1/2:0   Unnumbered

2            10.20.1.3
73729      Push        --          262136     2/1/2:0   Unnumbered

3            10.20.1.3
73730      Push        --          262140     1/1/2:0   Unnumbered

3            10.20.1.3
73730      Push        --          262135     2/1/2:0   Unnumbered

4            10.20.1.3
73731      Push        --          262139     1/1/2:0   Unnumbered

4            10.20.1.3
73731      Push        --          262134     2/1/2:0   Unnumbered
{...snip...}
```

```
*A:Dut-C# show router ldp bindings fec-type p2mp
=====
LDP LSR ID: 10.20.1.3
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
       WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
=====
LDP Generic P2MP Bindings
=====
P2MP-Id      RootAddr
Interface    Peer        IngLbl      EgrLbl      EgrIntf/   EgrNextHop
                LspId
-----
1            10.20.1.3
73728      10.20.1.2  --          262142     1/1/2:0   Unnumbered

1            10.20.1.3
73728      10.20.1.4  --          262137     2/1/2:0   Unnumbered

2            10.20.1.3
73729      10.20.1.2  --          262141     1/1/2:0   Unnumbered

2            10.20.1.3
73729      10.20.1.4  --          262136     2/1/2:0   Unnumbered
```

```

3          10.20.1.3
73730     10.20.1.2      --      262140 1/1/2:0      Unnumbered

3          10.20.1.3
73730     10.20.1.4      --      262135 2/1/2:0      Unnumbered

4          10.20.1.3
73731     10.20.1.2      --      262139 1/1/2:0      Unnumbered

4          10.20.1.3
73731     10.20.1.4      --      262134 2/1/2:0      Unnumbered

```

```
{...snip...}
```

```
*A:Dut-C# show router ldp bindings fec-type p2mp detail
```

```
=====
LDP LSR ID: 10.20.1.3
=====
```

```
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
=====
```

```
LDP Generic P2MP Bindings
=====
```

```
-----
P2MP Type      : 1                P2MP-Id        : 1
Root-Addr      : 10.20.1.3
-----
```

```
Ing Lbl        : --                Peer            : 10.20.1.2
Egr Lbl        : 262142
Egr Int/LspId  : 1/1/2:0
EgrNextHop     : Unnumbered
Egr. Flags     : None              Ing. Flags      : None
Egr If Name    : ip-10.180.3.3
Metric         : 1                Mtu             : 1496
-----
```

```
P2MP Type      : 1                P2MP-Id        : 1
Root-Addr      : 10.20.1.3
-----
```

```
Ing Lbl        : --                Peer            : 10.20.1.4
Egr Lbl        : 262137
Egr Int/LspId  : 2/1/2:0
EgrNextHop     : Unnumbered
Egr. Flags     : None              Ing. Flags      : None
Egr If Name    : ip-10.180.11.3
Metric         : 1                Mtu             : 1496
-----
```

```
P2MP Type      : 1                P2MP-Id        : 2
Root-Addr      : 10.20.1.3
-----
```

```
Ing Lbl        : --                Peer            : 10.20.1.2
Egr Lbl        : 262141
Egr Int/LspId  : 1/1/2:0
EgrNextHop     : Unnumbered
Egr. Flags     : None              Ing. Flags      : None
Egr If Name    : ip-10.180.3.3
Metric         : 1                Mtu             : 1496
-----
```

```
P2MP Type      : 1                P2MP-Id        : 2
Root-Addr      : 10.20.1.3
-----
```

Show Commands

```

Ing Lbl          :  --                      Peer          : 10.20.1.4
Egr Lbl          : 262136
Egr Int/LspId    : 2/1/2:0
EgrNextHop       : Unnumbered
Egr. Flags       : None                      Ing. Flags       : None
Egr If Name      : ip-10.180.11.3
Metric           : 1                          Mtu             : 1496

```

 {...snip...}

A:Dut-C# show router ldp bindings active

```

Legend: (S) - Static      (M) - Multi-homed Secondary Support
        (B) - BGP Next Hop (BU) - Alternate Next-hop for Fast Re-Route

```

 LDP Prefix Bindings (Active)

Prefix	Op	IngLbl	EgrLbl	EgrIntf/LspId	EgrNextHop
10.20.1.1/32	Push	--	262143	1/1/1	Unnumbered
10.20.1.1/32	Swap	262138	262143	1/1/1	Unnumbered
10.20.1.2/32	Push	--	262143	lag-1	Unnumbered
10.20.1.2/32	Swap	262139	262143	lag-1	Unnumbered
10.20.1.3/32	Pop	262143	--	--	--
10.20.1.4/32	Push	--	262143	2/1/2	Unnumbered
10.20.1.4/32	Swap	262142	262143	2/1/2	Unnumbered
10.20.1.5/32	Push	--	262143	2/1/1	Unnumbered
10.20.1.5/32	Swap	262141	262143	2/1/1	Unnumbered
10.20.1.6/32	Push	--	262140	2/1/2	Unnumbered
10.20.1.6/32	Swap	262140	262140	2/1/2	Unnumbered

 No. of Prefix Active Bindings: 11

 {...snip...}

*A:Dut-C# show router ldp bindings

 LDP LSR ID: 10.20.1.3

```

Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        S - Status Signaled Up, D - Status Signaled Down
        E - Epipe Service, V - VPLS Service, M - Mirror Service
        A - Apipe Service, F - Fpipe Service, I - IES Service, R - VPRN service
        P - Ipipe Service, WP - Label Withdraw Pending, C - Cpipe Service
        BU - Alternate For Fast Re-Route, TLV - (Type, Length: Value)

```

 LDP Prefix Bindings

Prefix Peer	IngLbl	EgrLbl	EgrIntf/ LspId	EgrNextHop
10.20.1.1/32	--	262143	1/1/1	Unnumbered
10.20.1.1				
10.20.1.1/32	262138U	262142	--	--
10.20.1.2				
10.20.1.1/32	262138U	262138	--	--
10.20.1.4				
10.20.1.1/32	262138U	262138	--	--
10.20.1.5				
10.20.1.2/32	262139U	262142	--	--

```

10.20.1.1
10.20.1.2/32      --      262143      lag-1      Unnumbered
10.20.1.2
10.20.1.2/32      262139U      262139      --      --
10.20.1.4
10.20.1.2/32      262139U      262139      --      --
10.20.1.5
10.20.1.3/32      262143U      --      --      --
10.20.1.1
10.20.1.3/32      262143U      --      --      --
10.20.1.2
10.20.1.3/32      262143U      --      --      --
10.20.1.4
10.20.1.3/32      262143U      --      --      --
10.20.1.5
10.20.1.4/32      262142U      262141      --      --
10.20.1.1
10.20.1.4/32      262142U      262141      --      --
10.20.1.2
10.20.1.4/32      --      262143      2/1/2      Unnumbered
10.20.1.4
10.20.1.4/32      262142U      262141      --      --
10.20.1.5
10.20.1.5/32      262141U      262138      --      --
10.20.1.1
10.20.1.5/32      262141U      262139      --      --
10.20.1.2
10.20.1.5/32      262141U      262141      --      --
10.20.1.4
10.20.1.5/32      --      262143      2/1/1      Unnumbered
10.20.1.5
10.20.1.6/32      262140U      262140      --      --
10.20.1.1
10.20.1.6/32      262140U      262138      --      --
10.20.1.2
10.20.1.6/32      262140N      262140      2/1/2      Unnumbered
10.20.1.4
10.20.1.6/32      262140U      262140      --      --
10.20.1.5

```

```
-----
No. of Prefix Bindings: 24
=====
```

```
{...snip...}
```

```
*A:Dut-C# show router ldp bindings detail
```

```
=====
LDP LSR ID: 10.20.1.3
=====
```

```
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
S - Status Signaled Up, D - Status Signaled Down
E - Epipe Service, V - VPLS Service, M - Mirror Service
A - Apipe Service, F - Fpipe Service, I - IES Service, R - VPRN service
P - Ipipe Service, WP - Label Withdraw Pending, C - Cpipe Service
BU - Alternate For Fast Re-Route, TLV - (Type, Length: Value)
=====
```

```
LDP Prefix Bindings
=====
```

```
-----
Prefix          : 10.20.1.1/32
-----
```

```
-----
Ing Lbl          :      --      Peer          : 10.20.1.1
-----
```

Show Commands

```

Egr Lbl          : 262143
Egr Int/LspId    : 1/1/1
EgrNextHop       : Unnumbered
Egr. Flags       : None           Ing. Flags       : None
Egr If Name      : ip-10.10.2.3
Metric           : 1000           Mtu              : 1500
-----
Prefix           : 10.20.1.1/32
-----
Ing Lbl          : 262138U        Peer              : 10.20.1.2
Egr Lbl          : 262142
Egr Int/LspId    : --
EgrNextHop       : --
Egr. Flags       : None           Ing. Flags       : None
Egr If Name      : n/a
-----
Prefix           : 10.20.1.1/32
-----
Ing Lbl          : 262138U        Peer              : 10.20.1.4
Egr Lbl          : 262138
Egr Int/LspId    : --
EgrNextHop       : --
Egr. Flags       : None           Ing. Flags       : None
Egr If Name      : n/a
-----
Prefix           : 10.20.1.1/32
-----
Ing Lbl          : 262138U        Peer              : 10.20.1.5
Egr Lbl          : 262138
Egr Int/LspId    : --
EgrNextHop       : --
Egr. Flags       : None           Ing. Flags       : None
Egr If Name      : n/a
-----
Prefix           : 10.20.1.2/32
-----
Ing Lbl          : 262139U        Peer              : 10.20.1.1
Egr Lbl          : 262142
Egr Int/LspId    : --
EgrNextHop       : --
Egr. Flags       : None           Ing. Flags       : None
Egr If Name      : n/a
-----
{...snip...}

*A:SRR# show router ldp bindings fec-type p2mp | match BU pre-lines 5 post-lines 5
=====
LDP LSR ID: 110.20.1.2
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
       WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
=====
LDP Generic P2MP Bindings
=====
P2MP-Id      RootAddr
Interface    Peer          IngLbl  EgrLbl  EgrIntf/  EgrNextHop
                               LspId
-----
8193         110.20.1.1
86055       20.20.1.1      254685BU  --      --      --

```

```

8193          110.20.1.1
86055         110.20.1.1      254318U  --  --  --

8194          110.20.1.1
86062         20.20.1.1      255758BU --  --  --

8194          110.20.1.1
86062         110.20.1.1      254317U  --  --  --

8195          110.20.1.1
86043         20.20.1.1      254692BU --  --  --
{....snip....}

```

```
*A:SRR# show router ldp bindings active fec-type p2mp | match BU pre-lines 5 post-
lines 5
```

P2MP-Id Interface	RootAddr Op	IngLbl	EgrLbl	EgrIntf/ LspId	EgrNextHop
8193 86055	110.20.1.1 Pop	254685BU	--	--	--
8193 86055	110.20.1.1 Pop	254318	--	--	--
8194 86062	110.20.1.1 Pop	255758BU	--	--	--
8194 86062	110.20.1.1 Pop	254317	--	--	--
8195 86043	110.20.1.1 Pop	254692BU	--	--	--
8195 86043	110.20.1.1 Pop	254316	--	--	--
8197 89107	110.20.1.1 Pop	254858	--	--	--

```

{....snip....}

```

```
*A:SRR# show router ldp bindings active fec-type p2mp p2mp-id 8193 root 110.20.1.1
```

```
=====
LDP Generic P2MP Bindings (Active)
=====
```

P2MP-Id Interface	RootAddr Op	IngLbl	EgrLbl	EgrIntf/ LspId	EgrNextHop
8193 86055	110.20.1.1 Pop	254685BU	--	--	--
8193 86055	110.20.1.1 Pop	254318	--	--	--

```
-----
No. of Generic P2MP Active Bindings: 2
=====
```

Show Commands

```
*A:SRR#
*A:SRR# show router ldp bindings fec-type p2mp p2mp-id 8193 root 110.20.1.1
=====
LDP LSR ID: 110.20.1.2
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
       WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
=====
LDP Generic P2MP Bindings
=====
P2MP-Id      RootAddr
Interface    Peer          IngLbl   EgrLbl EgrIntf/   EgrNextHop
                               LspId
-----
8193         110.20.1.1
86055       20.20.1.1     254685BU --      --        --

8193         110.20.1.1
86055       110.20.1.1     254318U  --      --        --

-----
No. of Generic P2MP Bindings: 2
=====
*A:SRR#

*A:SRR# show router ldp bindings fec-type p2mp p2mp-id 8193 root 110.20.1.1 detail
=====
LDP LSR ID: 110.20.1.2
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
       WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
=====
LDP Generic P2MP Bindings
=====
-----
P2MP Type      : 1                P2MP-Id      : 8193
Root-Addr      : 110.20.1.1
-----
Ing Lbl        : 254685BU      Peer          : 20.20.1.1
Egr Lbl        : --
Egr Int/LspId  : --
EgrNextHop     : --
Egr. Flags     : None         Ing. Flags    : None
-----
P2MP Type      : 1                P2MP-Id      : 8193
Root-Addr      : 110.20.1.1
-----
Ing Lbl        : 254318U      Peer          : 110.20.1.1
Egr Lbl        : --
Egr Int/LspId  : --
EgrNextHop     : --
Egr. Flags     : None         Ing. Flags    : None
=====
No. of Generic P2MP Bindings: 2
=====
*A:SRR#

*A:Dut-B# show router ldp bindings fec-type p2mp p2mp-id 8193 root 10.20.1.2
=====
```



```

LDP LSR ID: 10.20.1.2
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
       WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
=====
LDP Generic P2MP Bindings
=====
P2MP-Id      RootAddr
Interface    Peer          IngLbl    EgrLbl EgrIntf/    EgrNextHop
                               LspId
-----
8193         10.20.1.2
79733       10.20.1.3      --        258092 2/1/4:1    10.10.1.1

8193         10.20.1.2
79733       10.20.1.4      --        124027 1/1/1:1    10.10.2.2

8193         10.20.1.2
79733       10.20.1.5      --        125579 2/1/8:1    10.10.17.1
-----
No. of Generic P2MP Bindings: 3
=====

*A:Dut-B# show router ldp bindings fec-type p2mp p2mp-id 8193 root 10.20.1.2 session
10.20.1.3
=====
LDP LSR ID: 10.20.1.2
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
       WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
=====
LDP Generic P2MP Bindings
=====
P2MP-Id      RootAddr
Interface    Peer          IngLbl    EgrLbl EgrIntf/    EgrNextHop
                               LspId
-----
8193         10.20.1.2
79733       10.20.1.3      --        258092 2/1/4:1    10.10.1.1
-----
No. of Generic P2MP Bindings: 1
=====

*A:Dut-A# show router ldp bindings active
=====
Legend: (S) - Static          (M) - Multi-homed Secondary Support
       (B) - BGP Next Hop (BU) - Alternate Next-hop for Fast Re-Route
=====
LDP Prefix Bindings (Active)
=====
Prefix          Op   IngLbl    EgrLbl    EgrIntf/LspId  EgrNextHop
-----
10.20.1.1/32    Pop  131071    --         --              --
10.20.1.2/32    Push --        131071    1/1/1         10.10.1.2
10.20.1.2/32    Swap 131070    131071    1/1/1         10.10.1.2
10.20.1.2/32    Push --        262141BU 1/1/2         10.10.2.3
10.20.1.2/32    Swap 131070    262141BU 1/1/2         10.10.2.3
10.20.1.3/32    Push --        131069BU 1/1/1         10.10.1.2

```

Show Commands

```

10.20.1.3/32      Swap 131069      131069BU 1/1/1      10.10.1.2
10.20.1.3/32      Push  --          262143   1/1/2      10.10.2.3
10.20.1.3/32      Swap 131069      262143   1/1/2      10.10.2.3
10.20.1.4/32      Push  --          131068   1/1/1      10.10.1.2
10.20.1.4/32      Swap 131068      131068   1/1/1      10.10.1.2
10.20.1.4/32      Push  --          262140BU 1/1/2      10.10.2.3
10.20.1.4/32      Swap 131068      262140BU 1/1/2      10.10.2.3
10.20.1.5/32      Push  --          131067BU 1/1/1      10.10.1.2
10.20.1.5/32      Swap 131067      131067BU 1/1/1      10.10.1.2
10.20.1.5/32      Push  --          262139   1/1/2      10.10.2.3
10.20.1.5/32      Swap 131067      262139   1/1/2      10.10.2.3
10.20.1.6/32      Push  --          131066   1/1/1      10.10.1.2
10.20.1.6/32      Swap 131066      131066   1/1/1      10.10.1.2
10.20.1.6/32      Push  --          262138BU 1/1/2      10.10.2.3
10.20.1.6/32      Swap 131066      262138BU 1/1/2      10.10.2.3

```

No. of Prefix Active Bindings: 21

=====

LDP P2MP Bindings (Active)

```

=====
P2MP-Id      RootAddr
Interface    Op          IngLbl     EgrLbl EgrIntf/   EgrNextHop
                               LspId

```

No Matching Entries Found

=====

*A:Dut-A# show router ldp bindings

=====

LDP LSR ID: 10.20.1.1

```

=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
       S - Status Signaled Up, D - Status Signaled Down
       E - Epipe Service, V - VPLS Service, M - Mirror Service
       A - Apipe Service, F - Fpipe Service, I - IES Service, R - VPRN service
       P - Ipipe Service, WP - Label Withdraw Pending, C - Cpipe Service
       BU - Alternate Next-hop for Fast Re-Route, TLV - (Type, Length: Value)

```

=====

LDP Prefix Bindings

```

=====
Prefix      Peer      IngLbl     EgrLbl EgrIntf/   EgrNextHop
                               LspId

```

```

-----
10.20.1.1/32      10.20.1.2      131071U    --     --         --
10.20.1.1/32      10.20.1.3      131071U    --     --         --
10.20.1.2/32      10.20.1.2      --         131071 1/1/1    10.10.1.2
10.20.1.2/32      10.20.1.3      131070U    262141 1/1/2    10.10.2.3
10.20.1.3/32      10.20.1.2      131069U    131069 1/1/1    10.10.1.2
10.20.1.3/32      10.20.1.3      --         262143 1/1/2    10.10.2.3
10.20.1.4/32      10.20.1.2      131068N    131068 1/1/1    10.10.1.2
10.20.1.4/32      10.20.1.3      131068BU   262140 1/1/2    10.10.2.3
10.20.1.5/32      10.20.1.2      131067U    131067 1/1/1    10.10.1.2
10.20.1.5/32      10.20.1.3      131067N    262139 1/1/2    10.10.2.3
10.20.1.6/32      10.20.1.2      131066N    131066 1/1/1    10.10.1.2
10.20.1.6/32      10.20.1.3      131066BU   262138 1/1/2    10.10.2.3

```

No. of Prefix Bindings: 12

=====

LDP P2MP Bindings

=====

```

P2MP-Id          RootAddr
Interface        Peer          IngLbl  EgrLbl  EgrIntf/  EgrNextHop
                               LspId

-----
No Matching Entries Found
=====
LDP Service FEC 128 Bindings
=====
Type  VCId      SvcId      SDPIId    Peer          IngLbl  EgrLbl  LMTU  RMTU
-----
No Matching Entries Found
=====
LDP Service FEC 129 Bindings
=====
AGI                                     SAI
                                         TAI
Type          SvcId      SDPIId    Peer          IngLbl  EgrLbl  LMTU  RMTU
-----
No Matching Entries Found
=====

*A:SR1-A# show router ldp bindings service-id 100 detail
=====
LDP LSR ID: 1.1.1.1
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        S - Status Signaled Up, D - Status Signaled Down
        E - Epipe Service, V - VPLS Service, M - Mirror Service
        A - Apipe Service, F - Fpipe Service, I - IES Service, R - VPRN service
        P - Ipipe Service, WP - Label Withdraw Pending, C - Cpipe Service
        TLV - (Type, Length: Value)
=====
LDP Service Binding
=====
Type          : E-Eth          VcId          : 100
SvcId         : 100           SdpId         : 100
Peer Address  : 2.2.2.2       Vc-switching  : No
LMTU         : 986           RMTU          : 986
Egr. Lbl     : 130812D      Egr. Ctl Word : No
Egr. Flags   : None         Egr. Status Bits : Supported (0x1e)
Egr. Flow Label Tx : Yes/No  Egr. Flow Label Rx: Yes/No
Ing. Lbl     : 130808U      Ing. Ctl Word  : No
Ing. Flags   : None         Ing. Status Bits : Supported (0x16)
Ing. Flow Label Tx : Yes/No  Ing. Flow Label Tx: Yes/No
=====
No. of VC Labels: 1
...
=====
*A:SR1-A#

*A:SRU4>config>router>ldp# show router ldp bindings fec-type prefixes
=====
LDP LSR ID: 110.20.1.4
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        WP - Label Withdraw Pending
=====
LDP Prefix Bindings

```

Show Commands

```

=====
Prefix                Peer                IngLbl    EgrLbl  EgrIntf/  EgrNextHop
                   LspId
-----
1.1.1.0/24            10.20.1.22         --        1301    --        --
1.2.10.0/24           10.20.1.22         --        618     --        --
1.2.11.0/24           10.20.1.22         --        617     --        --
1.2.12.0/24           10.20.1.22         --        619     --        --
1.2.101.0/24          10.20.1.22         --        656     --        --
1.2.111.0/24          10.20.1.22         --        341     --        --
1.2.121.0/24          10.20.1.22         --        346     --        --
1.38.38.0/24          10.20.1.22         --        624     --        --
1.38.39.0/24          10.20.1.22         --        625     --        --
1.38.138.0/24         10.20.1.22         --        368     --        --
1.38.139.0/24         10.20.1.22         --        626     --        --
...
222.0.0.0/10          10.20.1.22         --        298     --        --
222.0.0.0/10          110.20.1.5         --        128741  --        --
222.9.69.0/24         10.20.1.22         --        1402    --        --
222.9.69.0/24         110.20.1.5         --        128778  --        --
-----
No. of Prefix Bindings: 2939
=====
*A:SRU4>config>router>ldp#

*A:SRU4# show router ldp bindings fec-type p2mp
=====
LDP LSR ID: 110.20.1.4
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
WP - Label Withdraw Pending
=====
LDP P2MP Bindings
=====
P2MP-Id RootAddr Interface Peer IngLbl EgrLbl EgrIntf/EgrNextHop LspId
-----
8193  110.20.1.4 74728  10.100.10.10 -- 5738 3/1/6 180.100.4.10
8194  110.20.1.4 74729  10.100.10.10 -- 7089 3/1/6 180.100.4.10
8195  110.20.1.4 74730  10.100.10.10 -- 5766 3/1/6 180.100.4.10
8196  110.20.1.4 74731  10.100.10.10 -- 6826 3/1/6 180.100.4.10
8197  110.20.1.4 74732  10.100.10.10 -- 7412 3/1/6 180.100.4.10
8198  110.20.1.4 74733  10.100.10.10 -- 6548 3/1/6 180.100.4.10
8199  110.20.1.4 74734  10.100.10.10 -- 6544 3/1/6 180.100.4.10
8200  110.20.1.4 74735  10.100.10.10 -- 5736 3/1/6 180.100.4.10
8201  110.20.1.4 74736  10.100.10.10 -- 7718 3/1/6 180.100.4.10
Press any key to continue (Q to quit)
A:both2# show router ldp bindings fec-type services
=====
LDP LSR ID: 1.1.1.57
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
       S - Status Signaled Up, D - Status Signaled Down
       E - Epipe Service, V - VPLS Service, M - Mirror Service
       A - Apipe Service, F - Fpipe Service, I - IES Service, R - VPRN service
       P - Ipipe Service, WP - Label Withdraw Pending, C - Cpipe Service
       TLV - (Type, Length: Value)
=====
LDP Service Bindings
=====
Type  VCId      SvcId      SDPId Peer                IngLbl  EgrLbl  LMTU  RMTU
-----

```

```

E-Eth 100      1      1      1.1.1.30      131067U 131068S 1500 1500
E-Eth 500      5      1      1.1.1.30      131066W 131066  3960 3960
-----
No. of VC Labels: 2
=====
A:both2#

*A:SRU4# show router ldp bindings session 10.8.100.15
=====
LDP LSR ID: 110.20.1.4
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
       S - Status Signaled Up, D - Status Signaled Down
       E - Epipe Service, V - VPLS Service, M - Mirror Service
       A - Apipe Service, F - Fpipe Service, I - IES Service, R - VPRN service
       P - Ipipe Service, WP - Label Withdraw Pending, C - Cpipe Service
       TLV - (Type, Length: Value)
=====
LDP Prefix Bindings
=====
Prefix          Peer          IngLbl    EgrLbl  EgrIntf    EgrNextHop
-----
No Matching Entries Found
=====
LDP Service FEC 128 Bindings
=====
Type  VCId      SvcId      SDPIId Peer          IngLbl  EgrLbl  LMTU  RMTU
-----
No Matching Entries Found
=====
LDP Service FEC 129 Bindings
=====
AGI
Type          SvcId      SDPIId Peer          IngLbl  EgrLbl  LMTU  RMTU
-----
No Matching Entries Found
=====
*A:SRU4#

A:ALA-12# show router ldp bindings ingress-label 2048 131071
=====
LDP LSR ID: 10.20.1.10
=====
Legend: U - Label In Use, R - Label Released
=====
LDP Prefix Bindings
=====
Prefix          Peer          IngLbl    EgrLbl  EgrIntf    EgrNextHop
-----
10.20.1.10/32   10.20.1.3    131069U   --      --          --
-----
No. of Prefix Bindings: 1
=====
LDP Service Bindings
=====
Type  VCId      SvcId      SDPIId Peer          IngLbl  EgrLbl  LMTU  RMTU
-----
No Matching Entries Found
=====

```

Show Commands

```
A:ALA-12#

*A:SRU4>config>router>ldp# show router ldp bindings active
=====
Legend: (S) - Static
=====
LDP Prefix Bindings (Active)
=====
Prefix                Op   IngLbl   EgrLbl   EgrIntf/LspId  EgrNextHop
-----
10.20.1.20/32         Push  --      0        3/2/1          10.100.30.20
10.20.1.20/32         Swap 131041  0        3/2/1          10.100.30.20
10.20.1.22/32         Push  --      0        3/2/7          160.60.60.2
10.20.1.22/32         Swap 131039  0        3/2/7          160.60.60.2
10.161.201.0/24      Push  --      0        3/2/7          160.60.60.2
10.161.201.0/24      Swap 131038  0        3/2/7          160.60.60.2
110.20.1.1/32        Push  --      441790    3/2/1          10.100.30.20
110.20.1.1/32        Swap 131045  441790    3/2/1          10.100.30.20
110.20.1.4/32        Pop   3        --        --             --
160.60.70.0/24       Push  --      0        3/2/7          160.60.60.2
160.60.70.0/24       Swap 131036  0        3/2/7          160.60.60.2
160.60.80.0/24       Push  --      0        3/2/7          160.60.60.2
160.60.80.0/24       Swap 129982  0        3/2/7          160.60.60.2
-----
No. of Prefix Bindings: 13
*A:SRU4>config>router>ldp#

A:Dut-B# show router ldp bindings prefix 10.20.1.3/32 detail
=====
LDP LSR ID: 10.20.1.2
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
      WP - Label Withdraw Pending
=====
LDP Prefix Binding
=====
Prefix                : 10.20.1.3/32
-----
Ing Lbl               : 130996U                               Peer           : 1.1.2.2
Egr Lbl               : 131017
Egr Int/LspId        : --
EgrNextHop           : --
Egr. Flags           : None                               Ing. Flags      : None
-----
Prefix                : 10.20.1.3/32
-----
Ing Lbl               : --                               Peer           : 3.3.2.2
Egr Lbl               : 3
Egr Int/LspId        : --
EgrNextHop           : --
Egr. Flags           : None                               Ing. Flags      : None
-----
Prefix                : 10.20.1.3/32
-----
Ing Lbl               : 130996U                               Peer           : 4.4.2.2
Egr Lbl               : 130969
Egr Int/LspId        : --
EgrNextHop           : --
Egr. Flags           : None                               Ing. Flags      : None
-----
```

```

Prefix          : 10.20.1.3/32
-----
Ing Lbl         : 130996U           Peer          : 5.5.2.2
Egr Lbl         : 131005
Egr Int/LspId   : --
EgrNextHop      : --
Egr. Flags     : None             Ing. Flags    : None
-----
Prefix          : 10.20.1.3/32
-----
Ing Lbl         : 130996U           Peer          : 6.6.2.2
Egr Lbl         : 130993
Egr Int/LspId   : --
EgrNextHop      : --
Egr. Flags     : None             Ing. Flags    : None
-----
Prefix          : 10.20.1.3/32
-----
Ing Lbl         : 130996U           Peer          : 10.10.1.1
Egr Lbl         : 131017
Egr Int/LspId   : --
EgrNextHop      : --
Egr. Flags     : None             Ing. Flags    : None
-----
Prefix          : 10.20.1.3/32
-----
Ing Lbl         : --                Peer          : 10.10.3.3
Egr Lbl         : 3
Egr Int/LspId   : --
EgrNextHop      : --
Egr. Flags     : None             Ing. Flags    : None
-----
Prefix          : 10.20.1.3/32
-----
Ing Lbl         : 130996U           Peer          : 10.10.4.4
Egr Lbl         : 130969
Egr Int/LspId   : --
EgrNextHop      : --
Egr. Flags     : None             Ing. Flags    : None
-----
Prefix          : 10.20.1.3/32
-----
Ing Lbl         : --                Peer          : 10.10.12.3
Egr Lbl         : 3
Egr Int/LspId   : lag-1
EgrNextHop      : 10.10.12.3
Egr. Flags     : None             Ing. Flags    : None
Metric          : 333              Mtu           : 1500
=====
No. of Prefix Bindings: 9
=====
*A:Dut-B#

```

discovery

- Syntax** `discovery [{peer [ip-address]} | {interface [ip-int-name]}] [state state] [detail] [adjacency-type type]`
- Context** show>router>ldp
- Description** This command displays the status of the interfaces participating in LDP discovery.
- Parameters**
- peer** *ip-address* — Specifies to display the IP address of the peer.
 - interface** *ip-int-name* — The name of an existing interface. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.
 - state** *state* — Specifies to display the current operational state of the adjacency.
 - Values** established, trying, down
 - detail** — Specifies to display detailed information.
 - adjacency-type** *type* — Specifies to display the adjacency type.
 - Values** link, targeted
- Output** **LDP Discovery Output** — The following table describes LDP discovery output fields.

Label	Description
Interface Name	The name of the interface.
Local Addr	The IP address of the originating (local) router.
Peer Addr	The IP address of the peer.
Adj Type	The adjacency type between the LDP peer and LDP session is targeted.
State	Established — The adjacency is established. Trying — The adjacency is not yet established.
No. of Hello Adjacencies	The total number of hello adjacencies discovered.
Up Time	The amount of time the adjacency has been enabled.
Hold-Time Remaining	The time left before a neighbor is declared to be down.
Hello Mesg Recv	The number of hello messages received for this adjacency.
Hello Mesg Sent	The number of hello messages that have been sent for this adjacency.
Remote Cfg Seq No	The configuration sequence number that was in the hello received when this adjacency started up. This configuration sequence number changes when there is a change of configuration.
Remote IP Address	The IP address used on the remote end for the LDP session.
Local Cfg Seq No	The configuration sequence number that was used in the hello sent when this adjacency started up. This configuration sequence number changes when there is a change of configuration.

Label	Description (Continued)
Local IP Address	The IP address used locally for the LDP session.

Sample Output

```
*A:Dut-A# \show router ldp discovery peer 10.20.1.2 detail
=====
LDP Hello Adjacencies (Detail)
=====
-----
Peer 10.20.1.2
-----
Local Address      : 10.20.1.1          Peer Address      : 10.20.1.2
Adjacency Type    : Targeted          State             : Established
Up Time           : 0d 00:01:01        Hold Time Remaining : 59
Hello Mesg Recv   : 9                Hello Mesg Sent    : 8
Local IP Address  : 10.20.1.1          Remote IP Address  : 10.20.1.2
Local Hello Timeout: 60              Remote Hello Timeout: 60
Local Cfg Seq No  : 3601982061        Remote Cfg Seq No  : 894145223
=====

*A:SRU4>config>router>ldp# show router ldp discovery
=====
LDP Hello Adjacencies
=====
-----
Interface Name          Local Addr      Peer Addr      AdjType State
-----
N/A                     110.20.1.4     10.8.100.15   Targ   Estab
N/A                     110.20.1.4     10.20.1.20    Targ   Estab
N/A                     110.20.1.4     10.20.1.22    Targ   Estab
N/A                     110.20.1.4     10.100.1.1    Targ   Trying
N/A                     110.20.1.4     110.20.1.1    Targ   Estab
N/A                     110.20.1.4     110.20.1.2    Targ   Trying
N/A                     110.20.1.4     110.20.1.3    Targ   Estab
N/A                     110.20.1.4     110.20.1.5    Targ   Estab
N/A                     110.20.1.4     110.20.1.6    Targ   Trying
N/A                     110.20.1.4     110.20.1.51   Targ   Trying
N/A                     110.20.1.4     110.20.1.52   Targ   Trying
N/A                     110.20.1.4     110.20.1.53   Targ   Trying
N/A                     110.20.1.4     110.20.1.55   Targ   Trying
N/A                     110.20.1.4     110.20.1.56   Targ   Trying
N/A                     110.20.1.4     110.20.1.110  Targ   Trying
N/A                     110.20.1.4     110.20.1.150  Targ   Trying
N/A                     110.20.1.4     220.220.1.6   Targ   Trying
aps-1                   110.20.1.4     110.20.1.3    Link   Estab
aps-2                   110.20.1.4     110.20.1.3    Link   Estab
aps-3                   110.20.1.4     110.20.1.3    Link   Estab
sr4-1                   110.20.1.4     110.20.1.3    Link   Estab
ess-7-1                 110.20.1.4     110.20.1.5    Link   Estab
ess-7-2                 110.20.1.4     110.20.1.5    Link   Estab
ess-7-3                 110.20.1.4     110.20.1.5    Link   Estab
ess-7-4                 110.20.1.4     110.20.1.5    Link   Estab
ess-7-5                 110.20.1.4     110.20.1.5    Link   Estab
hubA                    110.20.1.4     110.20.1.3    Link   Estab
hubA                    110.20.1.4     110.20.1.5    Link   Estab
```

Show Commands

```

hubA                110.20.1.4      200.0.0.1      Link  Estab
germ-1              110.20.1.4      110.20.1.110   Link  Estab
src-1.1             170.70.51.4     224.0.0.2      Link  Trying
src-1.2             170.70.52.4     224.0.0.2      Link  Trying
src-1.3             170.70.53.4     224.0.0.2      Link  Trying
src-1.4             170.70.54.4     224.0.0.2      Link  Trying
src-1.5             170.70.55.4     224.0.0.2      Link  Trying
src-1.6             170.70.56.4     224.0.0.2      Link  Trying
src-1.7             170.70.57.4     224.0.0.2      Link  Trying
src-1.8             170.70.58.4     224.0.0.2      Link  Trying
src-1.9             170.70.59.4     224.0.0.2      Link  Trying
src-1.10            170.70.60.4     224.0.0.2      Link  Trying
srl-1               110.20.1.4      33.66.32.1     Link  Estab
srl-3               110.20.1.4      33.66.33.1     Link  Estab
aps-8               110.20.1.4      33.66.34.1     Link  Estab
aps-9               110.20.1.4      33.66.35.1     Link  Estab
srr-1               110.20.1.4      11.22.10.2     Link  Estab
srr-2               110.20.1.4      11.22.11.2     Link  Estab
srr-3               110.20.1.4      1.1.1.1        Link  Estab
aps-11              110.20.1.4      11.22.13.2     Link  Estab
gsrl                110.20.1.4      10.8.100.15    Link  Estab
gsr2                110.20.1.4      10.20.1.22     Link  Estab
g7600               180.50.80.4     224.0.0.2      Link  Trying
ml60                110.20.1.4      10.20.1.20     Link  Estab

```

No. of Hello Adjacencies: 52

=====

```
*A:SRU4>config>router>ldp#
```

```
*A:SRU4>config>router>ldp# show router ldp discovery peer 10.8.100.15
```

=====

```
LDP Hello Adjacencies
```

```

-----
Interface Name          Local Addr      Peer Addr      AdjType State
-----
N/A                     110.20.1.4     10.8.100.15   Targ   Estab
-----

```

No. of Hello Adjacencies: 1

=====

```
*A:SRU4>config>router>ldp#
```

```
*A:SRU4>config>router>ldp# show router ldp discovery detail
```

=====

```
LDP Hello Adjacencies (Detail)
```

Peer 10.8.100.15

```

-----
Local Address      : 110.20.1.4          Peer Address       : 10.8.100.15
Adjacency Type    : Targeted           State              : Established
Up Time           : 0d 12:39:38       Hold Time Remaining : 43
Hello Mesg Recv   : 10403             Hello Mesg Sent    : 3243
Local IP Address  : 110.20.1.4       Remote IP Address  : 10.8.100.15
Local Hello Timeout: 45          Remote Hello Timeout: 90
Local Cfg Seq No  : 1828354504      Remote Cfg Seq No  : 0
-----

```

Peer 10.20.1.20

```

-----
Local Address      : 110.20.1.4          Peer Address       : 10.20.1.20
-----

```

```

Adjacency Type      : Targeted          State           : Established
Up Time             : 0d 12:39:57       Hold Time Remaining : 40
Hello Mesg Recv    : 7495              Hello Mesg Sent   : 3244
Local IP Address    : 110.20.1.4       Remote IP Address  : 10.20.1.20
Local Hello Timeout: 45                Remote Hello Timeout: 45
Local Cfg Seq No   : 572902976        Remote Cfg Seq No  : 1

```

```

-----
...
-----

```

```

Interface "gsr2"
-----

```

```

Local Address       : 110.20.1.4       Peer Address      : 10.20.1.22
Adjacency Type     : Link             State             : Established
Up Time            : 0d 12:40:41       Hold Time Remaining : 11
Hello Mesg Recv    : 10414            Hello Mesg Sent   : 11260
Local IP Address    : 160.60.60.4     Remote IP Address  : 160.60.60.2
Local Hello Timeout: 15                Remote Hello Timeout: 15
Local Cfg Seq No   : 1911286684      Remote Cfg Seq No  : 0

```

```

-----
Interface "g7600"
-----

```

```

Local Address       : 180.50.80.4     Peer Address      : 224.0.0.2
Adjacency Type     : Link             State             : Trying

```

```

-----
Interface "m160"
-----

```

```

Local Address       : 110.20.1.4       Peer Address      : 10.20.1.20
Adjacency Type     : Link             State             : Established
Up Time            : 0d 12:40:47       Hold Time Remaining : 14
Hello Mesg Recv    : 10450            Hello Mesg Sent   : 11262
Local IP Address    : 10.100.30.4     Remote IP Address  : 10.100.30.20
Local Hello Timeout: 15                Remote Hello Timeout: 15
Local Cfg Seq No   : 2523051834      Remote Cfg Seq No  : 1

```

```

=====
*A:SRU4>config>router>ldp#

```

interface

Syntax `interface [ip-int-name | ip-address] [detail]`

Context show>router>ldp

Description This command displays configuration information about LDP interfaces.

Parameters *ip-int-name* — The name of an existing interface. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

ip-address — The IP address of the LDP neighbor.

detail — Displays detailed information.

Output **LDP Interface Output** — The following table describes the LDP interface output fields.

Label	Description
Interface	Specifies the interface associated with the LDP instance.
Adm	Up — The LDP is administratively enabled. Down — The LDP is administratively disabled.
Opr	Up — The LDP is operationally enabled. Down — The LDP is operationally disabled.
Hello Factor	The value by which the hello timeout should be divided to give the hello time, for example, the time interval, in seconds, between LDP hello messages. LDP uses hello messages to discover neighbors and to detect loss of connectivity with its neighbors.
Hold Time	The hello time, also known as hold time. It is the time interval, in seconds, that LDP waits before declaring a neighbor to be down. Hello timeout is local to the system and is sent in the hello messages to a neighbor.
KA Factor	The value by which the keepalive timeout should be divided to give the keepalive time, for example, the time interval, in seconds, between LDP keepalive messages. LDP keepalive messages are sent to keep the LDP session from timing out when no other LDP traffic is being sent between the neighbors.
KA Timeout	The time interval, in seconds, that LDP waits before tearing down a session. If no LDP messages are exchanged during this time interval, the LDP session is torn down. Generally the value is configured to be 3 times the keepalive time (the time interval between successive LDP keepalive messages).
Auth	Enabled — Authentication using MD5 message based digest protocol is enabled. Disabled — No authentication is used.
No. of Interface	The total number of LDP interfaces.

Sample Output

```
*A:SICILY# show router ldp interface "srr" detail
=====
LDP Interfaces (Detail)
=====
-----
Interface "srr"
-----
Admin State       : Up                Oper State       : Down
Oper Down Reason  : noResources           iildp
Hold Time        : 140                Hello Factor     : 3
Oper Hold Time   : 140                Hello Reduction  *: 3
Hello Reduction  : Disabled           Keepalive Factor : 3
Keepalive Timeout : 140                Last Modified    : 07/11/13 02:27:53
Transport Addr   : System            Template Name    : N/A
Active Adjacencies : 0
Creator          : manual
Tunneling        : Disabled
Lsp Name         : None
Local LSR Type   : Interface
Local LSR        : None
BFD Status       : Disabled
Multicast Traffic : Enabled
=====
* indicates that the corresponding row element may have been truncated.
*A:SICILY#
```

```
*A:SICILY# show router ldp interface resource-failures
=====
LDP Interface Resource Failures
=====
srr
=====
*A:SICILY#
```

```
*A:SR4# show router ldp interface detail
=====
LDP Interfaces (Detail)
=====
-----
Interface "aps-1"
-----
Admin State       : Up                Oper State       : Down
Oper Down Reason  : instanceDown
Hold Time        : 45                Hello Factor     : 3
Keepalive Timeout : 30                Keepalive Factor : 3
Transport Addr   : System            Last Modified    : 07/26/11 02:09:50
Active Adjacencies : 0
Tunneling        : Disabled
Lsp Name         : None
Local LSR        : System
BFD Status       : Disabled
Multicast Traffic : Enabled
=====
```

```
*A:SRU4>config>router>ldp# show router ldp interface hubA
=====
```

Show Commands

```
LDP Interfaces
=====
Interface                               Adm Opr  Hello Hold  KA    KA    Transport
                               Factor Time  Factor Timeout Address
-----
hubA                                   Up  Up   3    15   3     30   System
-----
No. of Interfaces: 1
=====
*A:SRU4>config>router>ldp#
```

```
*A:SRU4>config>router>ldp# show router ldp interface hubA detail
=====
LDP Interfaces (Detail)
=====
-----
Interface "hubA"
-----
Admin State       : Up                Oper State       : Up
Hold Time        : 15                Hello Factor    : 3
Keepalive Timeout : 30                Keepalive Factor : 3
Transport Addr   : System            Last Modified   : 03/03/2010 19:47:34
Active Adjacencies : 3
Tunneling        : Disabled
Lsp Name         : None
Local LSR        : System
BFDD Status      : Disabled
=====
*A:SRU4>config>router>ldp#
```

Step 1: Configure loop back interface on router. For example, "ip-100.100.100.100"

Step 2: Assign this loopback to appropriate IGP Area/Level.

Step 3: Assign this loopback address as local-lsr-id using command under LDP interface (ldp interface parameters):

config>router>ldp>interface-parameters>interface ip-10.10.25.2\$local-lsr-id "ip-100.100.100.100"

```
*A:Dut-B>config>router>ldp# show router ldp interface "ip-10.10.25.2" detail
=====
LDP Interfaces (Detail)
=====
-----
Interface "ip-10.10.25.2"
-----
Admin State       : Up                Oper State       : Up
Hold Time        : 15                Hello Factor    : 3
Oper Hold Time   : 15
Hello Reduction  : Disabled          Hello Reduction *: 3
Keepalive Timeout : 30                Keepalive Factor : 3
Transport Addr   : System            Last Modified   : 01/24/13 23:45:42
Active Adjacencies : 1
Tunneling        : Disabled
Lsp Name         : None
Local LSR Type   : Interface
Local LSR        : ip-100.100.100.100 // Local LSR Id is Loopback.
BFDD Status      : Enabled
Multicast Traffic : Enabled
```

```

===== * indicates
that the corresponding row element may have been truncated.
*A:Dut-B>config>router>ldp#

```

fec-egress-stats

- Syntax** `fec-egress-stats [ip-prefix/mask]`
fec-egress-stats active
- Context** show>router>ldp
- Description** This command displays LDP prefix FECs egress statistics.
- Parameters** *ip-prefix* — Specify information for the specified IP prefix and mask length. Host bits must be 0.
mask — Specifies the 32-bit address mask used to indicate the bits of an IP address that are being used for the subnet address.
- Values** 0 — 32

fec-originate

- Syntax** **fec-originate** [*ip-prefix/mask*] [*operation-type*]
- Context** show>router>ldp
- Description** This command displays LDP static prefix FECs.
- Parameters**
- ip-prefix* — Specify information for the specified IP prefix and mask length. Host bits must be 0.
 - mask* — Specifies the 32-bit address mask used to indicate the bits of an IP address that are being used for the subnet address.
- Values** 0 — 32
- operation-type* — Specify the operation type to display.
- Values** pop, swap
- Output** **FEC Originate Output** — The following table describes the FEC originate parameters output fields.

Label	Description
Prefix	Specifies the static prefix FEC.
NH Type	Specifies the type of next-hop represented by this row entry. unknown — The next-hop type has not been set. IP Addr — The next-hop is an IP address. pop — There is no next-hop (pop the label and route).
NextHop	The IP address of the next-hop.
IngLabel	Specifies the label that is advertised to the upstream peer. If this variable is set to the default value of 4294967295, the ingress label will be dynamically assigned by the label manager.
EgrLabel	Specifies the egress label associated with this next-hop entry. The LSR will swap the incoming label with the configured egress label. If this egress label has a value of 4294967295, the LSR will pop the incoming label.
OperIngLabel	Specifies the actual or operational value of the label that was advertised to the upstream peer.

```
*A:SRU4>config>router>ldp# show router ldp fec-originate
=====
LDP Static Prefix FECs
=====
Prefix           NHType  NextHop      IngLabel    EgrLabel    OperIngLabel
-----
24.1.1.0.0/16    Pop     n/a          --          --          0
24.1.1.0.1/32    Pop     n/a          --          --          0
24.1.1.0.2/32    Pop     n/a          --          --          0
24.1.1.0.3/32    Pop     n/a          --          --          0
24.1.1.0.4/32    Pop     n/a          --          --          0
```



```

24.1.0.5/32      Pop      n/a      --      --      0
24.1.0.6/32      Pop      n/a      --      --      0
24.1.0.7/32      Pop      n/a      --      --      0
24.1.0.8/32      Pop      n/a      --      --      0
24.1.0.9/32      Pop      n/a      --      --      0
...
24.251.0.0/16    Pop      n/a      --      --      0
24.252.0.0/16    Pop      n/a      --      --      0
24.253.0.0/16    Pop      n/a      --      --      0
24.254.0.0/16    Pop      n/a      --      --      0

```

```
-----
No. of FECs: 508
=====
```

```
*A:SRU4>config>router>ldp#
```

parameters

Syntax	parameters
Context	show>router>ldp
Description	This command displays configuration information about LDP parameters.
Output	LDP Parameters Output — The following table describes the LDP parameters output fields.

Label	Description
Keepalive Timeout	The time interval, in seconds, that LDP waits before tearing down a session. If no LDP messages are exchanged during this time interval, the LDP session is torn down. Generally the value is configured to be 3 times the keepalive time (the time interval between successive LDP keepalive messages).
Timeout Factor	The value by which the keepalive timeout should be divided to give the keepalive time, for example, the time interval, in seconds, between LDP keepalive messages. LDP keepalive messages are sent to keep the LDP session from timing out when no other LDP traffic is being sent between the neighbors.
Hold Time	The hello time, also known as hold time. It is the time interval, in seconds, that LDP waits before declaring a neighbor to be down. Hello timeout is local to the system and is sent in the hello messages to a neighbor.
Hello Factor	The value by which the hello timeout should be divided to give the hello time, for example, the time interval, in seconds, between LDP hello messages. LDP uses hello messages to discover neighbors and to detect loss of connectivity with its neighbors.
Auth	Enabled — Authentication using MD5 message based digest protocol is enabled. Disabled — No authentication is used.
Admin Status	inService — The LDP is administratively enabled.

Label	Description (Continued)
	<code>outService</code> – The LDP is administratively disabled.
Deaggregated FECs	<p><code>False</code> – LDP aggregates multiple prefixes into a single Forwarding Equivalence Class (FEC) and advertises a single label for the FEC. This value is only applicable to LDP interfaces and not for targeted sessions.</p> <p><code>True</code> – LDP de-aggregates prefixes into multiple FECs.</p>
Propagate Policy	<p>The Propagate Policy value specifies whether the LSR should generate FECs and which FECs it should generate.</p> <p><code>system</code> – LDP will distribute label bindings only for the router's system IP address.</p> <p><code>interface</code> – LDP will distribute label bindings for all LDP interfaces.</p> <p><code>all</code> – LDP will distribute label bindings for all prefixes in the routing table.</p> <p><code>none</code> – LDP will not distribute any label bindings.</p>
Transport Address	<p><code>interface</code> – The interface's IP address is used to set up the LDP session between neighbors. If multiple interfaces exist between two neighbors, the 'interface' mode cannot be used since only one LDP session is actually set up between the two neighbors.</p> <p><code>system</code> – The system's IP address is used to set up the LDP session between neighbors.</p>
Label-Retention	<p><code>liberal</code> – All advertised label mappings are retained whether they are from a valid next hop or not. When the label distribution value is downstream unsolicited, a router may receive label bindings for the same destination for all its neighbors. Labels for the non-next hops for the FECs are retained in the software but not used. When a network topology change occurs where a non-nexthop becomes a true next hop, the label received earlier is then used.</p> <p><code>conservative</code> – Advertised label mappings are retained only if they will be used to forward packets; for example if the label came from a valid next hop. Label bindings received from non-next hops for each FEC are discarded.</p>
Control Mode	<p><code>ordered</code> – Label bindings are not distributed in response to a label request until a label binding has been received from the next hop for the destination.</p> <p><code>independent</code> – Label bindings are distributed immediately in response to a label request even if a label binding has not yet been received from the next hop for the destination.</p>

Label	Description (Continued)
Route Preference	The route preference assigned to LDP routes. When multiple routes are available to a destination, the route with the lowest preference will be used. This value is only applicable to LDP interfaces and not for targeted sessions.
Loop Detection	<p><code>none</code> – Loop detection is not supported on this router. This is the only valid value since Path Vector based loop detection is not supported.</p> <p><code>other</code> – Loop detection is supported but by a method other than <code>hopCount</code>, <code>pathVector</code>, or <code>hopCountAndPathVector</code>.</p> <p><code>hopCount</code> – Loop detection is supported by hop count only.</p> <p><code>pathVector</code> – Loop detection is supported by path vector only.</p> <p><code>hopCountAndPathVector</code> – Loop detection is supported by both path vector and hop count.</p>
Keepalive Timeout	The factor used to derive the Keepalive interval.
Keepalive Factor	The time interval, in seconds, that LDP waits before tearing down the session.
Hold-Time	The time left before a neighbor is declared to be down.
Hello Factor	The value by which the hello timeout should be divided to give the hello time, for example, the time interval, in seconds, between LDP hello messages. LDP uses hello messages to discover neighbors and to detect loss of connectivity with its neighbors.
Auth	<p><code>Enabled</code> – Authentication using MD5 message based digest protocol is enabled.</p> <p><code>Disabled</code> – No authentication is used.</p>
Passive-Mode	<p><code>true</code> – LDP responds only when it gets a connect request from a peer and will not attempt to actively connect to its neighbors.</p> <p><code>false</code> – LDP actively tries to connect to its peers.</p>
Targeted-Sessions	<p><code>true</code> – Targeted sessions are enabled.</p> <p><code>false</code> – Targeted sessions are disabled.</p>

Sample Output

```
*A:Dut-A>config>router>ldp# \show router ldp parameters
=====
LDP Parameters (LSR ID 10.20.1.1)
=====
-----
Graceful Restart Parameters
-----
Nbor Liveness Time : 120 sec           Max Recovery Time   : 120
```

Show Commands

```
-----
Interface Parameters
-----
Keepalive Timeout   : 30 sec           Keepalive Factor    : 3
Hold Time          : 15 sec           Hello Factor        : 3
Propagate Policy   : system           Transport Address   : system
Deaggregate FECs   : False            Route Preference    : 9
Label Distribution : downstreamUnsolic* Label Retention     : liberal
Control Mode       : ordered           Loop Detection      : none
-----

Targeted Session Parameters
-----
Keepalive Timeout   : 30 sec           Keepalive Factor    : 3
Hold Time          : 15 sec           Hello Factor        : 3
Hello Reduction     : Enabled          Hello Reduction Fctr: 3
Passive Mode        : False            Targeted Sessions   : Enabled
=====
* indicates that the corresponding row element may have been truncated.

*A:SRU4>config>router>ldp# show router ldp parameters
=====
LDP Parameters (LSR ID 110.20.1.4)
=====
-----
Graceful Restart Parameters
-----
Nbor Liveness Time : 5 sec             Max Recovery Time : 30
-----

Interface Parameters
-----
Keepalive Timeout   : 30 sec           Keepalive Factor    : 3
Hold Time          : 15 sec           Hello Factor        : 3
Propagate Policy   : system           Transport Address   : system
Deaggregate FECs   : False            Route Preference    : 9
Label Distribution : downstreamUnsolicited Label Retention     : liberal
Control Mode       : ordered           Loop Detection      : none
-----

Targeted Session Parameters
-----
Keepalive Timeout   : 40 sec           Keepalive Factor    : 4
Hold Time          : 45 sec           Hello Factor        : 3
Passive Mode        : False            Targeted Sessions   : Enabled
=====
*A:SRU4>config>router>ldp#
```

peer

- Syntax** `peer [ip-address] [detail]`
- Context** `show>router>ldp`
- Description** This command displays configuration information about LDP peers.
- Parameters** *ip-address* — The IP address of the LDP peer.
detail — Displays detailed information.

Output **LDP Peer Output** — The following table describes LDP peer output.

Label	Description
Peer	The IP address of the peer.
Adm	Up — The LDP is administratively enabled. Down — The LDP is administratively disabled.
Opr	Up — The LDP is operationally enabled. Down — The LDP is operationally disabled.
Hello Factor	The value by which the hello timeout should be divided to give the hello time, for example, the time interval, in seconds, between LDP hello messages. LDP uses hello messages to discover neighbors and to detect loss of connectivity with its neighbors.
Hold Time	The hello time or hold time. The time interval, in seconds, that LDP waits before declaring a neighbor to be down. Hello timeout is local to the system and is sent in the hello messages to a neighbor.
KA Factor	The value by which the keepalive timeout should be divided to give the keepalive time, for example, the time interval, in seconds, between LDP keepalive messages. LDP keepalive messages are sent to keep the LDP session from timing out when no other LDP traffic is being sent between the neighbors.
KA Timeout	The time interval, in seconds, that LDP waits before tearing down a session. If no LDP messages are exchanged during this time interval, the LDP session is torn down. Generally the value is configured to be 3 times the keepalive time (the time interval between successive LDP keepalive messages).
Auth	Enabled — Authentication using MD5 message based digest protocol is enabled. Disabled — No authentication is used.
Passive Mode	The mode used to set up LDP sessions. This value is only applicable to targeted sessions and not to LDP interfaces. True — LDP responds only when it gets a connect request from a peer and will not attempt to actively connect to its neighbors. False — LDP actively tries to connect to its peers.
Auto Create	Specifies if a targeted peer was automatically created through service manager. For an LDP interface, this value is always false.
No. of Peers	The total number of LDP peers.
Tunneling	Enabled — Tunneling is enabled. Disabled — No tunneling is used.

Label	Description (Continued)
LSP	The LSP name.

Sample Output

```
*A:SICILY# show router ldp peer 110.20.1.4 detail
=====
LDP Peers (Detail)
=====
-----
Peer 110.20.1.4
-----
Admin State      : Up           Oper State      : Down
Oper Down Reason : noResources
Hold Time       : 140          Hello Factor    : 3
Oper Hold Time  : 140
Hello Reduction  : Disabled     Hello Reduction Fact*: 3
Keepalive Timeout : 140        Keepalive Factor : 3
Passive Mode     : Disabled     Last Modified   : 07/11/13 00:36:09
Active Adjacencies : 0          Auto Created    : No
Creator         : manual       Template Name   : N/A
Tunneling       : Enabled
Lsp Name        : None
Local LSR       : None
BFD Status      : Enabled
Multicast Traffic : Disabled
=====
* indicates that the corresponding row element may have been truncated.
*A:SICILY#
```

```
*A:SR1-A# /show router ldp peer detail
=====
LDP Peers (Detail)
=====
-----
Peer 2.2.2.2
-----
Admin State      : Up           Oper State      : Up
Hold Time       : 45          Hello Factor    : 3
Oper Hold Time  : 45
Hello Reduction  : Disabled     Hello Reduction Fact*: 3
Keepalive Timeout : 40        Keepalive Factor : 4
Passive Mode     : Disabled     Last Modified   : 07/17/13 21:08:20
Active Adjacencies : 1          Auto Created    : No
Creator         : manual       Template Name   : N/A
Tunneling       : Disabled
Lsp Name        : None
Local LSR       : None
BFD Status      : Disabled
Multicast Traffic : Disabled
-----
Peer 3.3.3.3
-----
Admin State      : Up           Oper State      : Up
Hold Time       : 45          Hello Factor    : 3
Oper Hold Time  : 45
Hello Reduction  : Disabled     Hello Reduction Fact*: 3
```

```

Keepalive Timeout : 40           Keepalive Factor : 4
Passive Mode      : Disabled     Last Modified    : 07/17/13 21:08:20
Active Adjacencies : 1          Auto Created     : Yes
Creator          : svcMgr        Template Name    : N/A
Tunneling        : Disabled
Lsp Name         : None
Local LSR        : None
BFD Status       : Disabled
Multicast Traffic : Disabled

```

```
-----
Peer 4.4.4.4
-----
```

```

Admin State      : Up           Oper State       : Up
Hold Time        : 45           Hello Factor     : 3
Oper Hold Time   : 45
Hello Reduction  : Disabled     Hello Reduction Fact*: 3
Keepalive Timeout : 40         Keepalive Factor : 4
Passive Mode     : Disabled     Last Modified    : 07/17/13 21:11:29
Active Adjacencies : 0        Auto Created     : Yes
Creator          : template    Template Name    : templ1
Tunneling        : Disabled
Lsp Name         : None
Local LSR        : None
BFD Status       : Disabled
Multicast Traffic : Disabled

```

```

*A:Dut-A>config>router>ldp# \show router ldp peer 10.20.1.2 detail

```

```

=====
LDP Peers (Detail)
=====
-----
```

```
Peer 10.20.1.2
-----
```

```

Admin State      : Up           Oper State       : Up
Hold Time        : 15           Hello Factor     : 3
Oper Hold Time   : 120
Hello Reduction  : Enabled     Hello Reduction Fact*: 3
Keepalive Timeout : 30         Keepalive Factor : 3
Passive Mode     : Disabled     Last Modified    : 01/23/13 23:16:53
Active Adjacencies : 1        Auto Created     : No
Tunneling        : Enabled
Lsp Name         : None
Local LSR        : None
BFD Status       : Disabled
Multicast Traffic : Disabled

```

```

=====
* indicates that the corresponding row element may have been truncated.

```

```

*A:SR4# show router ldp peer detail

```

```

=====
LDP Peers (Detail)
=====
-----
```

```
Peer 10.8.100.15
-----
```

```

Admin State      : Up           Oper State       : Down
Oper Down Reason : instanceDown
Hold Time        : 45           Hello Factor     : 3

```

Show Commands

```
Keepalive Timeout : 40           Keepalive Factor : 4
Passive Mode      : Disabled     Last Modified    : 07/26/11 02:09:50
Active Adjacencies : 0         Auto Created     : No
Tunneling        : Enabled
Lsp Name         : None
Local LSR        : None
BFD Status       : Disabled
Multicast Traffic : Disabled
```

```
-----
*A:SRU4>config>router>ldp# show router ldp peer
```

```
=====
LDP Peers
```

```
=====
Peer           Adm  Opr  Hello  Hold   KA    KA    Passive  Auto
                Factor Time  Factor Time  Factor Timeout Mode    Created
-----
10.8.100.15    Up   Up   3      45     4     40     Disabled No
10.20.1.20     Up   Up   3      45     4     40     Disabled No
10.20.1.22     Up   Up   3      45     4     40     Disabled No
10.100.1.1     Up   Up   3      45     4     40     Disabled No
110.20.1.1     Up   Up   3      45     4     40     Disabled No
110.20.1.2     Up   Up   3      45     4     40     Disabled No
110.20.1.3     Up   Up   3      45     4     40     Disabled No
110.20.1.5     Up   Up   3      45     4     40     Disabled No
110.20.1.6     Up   Up   3      45     4     40     Disabled No
110.20.1.51    Up   Up   3      45     4     40     Disabled No
110.20.1.52    Up   Up   3      45     4     40     Disabled No
110.20.1.53    Up   Up   3      45     4     40     Disabled No
110.20.1.55    Up   Up   3      45     4     40     Disabled No
110.20.1.56    Up   Up   3      45     4     40     Disabled No
110.20.1.110   Up   Up   3      45     4     40     Disabled No
110.20.1.150   Up   Up   3      45     4     40     Disabled No
220.220.1.6    Up   Up   3      45     4     40     Disabled No
-----
```

```
No. of Peers: 17
```

```
=====
*A:SRU4>config>router>ldp#
```

```
*A:SRU4>config>router>ldp# show router ldp peer detail
```

```
=====
LDP Peers (Detail)
```

```
-----
Peer 10.8.100.15
```

```
-----
Admin State      : Up           Oper State       : Up
Hold Time        : 45           Hello Factor     : 3
Keepalive Timeout : 40         Keepalive Factor : 4
Passive Mode     : Disabled     Last Modified    : 03/03/2010 19:47:34
Active Adjacencies : 1         Auto Created     : No
Tunneling        : Disabled
Lsp Name         : None
Local LSR        : None
BFD Status       : Disabled
```

```
-----
Peer 10.20.1.20
```

```
-----
Admin State      : Up           Oper State       : Up
```



```

Hold Time           : 45           Hello Factor       : 3
Keepalive Timeout  : 40           Keepalive Factor   : 4
Passive Mode       : Disabled      Last Modified      : 03/03/2010 19:47:34
Active Adjacencies : 1           Auto Created       : No
Tunneling          : Disabled
Lsp Name           : None
Local LSR          : None
BFD Status         : Disabled
...

```

```
-----
Peer 220.220.1.6
-----
```

```

Admin State        : Up           Oper State         : Up
Hold Time          : 45           Hello Factor       : 3
Keepalive Timeout  : 40           Keepalive Factor   : 4
Passive Mode       : Disabled      Last Modified      : 03/03/2010 19:47:34
Active Adjacencies : 0           Auto Created       : No
Tunneling          : Disabled
Lsp Name           : None
Local LSR          : None
BFD Status         : Disabled

```

```
=====
*A:SRU4>config>router>ldp#
```

```
*A:SRU4>config>router>ldp# show router ldp peer 10.8.100.15 detail
```

```
=====
LDP Peers (Detail)
=====
```

```
-----
Peer 10.8.100.15
-----
```

```

Admin State        : Up           Oper State         : Up
Hold Time          : 45           Hello Factor       : 3
Keepalive Timeout  : 40           Keepalive Factor   : 4
Passive Mode       : Disabled      Last Modified      : 03/03/2010 19:47:34
Active Adjacencies : 1           Auto Created       : No
Tunneling          : Disabled
Lsp Name           : None
Local LSR          : None
BFD Status         : Disabled

```

```
=====
*A:SRU4>config>router>ldp#
```

peer-template

- Syntax** peer-template
- Context** show>router>ldp
- Description** This command displays the configured parameters of a peer-template

Sample Output

```
*A:SR1-A# show router ldp peer-template
```

Show Commands

```
=====
LDP Peer Template
=====
-----
Peer Template templ1
-----
Create Time      : 01/03/70 12:48:55  Last Modified   : 01/04/70 04:21:15
Admin State     : Up
Hello Timeout   : 45
Hello Reduction : Disabled
Keepalive Timeout : 40
Tunneling      : Disabled
Local LSR      : None
BFD Status     : Disabled
-----
Peer Template templ2
-----
Create Time      : 01/03/70 13:14:48  Last Modified   : 01/04/70 04:47:08
Admin State     : Up
Hello Timeout   : 45
Hello Reduction : Disabled
Keepalive Timeout : 40
Tunneling      : Disabled
Local LSR      : None
BFD Status     : Disabled
-----
Peer Template templ3
-----
Create Time      : 01/03/70 15:56:30  Last Modified   : 01/04/70 07:28:50
Admin State     : Up
Hello Timeout   : 45
Hello Reduction : Disabled
Keepalive Timeout : 40
Tunneling      : Disabled
Local LSR      : None
BFD Status     : Disabled
-----
Peer Template templ4
-----
Create Time      : 01/03/70 17:02:12  Last Modified   : 01/04/70 08:34:32
Admin State     : Up
Hello Timeout   : 45
Hello Reduction : Disabled
Keepalive Timeout : 40
Tunneling      : Disabled
Local LSR      : None
BFD Status     : Disabled
=====
```

peer-template-map

- Syntax** peer-template-map [tldp-peers]
- Context** show>router>ldp
- Description** This command displays peer template mappings to prefix policy.

Sample Output

```

*A:SR1-A# /show router ldp peer-template-map
=====
LDP Peer Template Map
=====
-----
Peer Template templ1
-----
Peer Policy 1           : policy1
-----
Peer Template templ2
-----
Peer Policy 1           : policy1
Peer Policy 2           : policy2
Peer Policy 3           : policy3
-----
Peer Template templ3
-----
Peer Policy 1           : policy2
=====

*A:SR1-A# /show router ldp peer-template-map tldp-peers
=====
LDP Peer Template Map TLDP Peers
=====
-----
Peer Template templ1
-----
10.0.10.1               10.0.10.2
10.0.10.3               10.0.10.4
10.0.10.5               10.0.10.6
10.0.10.7               10.0.10.8
10.0.10.9               10.0.10.10
10.0.10.11              10.0.10.12
10.0.10.13              10.0.10.14
10.0.10.15              10.0.10.16
10.0.10.17              10.0.10.18
10.0.10.19              10.0.10.20
10.0.10.21              10.0.10.22
10.0.10.23              10.0.10.24
10.0.10.25              10.0.10.26
10.0.10.27              10.0.10.28
10.0.10.29              10.0.10.30
10.0.10.31              10.0.10.32
10.0.10.33              10.0.10.34
10.0.10.35              10.0.10.36
10.0.10.37              10.0.10.38
10.0.10.39              10.0.10.40
10.0.10.41              10.0.10.42
10.0.10.43              10.0.10.44
10.0.10.45              10.0.10.46
10.0.10.47              10.0.10.48
10.0.10.49              10.0.10.50
-----
Peer Template templ3
-----
30.1.3.5                30.1.3.6
30.1.3.7                30.1.3.8

```

Show Commands

```
30.1.3.9          30.1.3.10
30.1.3.11        30.1.3.12
30.1.3.13        30.1.3.14
=====
```

peer-parameters

Syntax `peer-parameters peer-ip-address`

Context `show>router>ldp`

Description This command displays LDP peer information.

Parameters *peer-ip-address* — Specify the peer IP address.

LDP peer-parameters output — The following table describes LDP peer-parameters output.

Label	Description
Peer	The IP address of the peer.
TTL security	Enabled — LDP peering sessions protected. Disabled — LDP peering sessions unprotected.
Min-TTL-Value	Displays the minimum TTL value for an incoming packet.
Auth	Enabled — Authentication using MD5 message based digest protocol is enabled. Disabled — No authentication is used.

Sample Output

```
*A:SRR# show router ldp peer-parameters
=====
LDP Peers
=====
Peer : 10.8.100.15
-----
TTL Security      : Disabled      Min-TTL          : n/a
Authentication Key : Enabled        DOD              : Disabled
Auth key chain    : n/a
FEC129 Cisco Inter* : Disabled      Path MTU Discovery:: Disabled
Import Policies   : None          Export Policies   : None
-----
Peer : 10.20.1.20
-----
TTL Security      : Disabled      Min-TTL          : n/a
Authentication Key : Disabled      DOD              : Disabled
Auth key chain    : n/a
```

```
FEC129 Cisco Inter*: Disabled          Path MTU Discovery:: Disabled
Import Policies      : None             Export Policies     : None
```

```
-----
*A:SRU4>config>router>ldp# show router ldp peer-parameters
```

```
=====
LDP Peers
=====
```

```
-----
Peer : 10.8.100.15
```

```
-----
TTL Security      : Disabled          Min-TTL           : n/a
Authentication Key : Disabled          DOD                : Disabled
Auth key chain    : n/a
FEC129 Cisco Inter*: Disabled
Import Policies   : None             Export Policies    : None
```

```
-----
Peer : 10.20.1.20
```

```
-----
TTL Security      : Disabled          Min-TTL           : n/a
Authentication Key : Disabled          DOD                : Disabled
Auth key chain    : n/a
FEC129 Cisco Inter*: Disabled
Import Policies   : None             Export Policies    : None
```

```
-----
Peer : 220.220.1.6
```

```
-----
TTL Security      : Enabled           Min-TTL           : 255
Authentication Key : Enabled          DOD                : Disabled
Auth key chain    : n/a
FEC129 Cisco Inter*: Disabled
Import Policies   : None             Export Policies    : None
```

```
-----
No. of Peers: 17
```

```
=====
* indicates that the corresponding row element may have been truncated.
```

```
*A:SRU4>config>router>ldp#
```

session

- Syntax** `session [ip-addr[:label-space]] [detail | statistics [packet-type]] [session-type]`
- Context** show>router>ldp
- Description** This command displays configuration information about LDP sessions.
- Parameters**
- ip-address* — Specify the IP address of the LDP peer.
 - label-space* — Specifies the label space identifier that the router is advertising on the interface.
 - Values** 0 — 65535
 - detail** — Displays detailed information.
 - statistics** *packet-type* — Specify the packet type.
 - Values** hello, keepalive, init, label, notification, address
 - session-type* — Specifies to display the session type.
 - Values** link, targeted, both
- Output** **LDP Session Output** — The following table describes LDP session output fields.

Label	Description
Peer LDP ID	The IP address of the LDP peer.
Adj Type	The adjacency type between the LDP peer and LDP session is targeted. Link — Specifies that this adjacency is a result of a link hello. Targeted — Specifies that this adjacency is a result of a targeted hello.
State	Established — The adjacency is established. Trying — The adjacency is not yet established.
Mesg Sent	The number of messages sent.
Mesg Rcvd	The number of messages received.
Up Time	The amount of time the adjacency has been enabled.

Sample Output

```
*B:SRR# show router ldp session overload
=====
LDP Session Overload
=====
-----
Session with peer 10.8.100.15:0, Local 110.20.1.2:0
-----
Overload sent for      : 16155
Overload received for  : 0
```

```
-----
Session with peer 10.20.1.22:0, Local 110.20.1.2:0
-----
```

```
Overload sent for      : 0
Overload received for  : 16155
-----
```

```
*B:SRR# show router ldp session detail
```

```
=====
LDP Sessions (Detail)
=====
```

```
Legend: DoD - Downstream on Demand (for address FEC's only)
        DU  - Downstream Unsolicited
=====
```

```
-----
Session with Peer 10.8.100.15:0, Local 110.20.1.2:0
-----
```

```
Adjacency Type      : Targeted      State           : Established
Up Time             : 0d 00:48:26
Max PDU Length      : 4096          KA/HOLD Time Remaining : 131
Link Adjacencies    : 0            Targeted Adjacencies  : 1
Local Address       : 110.20.1.2    Peer Address       : 10.8.100.15
Local TCP Port      : 53718         Peer TCP Port      : 646
Local KA Timeout    : 140          Peer KA Timeout    : 180
Msg Sent            : 37531         Msg Recv           : 22526
FECs Sent           : 0            FECs Recv          : 1688
Addrs Sent          : 417           Addrs Recv         : 5
GR State            : Capable        Label Distribution   : DU
Nbr Liveness Time   : 5            Max Recovery Time   : 0
Number of Restart   : 0            Last Restart Time   : Never
P2MP                : Not Capable    MP MBB              : Not Capable
Dynamic Capability  : Not Capable    LSR Overload        : Not Capable
Advertise           : Service
Addr FEC OverLoad Sent : No          Addr FEC OverLoad Recv : No
Mcast FEC Overload Sent: No          Mcast FEC Overload Recv: No
Serv FEC Overload Sent : No          Serv FEC Overload Recv : No
-----
```

```
-----
Session with Peer 110.20.1.1:0, Local 110.20.1.2:0
-----
```

```
Adjacency Type      : Link          State           : Established
Up Time             : 0d 00:13:45
Max PDU Length      : 4096          KA/HOLD Time Remaining : 140
Link Adjacencies    : 2            Targeted Adjacencies  : 0
Local Address       : 110.20.1.2    Peer Address       : 110.20.1.1
Local TCP Port      : 53914         Peer TCP Port      : 646
Local KA Timeout    : 140          Peer KA Timeout    : 140
Msg Sent            : 8032          Msg Recv           : 6325
FECs Sent           : 1813         FECs Recv          : 1196
Addrs Sent          : 417           Addrs Recv         : 1054
GR State            : Capable        Label Distribution   : DU
Nbr Liveness Time   : 0            Max Recovery Time   : 0
Number of Restart   : 0            Last Restart Time   : Never
P2MP                : Capable        MP MBB              : Capable
Dynamic Capability  : Not Capable    LSR Overload        : Capable
Advertise           : Address
Addr FEC OverLoad Sent : No          Addr FEC OverLoad Recv : No
Mcast FEC Overload Sent: No          Mcast FEC Overload Recv: No
Serv FEC Overload Sent : No          Serv FEC Overload Recv : No
-----
```

Show Commands

{...snip...}

```
*A:Dut-C# show router ldp session local-addresses
=====
LDP Session Local-Addresses
=====
-----
Session with Peer 10.20.1.2:0, Local 10.20.1.3:0
-----
Sent Addresses: 1.1.1.1      10.10.12.3    10.10.22.3    10.20.1.3
                10.180.2.3    10.180.3.3    10.180.5.3    10.180.11.3
                10.181.2.3    10.181.3.3    10.181.5.3    10.181.11.3
                10.182.2.3    10.182.3.3    10.182.5.3    10.182.11.3

Recv Addresses: 2.2.2.2      10.10.12.2    10.20.1.2     10.180.1.2
                10.180.3.2    10.180.4.2    10.181.1.2    10.181.3.2
                10.181.4.2    10.182.1.2    10.182.3.2    10.182.4.2
-----
Session with Peer 10.20.1.4:0, Local 10.20.1.3:0
-----
Sent Addresses: 1.1.1.1      10.10.12.3    10.10.22.3    10.20.1.3
                10.180.2.3    10.180.3.3    10.180.5.3    10.180.11.3
                10.181.2.3    10.181.3.3    10.181.5.3    10.181.11.3
                10.182.2.3    10.182.3.3    10.182.5.3    10.182.11.3

Recv Addresses: 10.10.22.4   10.20.1.4     10.180.4.4    10.180.6.4
                10.180.9.4   10.180.11.4   10.181.4.4    10.181.6.4
                10.181.9.4   10.181.11.4   10.182.4.4    10.182.6.4
                10.182.9.4   10.182.11.4
-----
Session with Peer 10.20.1.5:0, Local 10.20.1.3:0
-----
Sent Addresses: 1.1.1.1      10.10.12.3    10.10.22.3    10.20.1.3
                10.180.2.3    10.180.3.3    10.180.5.3    10.180.11.3
                10.181.2.3    10.181.3.3    10.181.5.3    10.181.11.3
                10.182.2.3    10.182.3.3    10.182.5.3    10.182.11.3

Recv Addresses: 10.20.1.5   10.180.5.5    10.180.6.5    10.180.10.5
                10.181.5.5   10.181.6.5    10.181.10.5   10.182.5.5
                10.182.6.5   10.182.10.5
=====
*A:Dut-C#
```

```
*A:Dut-C# show router ldp session local-addresses
=====
LDP Session Local-Addresses
=====
-----
Session with Peer 10.20.1.2:0, Local 10.20.1.3:0
-----
Sent Addresses: 1.1.1.1      10.10.12.3    10.10.22.3    10.20.1.3
                10.180.2.3    10.180.3.3    10.180.5.3    10.180.11.3
                10.181.2.3    10.181.3.3    10.181.5.3    10.181.11.3
                10.182.2.3    10.182.3.3    10.182.5.3    10.182.11.3

Recv Addresses: 2.2.2.2      10.10.12.2    10.20.1.2     10.180.1.2
                10.180.3.2    10.180.4.2    10.181.1.2    10.181.3.2
                10.181.4.2    10.182.1.2    10.182.3.2    10.182.4.2
-----
```


Session with Peer 10.20.1.4:0, Local 10.20.1.3:0

```
-----
Sent Addresses: 1.1.1.1      10.10.12.3    10.10.22.3    10.20.1.3
                10.180.2.3    10.180.3.3    10.180.5.3    10.180.11.3
                10.181.2.3    10.181.3.3    10.181.5.3    10.181.11.3
                10.182.2.3    10.182.3.3    10.182.5.3    10.182.11.3

```

```
Recv Addresses: 10.10.22.4    10.20.1.4     10.180.4.4    10.180.6.4
                10.180.9.4    10.180.11.4   10.181.4.4    10.181.6.4
                10.181.9.4    10.181.11.4  10.182.4.4    10.182.6.4
                10.182.9.4    10.182.11.4

```

Session with Peer 10.20.1.5:0, Local 10.20.1.3:0

```
-----
Sent Addresses: 1.1.1.1      10.10.12.3    10.10.22.3    10.20.1.3
                10.180.2.3    10.180.3.3    10.180.5.3    10.180.11.3
                10.181.2.3    10.181.3.3    10.181.5.3    10.181.11.3
                10.182.2.3    10.182.3.3    10.182.5.3    10.182.11.3

```

```
Recv Addresses: 10.20.1.5     10.180.5.5    10.180.6.5    10.180.10.5
                10.181.5.5    10.181.6.5    10.181.10.5   10.182.5.5
                10.182.6.5    10.182.10.5

```

=====
*A:Dut-C#

*A:SRU4>config>router>ldp# show router ldp session

=====
LDP Sessions

```
-----
Peer LDP Id      Adj Type  State      Msg Sent  Msg Recv  Up Time
-----
1.1.1.1:0       Link     Nonexistent  2         1         0d 00:00:04
10.8.100.15:0   Both    Nonexistent 14653     21054     0d 12:48:25
10.20.1.20:0    Both    Established 105187    84837     0d 12:48:27
10.20.1.22:0    Both    Established 144586    95148     0d 12:48:23
11.22.10.2:0    Link     Nonexistent  4         2         0d 00:00:16
11.22.11.2:0    Link     Nonexistent  4         4         0d 00:00:14
11.22.13.2:0    Link     Nonexistent  5         6         0d 00:00:20
33.66.33.1:0    Link     Nonexistent  6         7         0d 00:00:25
33.66.34.1:0    Link     Nonexistent  2         2         0d 00:00:05
33.66.35.1:0    Link     Nonexistent  4         4         0d 00:00:14
110.20.1.1:0    Targeted Nonexistent  0         1         0d 00:00:04
110.20.1.3:0    Both    Established  94        97        0d 00:00:55
110.20.1.5:0    Both    Established 230866    286216    0d 12:48:27
110.20.1.110:0  Link     Nonexistent  2         2         0d 00:00:05
200.0.0.1:0     Link     Nonexistent  2         2         0d 00:00:05

```

No. of Sessions: 15

=====
*A:SRU4>config>router>ldp#

*A:SRU4>config>router>ldp# show router ldp session 10.20.1.20:0

=====
LDP Sessions

```
-----
Peer LDP Id      Adj Type  State      Msg Sent  Msg Recv  Up Time
-----
10.20.1.20:0    Both    Established 105204    84859     0d 12:49:05

```

Show Commands

```
No. of Sessions: 1
=====
*A:SRU4>config>router>ldp#

*A:SRU4# show router ldp session detail
=====
LDP Sessions (Detail)
=====
Legend: DoD - Downstream on Demand (for address FEC's only)
        DU - Downstream Unsolicited
=====
-----
Session with Peer 10.8.100.15:0
-----
Adjacency Type      : Both                State                : Nonexistent
Up Time             : 0d 00:50:30
Max PDU Length      : 4096                 KA/Hold Time Remaining : 0
Link Adjacencies    : 1                   Targeted Adjacencies  : 1
Local Address       : 110.20.1.4           Peer Address          : 10.8.100.15
Local TCP Port      : 0                   Peer TCP Port         : 0
Local KA Timeout    : 30                   Peer KA Timeout       : 30
Mesg Sent           : 951                   Mesg Recv             : 1388
FECs Sent           : 0                   FECs Recv             : 0
GR State            : Capable               Label Distribution    : DU
Nbr Liveness Time   : 0                   Max Recovery Time     : 0
Number of Restart   : 0                   Last Restart Time     : Never
P2MP                : Not Capable          MP MBB               : Not Capable
Dynamic Capability  : Not Capable
Advertise           : Address/Servi*
-----
Session with Peer 10.20.1.20:0
-----
*A:SRU4>config>router>ldp# show router ldp session detail
=====
LDP Sessions (Detail)
=====
Legend: DoD - Downstream on Demand (for address FEC's only)
        DU - Downstream Unsolicited
=====
-----
Session with Peer 1.1.1.1:0
-----
Adjacency Type      : Link                State                : Nonexistent
Up Time             : 0d 00:00:22
Max PDU Length      : 4096                 KA/Hold Time Remaining : 0
Link Adjacencies    : 1                   Targeted Adjacencies  : 0
Local Address       : 110.20.1.4           Peer Address          : 1.1.1.1
Local TCP Port      : 0                   Peer TCP Port         : 0
Local KA Timeout    : 30                   Peer KA Timeout       : 30
Mesg Sent           : 5                   Mesg Recv             : 2
FECs Sent           : 0                   FECs Recv             : 0
GR State            : Capable               Label Distribution    : DU
Nbr Liveness Time   : 0                   Max Recovery Time     : 0
Number of Restart   : 0                   Last Restart Time     : Never
Advertise           : Address
-----
Session with Peer 10.8.100.15:0
-----
Adjacency Type      : Both                State                : Nonexistent
Up Time             : 0d 12:49:26
```

```

Max PDU Length      : 4096           KA/Hold Time Remaining: 0
Link Adjacencies   : 1               Targeted Adjacencies  : 1
Local Address      : 110.20.1.4      Peer Address          : 10.8.100.15
Local TCP Port     : 0               Peer TCP Port         : 0
Local KA Timeout   : 30              Peer KA Timeout       : 30
Mesg Sent          : 14672           Mesg Recv             : 21081
FECs Sent         : 0               FECs Recv             : 0
GR State          : Capable          Label Distribution    : DU
Nbr Liveness Time : 0               Max Recovery Time     : 0
A:cpm-a#
...

```

```
-----
Session with Peer 200.0.0.1:0
-----
```

```

Adjacency Type      : Link           State                  : Nonexistent
Up Time            : 0d 00:00:02
Max PDU Length     : 4096           KA/Hold Time Remaining: 28
Link Adjacencies   : 1               Targeted Adjacencies  : 0
Local Address      : 110.20.1.4      Peer Address          : 200.0.0.1
Local TCP Port     : 0               Peer TCP Port         : 0
Local KA Timeout   : 30              Peer KA Timeout       : 30
Mesg Sent          : 1               Mesg Recv             : 1
FECs Sent         : 0               FECs Recv             : 0
GR State          : Capable          Label Distribution    : DU
Nbr Liveness Time : 0               Max Recovery Time     : 0
Number of Restart  : 0               Last Restart Time     : Never
Advertise          : Address

```

```
=====
* indicates that the corresponding row element may have been truncated.

```

```
*A:SRU4>config>router>ldp#
```

status

- Syntax** **status**
- Context** show>router>ldp
- Description** This command displays LDP status information.
- Output** **LDP Status Output** — The following table describes LDP status output fields.

Label	Description
Admin State	Up — The LDP is administratively enabled. Down — The LDP is administratively disabled.
Oper State	Up — The LDP is operationally enabled. Down — The LDP is operationally disabled.
Created at	The date and time when the LDP instance was created.
Up Time	The time, in hundredths of seconds, that the LDP instance has been operationally up.
Last Change	The date and time when the LDP instance was last modified.
Oper Down Events	The number of times the LDP instance has gone operationally down since the instance was created.
Active Adjacencies	The number of active adjacencies (established sessions) associated with the LDP instance.
Active Sessions	The number of active sessions (session in some form of creation) associated with the LDP instance.
Active Interfaces	The number of active (operationally up) interfaces associated with the LDP instance.
Inactive Interfaces	The number of inactive (operationally down) interfaces associated with the LDP instance.
Active Peers	The number of active LDP peers.
Inactive Peers	The number of inactive LDP peers.
Addr FECs Sent	The number of labels that have been sent to the peer associated with this FEC.
Addr FECs Recv	The number of labels that have been received from the peer associated with this FEC.
Serv FECs Sent	The number of labels sent to the peer associated with this FEC.
Serv FECs Recv	The number of labels received from the peer associated with this FEC.
Attempted Sessions	The total number of attempted sessions for this LDP instance.

Label	Description (Continued)
No Hello Err	The total number of “Session Rejected” or “No Hello Error” notification messages sent or received by this LDP instance.
Param Adv Err	The total number of “Session Rejected” or “Parameters Advertisement Mode Error” notification messages sent or received by this LDP instance.
Max PDU Err	The total number of “Session Rejected” or “Parameters Max PDU Length Error” notification messages sent or received by this LDP instance.
Label Range Err	The total number of “Session Rejected” or “Parameters Label Range Error” notification messages sent or received by this LDP instance.
Bad LDP Id Err	The number of bad LDP identifier fatal errors detected for sessions associated with this LDP instance.
Bad PDU Len Err	The number of bad PDU length fatal errors detected for sessions associated with this LDP instance.
Bad Mesg Len Err	The number of bad message length fatal errors detected for sessions associated with this LDP instance.
Bad TLV Len Err	The number of bad TLV length fatal errors detected for sessions associated with this LDP instance.
Malformed TLV Err	The number of malformed TLV value fatal errors detected for sessions associated with this LDP instance.
Shutdown Notif Sent	The number of shutdown notifications sent related to sessions associated with this LDP instance.
Keepalive Expired Err	The number of session Keepalive timer expired errors detected for sessions associated with this LDP instance.
Shutdown Notif Recv	The number of shutdown notifications received related to sessions associated with this LDP instance.

Sample Output

```
*B:SRR# show router ldp status
=====
LDP Status for LSR ID 110.20.1.2
=====
Admin State      : Up                Oper State      : Up
Created at      : 07/11/13 01:17:50  Up Time        : 0d 00:10:45
Oper Down Reason : n/a                Oper Down Events : 1
Last Change     : 07/11/13 01:23:46  Tunn Down Damp Time : 20 sec
Label Withdraw Del*: 0 sec          Implicit Null Label : Disabled
Short. TTL Prop Lo*: Disabled      Short. TTL Prop Tran*: Disabled
Import Policies :                   Export Policies  :
    Import-LDP                       Import-LDP
Tunl Exp Policies :
    from-proto-bgp
```

Show Commands

```
Aggregate Prefix      : True
FRR                  : Enabled
Dynamic Capability   : False
MP MBB Capability    : True
Overload Capability  : True
Active Adjacencies  : 26
Active Interfaces    : 38
Active Peers        : 11
Addr FECs Sent      : 16155
Serv FECs Sent      : 0
P2MP FECs Sent      : 971
Attempted Sessions  : 17
No Hello Err        : 0
Max PDU Err         : 0
Bad LDP Id Err      : 0
Bad Mesg Len Err    : 0
Unknown TLV Err     : 0
Malformed TLV Err   : 0
Shutdown Notif Sent: 0

Agg Prefix Policies  : None
Mcast Upstream FRR  : Enabled
P2MP Capability      : True
MP MBB Time          : 3

Active Sessions      : 11
Inactive Interfaces  : 0
Inactive Peers       : 1
Addr FECs Recv      : 13672
Serv FECs Recv       : 0
P2MP FECs Recv      : 2448

Param Adv Err        : 0
Label Range Err     : 0
Bad PDU Len Err     : 0
Bad TLV Len Err     : 1

Keepalive Expired Err: 0
Shutdown Notif Recv  : 1
=====
* indicates that the corresponding row element may have been truncated.
*B:SRR#
```

```
*A:SRU4>config>router>ldp# show router ldp status
=====
LDP Status for LSR ID 110.20.1.4
=====
Admin State          : Up
Created at           : 07/11/13 00:37:31
Oper Down Reason     : n/a
Last Change          : 07/11/13 01:11:14
Label Withdraw Del* : 0 sec
Short. TTL Prop Lo* : Enabled
Import Policies      :
  Import-LDP
Tunl Exp Policies    :
  from-proto-bgp
Aggregate Prefix     : True
FRR                  : Enabled
Dynamic Capability   : False
MP MBB Capability    : True
Overload Capability  : True
Active Adjacencies  : 38
Active Interfaces    : 36
Active Peers        : 19
Addr FECs Sent      : 36206
Serv FECs Sent      : 108
P2MP FECs Sent      : 2411
Attempted Sessions  : 166
No Hello Err        : 0
Max PDU Err         : 0
Bad LDP Id Err      : 2
Bad Mesg Len Err    : 0
Unknown TLV Err     : 0
Malformed TLV Err   : 0
Shutdown Notif Sent: 0

Oper State           : Up
Up Time              : 0d 00:44:09
Oper Down Events     : 1
Tunn Down Damp Time  : 20 sec
Implicit Null Label  : Enabled
Short. TTL Prop Tran*: Enabled
Export Policies      :
  Import-LDP

Agg Prefix Policies  : None
Mcast Upstream FRR  : Enabled
P2MP Capability      : True
MP MBB Time          : 10

Active Sessions      : 16
Inactive Interfaces  : 0
Inactive Peers       : 0
Addr FECs Recv      : 23413
Serv FECs Recv       : 108
P2MP FECs Recv      : 370

Param Adv Err        : 0
Label Range Err     : 0
Bad PDU Len Err     : 0
Bad TLV Len Err     : 19

Keepalive Expired Err: 21
Shutdown Notif Recv  : 13
=====
* indicates that the corresponding row element may have been truncated.
*A:SRU4>config>router>ldp#
```

```

*A:SRR# show router ldp status
=====
LDP Status for LSR ID 110.20.1.2
=====
Admin State      : Up                Oper State       : Up
Created at      : 07/26/11 05:49:52  Up Time         : 0d 15:20:48
Oper Down Reason : n/a                Oper Down Events : 12
Last Change     : 07/26/11 06:05:46  Tunn Down Damp Time : 20 sec
Label Withdraw Del*: 120 sec          Implicit Null Label : Disabled
Short. TTL Prop Lo*: Disabled        Short. TTL Prop Tran*: Disabled
Import Policies : None                Export Policies  : None
Tunl Exp Policies :
    from-proto-bgp
Aggregate Prefix : True                Agg Prefix Policies : None
FRR              : Enabled
Dynamic Capability : False           P2MP Capability    : True
MP MBB Capability : True              MP MBB Time       : 3
Active Adjacencies : 10              Active Sessions    : 6
Active Interfaces  : 34              Inactive Interfaces : 3
Active Peers       : 18              Inactive Peers     : 1
Addr FECs Sent     : 3066            Addr FECs Recv     : 3066
Serv FECs Sent     : 0                Serv FECs Recv     : 0
P2MP FECs Sent     : 600             P2MP FECs Recv    : 1200
Attempted Sessions : 8575
No Hello Err      : 24                Param Adv Err      : 0
Max PDU Err       : 0                  Label Range Err    : 0
Bad LDP Id Err    : 18020            Bad PDU Len Err   : 0
Bad Mesg Len Err  : 0                  Bad TLV Len Err   : 0
Unknown TLV Err   : 0
Malformed TLV Err : 0                  Keepalive Expired Err: 1751
Shutdown Notif Sent: 10              Shutdown Notif Recv : 0
=====
* indicates that the corresponding row element may have been truncated.
*A:SRR#

```

```

*B:Dut-B# show router ldp status
=====
LDP Status for LSR ID 10.20.1.2
=====
Admin State      : Up                Oper State       : Up
Created at      : 11/19/2010 23:45:01  Up Time         : 68d 01:00:07
Oper Down Reason : n/a                Oper Down Events : 0
Last Change     : 11/19/2010 23:45:01  Tunn Down Damp Time : 3 sec
Label Withdraw Del*: 0 sec          Implicit Null Label : Disabled
Short. TTL Prop Lo*: Enabled        Short. TTL Prop Tran*: Enabled
Import Policies : None                Export Policies  : None
Tunl Exp Policies : None
Aggregate Prefix : False
Agg Prefix Policies: None
Dynamic Capability : False           P2MP Capability    : True
MP MBB Capability : True              MP MBB Time       : 3
Active Adjacencies : 3                Active Sessions    : 1
Active Interfaces  : 3                Inactive Interfaces : 9
Active Peers       : 0                Inactive Peers     : 0
Addr FECs Sent     : 1                Addr FECs Recv     : 1
Serv FECs Sent     : 0                Serv FECs Recv     : 0
Attempted Sessions : 0
No Hello Err      : 0                  Param Adv Err      : 0
Max PDU Err       : 0                  Label Range Err    : 0

```

Show Commands

```
Bad LDP Id Err      : 0                Bad PDU Len Err      : 0
Bad Mesg Len Err    : 0                Bad TLV Len Err      : 0
Unknown TLV Err     : 0
Malformed TLV Err   : 0                Keepalive Expired Err: 0
Shutdown Notif Sent: 0                Shutdown Notif Recv  : 0
=====
* indicates that the corresponding row element may have been truncated.
*B:Dut-B#

*A:SRU4# show router ldp statistics-summary
=====
Statistics Summary
=====
LDP FEC Prefix egress statistics : 0
=====
*A:SRU4#
```

Clear Commands

fec-egress-statistics

Syntax	fec-egress-statistics [<i>ip-prefix/mask</i>]
Context	clear>router>ldp
Description	This command clears LDP FEC egress statistics.. <i>ip-prefix</i> — Specify information for the specified IP prefix and mask length. Host bits must be 0. <i>mask</i> — Specifies the 32-bit address mask used to indicate the bits of an IP address that are being used for the subnet address. Values 0 — 32

instance

Syntax	instance
Context	clear>router>ldp
Description	This command resets the LDP instance.

interface

Syntax	interface [<i>ip-int-name</i>]
Context	clear>router>ldp
Description	This command restarts or clears statistics for LDP interfaces.
Parameters	<i>ip-int-name</i> — The name of an existing interface. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

peer

Syntax	peer [<i>ip-address</i>] [statistics]
Context	clear>router>ldp
Description	This command restarts or clears statistics for LDP targeted peers.
Parameters	<i>ip-address</i> — The IP address of a targeted peer. statistics — Clears only the statistics for a targeted peer

session

- Syntax** **session** [*ip-addr*:*label-space*] [**statistics**]
- Context** clear>router>ldp
- Description** This command restarts or clears statistics for LDP sessions.
- Parameters** *label-space* — Specifies the label space identifier that the router is advertising on the interface.
- Values** 0 — 65535
- statistics** — Clears only the statistics for a session.

statistics

- Syntax** **statistics**
- Context** clear>router>ldp
- Description** This command clears LDP instance statistics.

Debug Commands

The following output shows debug LDP configurations discussed in this section.

```
A:ALA-12# debug router ldp peer 10.10.10.104
A:ALA-12>debug>router>ldp# show debug ldp
debug
  router "Base"
    ldp peer 10.10.10.104
      event
        bindings
        messages
      exit
      packet
        hello
        init
        keepalive
        label
      exit
    exit
  exit
exit
A:ALA-12>debug>router>ldp#
```

ldp

Syntax	[no] ldp
Context	debug>router
Description	Use this command to configure LDP debugging.

interface

Syntax	[no] interface <i>interface-name</i>
Context	debug>router>ldp
Description	Use this command for debugging an LDP interface.
Parameters	<i>interface-name</i> — The name of an existing interface.

peer

Syntax	[no] peer <i>ip-address</i>
Context	debug>router>ldp
Description	Use this command for debugging an LDP peer.

Show Commands

Parameters *ip-address* — The IP address of the LDP peer.

event

Syntax **[no] event**

Context debug>router>ldp>if
debug>router>ldp>peer

Description This command configures debugging for specific LDP events.

bindings

Syntax **[no] bindings**

Context debug>router>ldp>peer>event

Description This command displays debugging information about addresses and label bindings learned from LDP peers for LDP bindings.

The **no** form of the command disables the debugging output.

messages

Syntax **[no] messages**

Context debug>router>ldp>if>event
debug>router>ldp>peer>event

Description This command displays specific information (for example, message type, source, and destination) regarding LDP messages sent to and received from LDP peers.

The **no** form of the command disables debugging output for LDP messages.

packet

Syntax **packet [detail]**
no packet

Context debug>router>ldp>if
debug>router>ldp>peer

Description This command enables debugging for specific LDP packets.

The **no** form of the command disables the debugging output.

Parameters **detail** — Displays detailed information.

hello

Syntax	hello [detail] no hello
Context	debug>router>ldp>if>packet debug>router>ldp>peer>packet
Description	This command enables debugging for LDP hello packets. The no form of the command disables the debugging output.
Parameters	detail — Displays detailed information.

init

Syntax	init [detail] no init
Context	debug>router>ldp>peer>packet
Description	This command enables debugging for LDP Init packets. The no form of the command disables the debugging output.
Parameters	detail — Displays detailed information.

keepalive

Syntax	[no] keepalive
Context	debug>router>ldp>peer>packet
Description	This command enables debugging for LDP Keepalive packets. The no form of the command disables the debugging output.

label

Syntax	label [detail] no label
Context	debug>router>ldp>peer>packet
Description	This command enables debugging for LDP Label packets. The no form of the command disables the debugging output.
Parameters	detail — Displays detailed information.

Show Commands

Standards and Protocol Support

Ethernet Standards

IEEE 802.1ab-REV/D3 Station and Media Access Control Connectivity Discovery
IEEE 802.1d Bridging
IEEE 802.1p/Q VLAN Tagging
IEEE 802.1s Multiple Spanning Tree
IEEE 802.1w Rapid Spanning Tree Protocol
IEEE 802.1x Port Based Network Access Control
IEEE 802.1ad Provider Bridges
IEEE 802.1ah Provider Backbone Bridges
IEEE 802.1ag Service Layer OAM
IEEE 802.3ah Ethernet in the First Mile
IEEE 802.1ak Multiple MAC Registration Protocol
IEEE 802.3 10BaseT
IEEE 802.3ad Link Aggregation
IEEE 802.3ae 10Gbps Ethernet
IEEE 802.3ah Ethernet OAM
IEEE 802.3u 100BaseTX
IEEE 802.3x Flow Control
IEEE 802.3z 1000BaseSX/LX
ITU-T Y.1731 OAM functions and mechanisms for Ethernet based networks
ITU-T G.8031 Ethernet linear protection switching
ITU-T G.8032 Ethernet Ring Protection Switching (version 2)

OSPF

RFC 1765 OSPF Database Overflow
RFC 2328 OSPF Version 2
RFC 2370 Opaque LSA Support
RFC 2740 OSPF for IPv6 (OSPFv3)
draft-ietf-ospf-ospfv3-update-14.txt
RFC 3101 OSPF NSSA Option
RFC 3137 OSPF Stub Router Advertisement
RFC 3623 Graceful OSPF Restart – GR helper

RFC 3630 Traffic Engineering (TE) Extensions to OSPF Version 2
RFC 4203 OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS) - (support of Link Local/Remote Identifiers and SRLG sub-TLVs)
RFC 5185 OSPF Multi-Area Adjacency
RFC5243 OSPF Database Summary List Optimization

BGP

RFC 1397 BGP Default Route Advertisement
RFC 1772 Application of BGP in the Internet
RFC 1965 Confederations for BGP
RFC 1997 BGP Communities Attribute
RFC 2385 Protection of BGP Sessions via MD5
RFC 2439 BGP Route Flap Dampening
RFC 2558 Multiprotocol Extensions for BGP-4
RFC 2918 Route Refresh Capability for BGP-4
RFC 3107 Carrying Label Information in BGP-4
RFC 3392 Capabilities Advertisement with BGP4
RFC 4271 BGP-4 (previously RFC 1771)
RFC 4360 BGP Extended Communities Attribute
RFC 4364 BGP/MPLS IP Virtual Private Networks (VPNs)(previously RFC 2547bis BGP/MPLS VPNs)
RFC 4456 BGP Route Reflection: Alternative to Full-mesh IBGP (previously RFC 1966 & 2796)
RFC 4486 Subcodes for BGP Cease Notification Message
RFC 4577 OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs)
RFC 4659 BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN

RFC 4684 Constrained Route Distribution for Border Gateway Protocol/MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs)
RFC 4724 Graceful Restart Mechanism for BGP – GR helper
RFC 4760 Multi-protocol Extensions for BGP
RFC 4798 Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE)
RFC 4893 BGP Support for Four-octet AS Number Space
RFC 5004 Avoid BGP Best Path Transitions from One External to Another
RFC 5065 Confederations for BGP (obsoletes 3065)
RFC 5291 Outbound Route Filtering Capability for BGP-4
RFC 5575 Dissemination of Flow Specification Rules
RFC 5668 4-Octet AS Specific BGP Extended Community
draft-ietf-idr-add-paths
draft-ietf-idr-best-external

IS-IS

ISO/IEC 10589:2002, Second Edition Intermediate System to Intermediate System Intra-Domain Routeing Information Exchange Protocol
RFC 1195 Use of OSI IS-IS for Routing in TCP/IP and Dual Environments
RFC 2973 IS-IS Mesh Groups
RFC 3359 Reserved Type, Length and Value (TLV) Codepoints in Intermediate System to Intermediate System
RFC 3719 Recommendations for Interoperable Networks using Intermediate System to Intermediate System (IS-IS)
RFC 3787 Recommendations for Interoperable IP Networks using

Standards and Protocols

Intermediate System to Intermediate System (IS-IS)
RFC 5120 M-ISIS: Multi Topology (MT) Routing in IS-IS (Partial)
RFC 5301 Dynamic Hostname Exchange Mechanism for IS-IS
RFC 5302 Domain-wide Prefix Distribution with Two-Level IS-IS
RFC 5303 Three-Way Handshake for IS-IS Point-to-Point Adjacencies
RFC 5304 IS-IS Cryptographic Authentication
RFC 5305 IS-IS Extensions for Traffic Engineering TE
RFC 5306 Restart Signaling for IS-IS
RFC 5307 IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)
RFC 5309 Point-to-Point Operation over LAN in Link State Routing Protocols
RFC 6329 IS-IS Extensions Supporting IEEE 802.1aq Shortest Path Bridging (Partial)
draft-ietf-isis-mi IS-IS Multi-Instance

IPSec

RFC 2401 Security Architecture for the Internet Protocol
RFC 2406 IP Encapsulating Security Payload (ESP)
RFC 2409 The Internet Key Exchange (IKE)
RFC 2560 X.509 Internet Public Key Infrastructure
 Online Certificate Status Protocol - OCSP
RFC 3706 IKE Dead Peer Detection
RFC 3947 Negotiation of NAT-Traversal in the IKE
RFC 3948 UDP Encapsulation of IPsec ESP Packets
RFC 4210 Internet X.509 Public Key Infrastructure
 Certificate Management Protocol (CMP)
RFC 4211 Internet X.509 Public Key Infrastructure
 Certificate Request Message Format (CRMF)
RFC 5996 Internet Key Exchange Protocol Version 2 (IKEv2)

RFC 5998 An Extension for EAP-Only Authentication in IKEv2
draft-ietf-ipsec-isakmp-xauth-06.txt – Extended Authentication within ISAKMP/Oakley (XAUTH)
draft-ietf-ipsec-isakmp-modecfg-05.txt – The ISAKMP Configuration Method

IPv6

RFC 1981 Path MTU Discovery for IPv6
RFC 2375 IPv6 Multicast Address Assignments
RFC 2460 Internet Protocol, Version 6 (IPv6) Specification
RFC 2461 Neighbor Discovery for IPv6
RFC 2462 IPv6 Stateless Address Auto configuration
RFC 2463 Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 Specification
RFC 2464 Transmission of IPv6 Packets over Ethernet Networks
RFC 2529 Transmission of IPv6 over IPv4 Domains without Explicit Tunnels
RFC 2545 Use of BGP-4 Multiprotocol Extension for IPv6 Inter-Domain Routing
RFC 2710 Multicast Listener Discovery (MLD) for IPv6
RFC 2740 OSPF for IPv6
RFC 3306 Unicast-Prefix-based IPv6 Multicast Addresses
RFC 3315 Dynamic Host Configuration Protocol for IPv6
RFC 3587 IPv6 Global Unicast Address Format
RFC3590 Source Address Selection for the Multicast Listener Discovery (MLD) Protocol
RFC 3810 Multicast Listener Discovery Version 2 (MLDv2) for IPv6
RFC 4007 IPv6 Scoped Address Architecture
RFC 4193 Unique Local IPv6 Unicast Addresses
RFC 4291 IPv6 Addressing Architecture
RFC 4552 Authentication/Confidentiality for OSPFv3
RFC 4659 BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN

RFC 5072 IP Version 6 over PPP
RFC 5095 Deprecation of Type 0 Routing Headers in IPv6
RFC 5308 Routing IPv6 with IS-IS

Multicast

RFC 1112 Host Extensions for IP Multicasting (Snooping)
RFC 2236 Internet Group Management Protocol, (Snooping)
RFC 3376 Internet Group Management Protocol, Version 3 (Snooping)
RFC 2362 Protocol Independent Multicast-Sparse Mode (PIMSM)
RFC 3618 Multicast Source Discovery Protocol (MSDP)
RFC 3446 Anycast Rendezvous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP)
RFC 4601 Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)
RFC 4604 Using IGMPv3 and MLDv2 for Source-Specific Multicast
RFC 4607 Source-Specific Multicast for IP
RFC 4608 Source-Specific Protocol Independent Multicast in 232/8
RFC 4610 Anycast-RP Using Protocol Independent Multicast (PIM)
draft-ietf-pim-sm-bsr-06.txt
draft-rosen-vpn-mcast-15.txt Multicast in MPLS/BGP IP VPNs
draft-ietf-mboned-msdp-mib-01.txt
draft-ietf-l3vpn-2547bis-mcast-07: Multicast in MPLS/BGP IP VPNs
draft-ietf-l3vpn-2547bis-mcast-bgp-05: BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs
RFC 3956: Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address

MPLS

RFC 2430 A Provider Architecture DiffServ & TE
RFC 2474 Definition of the DS Field the IPv4 and IPv6 Headers (Rev)

RFC 2597 Assured Forwarding PHB Group (rev3260)
 RFC 2598 An Expedited Forwarding PHB
 RFC 3031 MPLS Architecture
 RFC 3032 MPLS Label Stack Encoding
 RFC 3443 Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks
 RFC 4182 Removing a Restriction on the use of MPLS Explicit NULL
 RFC 3140 Per-Hop Behavior Identification Codes
 RFC 5332 MPLS Multicast Encapsulations

MPLS — LDP

RFC 3037 LDP Applicability
 RFC 3478 Graceful Restart Mechanism for LDP – GR helper
 RFC 5036 LDP Specification
 RFC 5283 LDP extension for Inter-Area LSP
 RFC 5443 LDP IGP Synchronization
 RFC 6388 LDP Extensions for Point-to-Multipoint and Multipoint-to-Multipoint LSP
 RFC 6826 Multipoint LDP in-band signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths
 draft-pdutta-mpls-ldp-hello-reduce-04.txt, Targeted LDP Hello Reduction

MPLS/RSVP-TE

RFC 2702 Requirements for Traffic Engineering over MPLS
 RFC2747 RSVP Cryptographic Authentication
 RFC 2961 RSVP Refresh Overhead Reduction Extensions
 RFC3097 RSVP Cryptographic Authentication - Updated Message Type Value
 RFC 3209 Extensions to RSVP for Tunnels
 RFC 3473 Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Extensions –

(support of of IF_ID RSVP_HOP object with unnumbered interface and RSVP-TE Graceful Restart Helper Procedures)
 RFC 3477 Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)
 RFC 3564 Requirements for Diff-Serv-aware TE
 RFC 3906 Calculating Interior Gateway Protocol (IGP) Routes Over Traffic Engineering Tunnels
 RFC 4090 Fast reroute Extensions to RSVP-TE for LSP Tunnels
 RFC 4124 Protocol Extensions for Support of Diffserv-aware MPLS Traffic Engineering
 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
 draft-newton-mpls-te-dynamic-overbooking-00 A Diffserv-TE Implementation Model to dynamically change booking factors during failure events
 RFC 4561 Definition of a RRO Node-Id Sub-Object
 RFC 4875 Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)
 RFC 5151 Inter-domain MPLS and GMPLS Traffic Engineering – RSVP-TE Extensions
 RFC 5712 MPLS Traffic Engineering Soft Preemption
 RFC 5817 Graceful Shutdown in GMPLS Traffic Engineering Networks

MPLS — OAM

RFC 4379 Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures
 RFC 6425 Detecting Data Plane Failures in Point-to-Multipoint Multiprotocol

MPLS-TP

RFC 5586 MPLS Generic Associated Channel
 RFC 5921 A Framework for MPLS in Transport Networks
 RFC 5960 MPLS Transport Profile Data Plane Architecture
 RFC 6370 MPLS-TP Identifiers
 RFC 6378 MPLS-TP Linear Protection
 RFC 6428 Proactive Connectivity Verification, Continuity Check and Remote Defect indication for MPLS Transport Profile
 RFC 6426 MPLS On-Demand Connectivity and Route Tracing
 RFC 6478 Pseudowire Status for Static Pseudowires
 draft-ietf-mpls-tp-ethernet-addressing-02 MPLS-TP Next-Hop Ethernet Addressing

RIP

RFC 1058 RIP Version 1
 RFC 2082 RIP-2 MD5 Authentication
 RFC 2453 RIP Version 2

TCP/IP

RFC 768 UDP
 RFC 1350 The TFTP Protocol (Rev.
 RFC 791 IP
 RFC 792 ICMP
 RFC 793 TCP
 RFC 826 ARP
 RFC 854 Telnet
 RFC 951 BootP (rev)
 RFC 1519 CIDR
 RFC 1542 Clarifications and Extensions for the Bootstrap Protocol
 RFC 1812 Requirements for IPv4 Routers
 RFC 2347 TFTP option Extension
 RFC 2328 TFTP Blocksize Option
 RFC 2349 TFTP Timeout Interval and Transfer Size option
 RFC 2401 Security Architecture for Internet Protocol
 RFC 2428 FTP Extensions for IPv6 and NATs
 RFC 3596 DNS Extensions to Support IP version 6

Standards and Protocols

draft-ietf-bfd-mib-00.txt Bidirectional Forwarding Detection Management Information Base
RFC 5880 Bidirectional Forwarding Detection
RFC 5881 BFD IPv4 and IPv6 (Single Hop)
RFC 5883 BFD for Multihop Paths
RFC 5286 Basic Specification for IP Fast Reroute: Loop-Free Alternates

VRRP

RFC 2787 Definitions of Managed Objects for the Virtual Router Redundancy Protocol
RFC 3768 Virtual Router Redundancy Protocol
draft-ietf-vrrp-unified-spec-02: Virtual Router Redundancy Protocol Version 3 for IPv4 and IPv6

PPP

RFC 1332 PPP IPCP
RFC 1377 PPP OSINLCP
RFC 1638/2878 PPP BCP
RFC 1661 PPP (rev RFC2151)
RFC 1662 PPP in HDLC-like Framing
RFC 1877 PPP Internet Protocol Control Protocol Extensions for Name Server Addresses
RFC 1989 PPP Link Quality Monitoring
RFC 1990 The PPP Multilink Protocol (MP)
RFC 1994 "PPP Challenge Handshake Authentication Protocol (CHAP)
RFC 2516 A Method for Transmitting PPP Over Ethernet
RFC 2615 PPP over SONET/SDH
RFC 2686 The Multi-Class Extension to Multi-Link PPP

Frame Relay

FRF.1.2 - PVC User-to-Network Interface (UNI) Implementation Agreement
FRF.5 - Frame Relay/ATM PVC Network Interworking Implementation
ANSI T1.617 Annex D, DSS1 — Signalling Specification For Frame Relay Bearer Service.
FRF2.2 - PVC Network-to-Network Interface (NNI) Implementation Agreement.

FRF.12 Frame Relay Fragmentation Implementation Agreement
FRF.16.1 Multilink Frame Relay UNI/ NNI Implementation Agreement
ITU-T Q.933 Annex A- Additional procedures for Permanent Virtual Connection (PVC) status management

ATM

RFC 1626 Default IP MTU for use over ATM AAL5
RFC 2514 Definitions of Textual Conventions and OBJECT_IDENTITIES for ATM Management
RFC 2515 Definition of Managed Objects for ATM Management
RFC 2684 Multiprotocol Encapsulation over ATM Adaptation Layer 5
AF-TM-0121.000 Traffic Management Specification Version 4.1
ITU-T Recommendation I.610 - B-ISDN Operation and Maintenance Principles and Functions version 11/95
ITU-T Recommendation I.432.1 – BISDN user-network interface – Physical layer specification: General characteristics
GR-1248-CORE - Generic Requirements for Operations of ATM Network Elements (NEs). Issue 3
GR-1113-CORE - Bellcore, Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements, Issue 1
AF-ILMI-0065.000 Integrated Local Management Interface (ILMI) Version 4.0
AF-TM-0150.00 Addendum to Traffic Management v4.1 optional minimum desired cell rate indication for UBR
AF-PHY-0086.001, Inverse Multiplexing for ATM (IMA) Specification Version 1.1

DHCP

RFC 2131 Dynamic Host Configuration Protocol (REV)
RFC 3046 DHCP Relay Agent Information Option (Option 82)

RFC 1534 Interoperation between DHCP and BOOTP

Policy Management and Credit Control

3GPP TS 29.212 - Policy and Charging Control (PCC) over Gx/Sd Reference Point (Release 11) - Gx support as it applies to wireline environment (BNG)
RFC 3588 – Diameter Base Protocol
RFC 4006 – Diameter Credit Control Application

NAT

RFC 6333 Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion
RFC 6334 Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Option for Dual-Stack Lite
RFC 6888 Common Requirements For Carrier-Grade NATs (CGNs)
RFC 5508 NAT Behavioral Requirements for ICMP
RFC 5382 NAT Behavioral Requirements for TCP
RFC 6146 Stateful NAT64

VPLS

RFC 4762 Virtual Private LAN Services Using LDP
RFC 5501: Requirements for Multicast Support in Virtual Private LAN Services (previously draft-ietf-l2vpn-vpls-mcast-reqts-04)
RFC 6074: Provisioning, Auto-Discovery, and Signaling in Layer 2 Virtual Private Networks (L2VPNs) (previously draft-ietf-l2vpn-signaling-08)
draft-ietf-l2vpn-vpls-mcast-13.txt Multicast in VPLS
RFC 7041 Extensions to the Virtual Private LAN Service (VPLS) Provider Edge (PE) Model for Provider Backbone Bridging
RFC 7117 Multicast in Virtual Private LAN Service (VPLS)

Pseudowire

RFC 3985 Pseudo Wire Emulation Edge-to-Edge (PWE3)

RFC 4385 Pseudo Wire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN

RFC 3916 Requirements for Pseudo-Wire Emulation Edge-to-Edge (PWE3)

RFC 4717 Encapsulation Methods for Transport ATM over MPLS Networks (draft-ietf-pwe3-atm-encap-10.txt)

RFC 4816 PWE3 ATM Transparent Cell Transport Service (draft-ietf-pwe3-cell-transport-04.txt)

RFC 4448 Encapsulation Methods for Transport of Ethernet over MPLS Networks (draft-ietf-pwe3-ethernet-encap-11.txt)

RFC 4619 Encapsulation Methods for Transport of Frame Relay over MPLS Networks (draft-ietf-pwe3-frame-relay-07.txt)

RFC 4446 IANA Allocations for PWE3

RFC 4447 Pseudowire Setup and Maintenance Using LDP (draft-ietf-pwe3-control-protocol-17.txt)

RFC 5085, Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires

RFC 5659 An Architecture for Multi-Segment Pseudowire Emulation Edge-to-Edge
draft-ietf-l2vpn-vpws-iw-oam-03.txt, OAM Procedures for VPWS Interworking

draft-ietf-pwe3-mpls-eth-oam-iwk-07.txt, MPLS and Ethernet OAM Interworking RFC6073, Segmented Pseudowire

draft-ietf-pwe3-dynamic-ms-pw-16.txt , Dynamic Placement of Multi Segment Pseudo Wires

RFC 6310, Pseudowire (PW) OAM Message Mapping

RFC6391 Flow Aware Transport of Pseudowires over an MPLS PSN

RFC 6575, ARP Mediation for IP Interworking of Layer 2 VPN

RFC 6718, Pseudowire Redundancy

RFC 6870, Pseudowire Preferential Forwarding Status bit

MFA Forum 9.0.0 The Use of Virtual trunks for ATM/MPLS Control Plane Interworking

MFA Forum 12.0.0 Multiservice Interworking - Ethernet over MPLS

MFA forum 13.0.0 - Fault Management for Multiservice Interworking v1.0

MFA Forum 16.0.0 – Multiservice Interworking - IP over MPLS

ANCP/L2CP

RFC5851 ANCP framework
draft-ietf-ancp-protocol-02.txt ANCP Protocol

Voice /Video Performance:

ITU-T G.107 The E Model- A computational model for use in planning.

ETSI TS 101 329-5 Annex E extensions- QoS Measurement for VoIP - Method for determining an Equipment Impairment Factor using Passive Monitoring

ITU-T Rec. P.564 - Conformance testing for voice over IP transmission quality assessment models

ITU-T G.1020 - Appendix I- Performance Parameter Definitions for Quality of Speech and other Voiceband Applications Utilizing IP Networks- Mean Absolute Packet Delay Variation.& Markov Models.

RFC 3550 Appendix A.8- RTP: A Transport Protocol for Real-Time Applications- Estimating the Interarrival Jitter.

Circuit Emulation

RFC 4553 Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP)

RFC 5086 Structure-Aware Time Division Multiplexed (TDM) Circuit Emulation Service over Packet Switched Network (CESoPSN)

MEF-8 Implementation Agreement for the Emulation of PDH Circuits over Metro Ethernet Networks, October 2004

RFC 5287 Control Protocol Extensions for the Setup of Time-Division Multiplexing (TDM) Pseudowires in MPLS Networks

SONET/SDH

ITU-G.841 Telecommunication Standardization Section of ITU, Types and Characteristics of SDH Networks Protection Architecture, issued in October 1998 and as augmented by Corrigendum1 issued in July 2002

AAA

RFC 2865 Remote Authentication Dial In User Service

RFC 2866 RADIUS Accounting
draft-grant-tacacs-02.txt

SSH

RFC 4250 The Secure Shell (SSH) Protocol Protocol Assigned Numbers

RFC 4251 The Secure Shell (SSH) Protocol Architecture

RFC 4254 The Secure Shell (SSH) Connection Protocol

OpenFlow

ONF OpenFlow Switch Specification version 1.3.1 (Hybrid-switch/ FlowTable)

Timing

GR-253-CORE SONET Transport Systems: Common Generic Criteria. Issue 3, September 2000

ITU-T G.781 Telecommunication Standardization Section of ITU, Synchronization layer functions, issued 09/2008

ITU-T G.813 Telecommunication Standardization Section of ITU, Timing characteristics of SDH equipment slave clocks (SEC), issued 03/2003.

GR-1244-CORE Clocks for the Synchronized Network: Common Generic Criteria, Issue 3, May 2005

ITU-T G.8261 Telecommunication Standardization Section of ITU, Timing and synchronization aspects in packet networks, issued 04/2008.

ITU-T G.8262 Telecommunication Standardization Section of ITU, Timing characteristics of synchronous Ethernet equipment slave clock (EEC), issued 08/2007.

Standards and Protocols

ITU-T G.8264 Telecommunication Standardization Section of ITU, Distribution of timing information through packet networks, issued 10/2008.

ITU-T G.8265.1 Telecommunication Standardization Section of ITU, Precision time protocol telecom profile for frequency synchronization, issued 10/2010.

IEEE 1588-2008, IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems

Network Management

ITU-T X.721: Information technology-OSI-Structure of Management Information

ITU-T X.734: Information technology-OSI-Systems Management: Event Report Management Function

M.3100/3120 Equipment and Connection Models

TMF 509/613 Network Connectivity Model

RFC 1157 SNMPv1

RFC 1215 A Convention for Defining Traps for use with the SNMP

RFC 1657 BGP4-MIB

RFC 1724 RIPv2-MIB

RFC 1850 OSPF-MIB

RFC 1907 SNMPv2-MIB

RFC 2011 IP-MIB

RFC 2138 RADIUS

RFC 2206 RSVP-MIB

RFC 2452 IPv6 Management Information Base for the Transmission Control Protocol

RFC 2465 Management Information Base for IPv6: Textual Conventions and General Group

RFC 2558 SONET-MIB

RFC 2571 SNMP-FRAMEWORKMIB

RFC 2572 SNMP-MPD-MIB

RFC 2573 SNMP-TARGET-&-NOTIFICATION-MIB

RFC 2574 SNMP-USER-BASED-SMMIB

RFC 2575 SNMP-VIEW-BASED-ACM-MIB

RFC 2576 SNMP-COMMUNITY-MIB

RFC 2578 Structure of Management Information Version 2 (SMIv2)

RFC 2665 EtherLike-MIB

RFC 2819 RMON-MIB

RFC 2863 IF-MIB

RFC 2864 INVERTED-STACK-MIB

RFC 2987 VRRP-MIB

RFC 3014 NOTIFICATION-LOGMIB

RFC 3019 IP Version 6 Management Information Base for The Multicast Listener Discovery Protocol

RFC 3164 Syslog

RFC 3273 HCRMON-MIB

RFC 3411 An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks

RFC 3412 - Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)

RFC 3413 - Simple Network Management Protocol (SNMP) Applications

RFC 3414 - User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)

RFC 3418 - SNMP MIB

RFC 3826 The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model

RFC 4113 Management Information Base for the User Datagram Protocol (UDP)

RFC 4292 IP-FORWARD-MIB

RFC 4293 MIB for the Internet Protocol

RFC 5101 - Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information

RFC 6241: NETCONF Configuration Protocol

RFC 6242 Using the NETCONF Protocol over Secure Shell (SSH)

draft-ietf-ospf-mib-update-04.txt

draft-ietf-mpls-lsr-mib-06.txt

draft-ietf-mpls-te-mib-04.txt

draft-ietf-mpls-ldp-mib-07.txt

draft-ietf-isis-wg-mib-06 Management Information Base for Intermediate System to Intermediate System (IS-

IS)
IANA-IFType-MIB
IEEE8023-LAG-MIB

INDEX

L

LDP

- overview 430
 - architecture 431
 - execution flow 434
 - label exchange 436
 - LDP and MPLS 430
 - subsystems 432
 - label manager 433
 - LDP configuration 433
 - logger 433
 - memory manager 433
 - service manager 433
- configuring 506
 - basic 507
 - command reference 521
 - enabling 508
 - interface 513
 - LDP services signaling 515
 - management tasks 518
 - peer 514
 - targeted sessions 512

M

MPLS

- overview 21
 - label stack 22
 - label switched routers 25
 - label values 24
 - LSP types 25
 - reservation types 89
 - SDPs 171
 - services 171
 - signaling protocols 177
 - traffic engineering 102
- configuring
 - admin-group 180
 - basic 178
 - command reference 197, 200, 201
 - interface 181
 - LSPs 176, 183

- management tasks 190
- overview 176
- paths 176, 182
- router interface 177
- signaling protocol 177
- static LSP 184

R

RSVP

- overview 85
 - for MPLS 87
 - hellos 87
 - message pacing 89
 - traffic engineering 87
- configuring
 - command reference 205
 - message pacing 188
 - management tasks 195
 - RSVP parameters 187

S

- shared risk link groups 140

T

- T-LDP 449

