



7210 SAS M, X OS MPLS Guide

Software Version: 7210 SAS OS 5.0 Rev. 05

April 2013

Document Part Number: 93-0427-01-05



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 2011 Alcatel-Lucent. All rights reserved.

TABLE OF CONTENTS

Preface

Alcatel-Lucent 7210 SAS M Router Configuration Process	15
--	----

MPLS and RSVP

MPLS	18
MPLS Label Stack	19
Label Values	21
Label Switching Routers	22
LSP Types	22
MPLS Fast Re-Route (FRR)	24
Manual Bypass LSP	24
PLR Bypass LSP Selection Rules	25
FRR Node-Protection (Facility)	27
Uniform FRR Failover Time	27
RSVP	29
Using RSVP for MPLS	31
RSVP Traffic Engineering Extensions for MPLS	31
Hello Protocol	31
MD5 Authentication of RSVP Interface	32
Reservation Styles	33
RSVP Message Pacing	33
RSVP Overhead Refresh Reduction	34
Configuring Implicit Null	35
Traffic Engineering	36
TE Metric (IS-IS and OSPF)	37
Maintenance of TE links and Nodes	37
Advanced MPLS/RSVP Features	38
TE Graceful Shutdown	39
MPLS Service Usage	40
MPLS/RSVP Configuration Process Overview	41
Configuration Notes	42
Configuring MPLS and RSVP with CLI	43
MPLS Configuration Overview	44
LSPs	44
Paths	44
Router Interface	45
Choosing the Signaling Protocol	45
Basic MPLS Configuration	46
Common Configuration Tasks	47
Configuring MPLS Components	48
Configuring Global MPLS Parameters	48
Configuring an MPLS Interface	49
Configuring MPLS Paths	50
Configuring an MPLS LSP	51
Configuring a Static LSP	52
Configuring Manual Bypass Tunnels	53
Configuring RSVP Parameters	55

Table of Contents

Configuring RSVP Message Pacing Parameters	56
Configuring Graceful Shutdown	57
MPLS Configuration Management Tasks	58
Modifying MPLS Parameters	58
Modifying an MPLS LSP	59
Modifying MPLS Path Parameters	60
Modifying MPLS Static LSP Parameters	61
Deleting an MPLS Interface	62
RSVP Configuration Management Tasks	63
Modifying RSVP Parameters	63
Modifying RSVP Message Pacing Parameters	64
Deleting an Interface from RSVP	64
MPLS/RSVP Command Reference	65
Command Hierarchies	65
MPLS Commands	65
LSP Commands	66
MPLS Path Commands	68
RSVP Commands	69
Show Commands	70
Tools Commands	71
Clear Commands	71
Debug Commands	72
Label Distribution Protocol	
Label Distribution Protocol	154
LDP and MPLS	154
LDP Architecture	155
Subsystem Interrelationships	156
Memory Manager and LDP	157
Label Manager	157
LDP Configuration	157
Logger	157
Service Manager	157
Execution Flow	158
Initialization	158
Session Lifetime	158
Session Establishment	158
Label Exchange	159
Other Reasons for Label Actions	159
Cleanup	159
LDP Filters	160
LDP over RSVP Tunnels	161
Signaling and Operation	164
Rerouting Around Failures	166
Configuring Implicit Null Label	167
Multi-Area and Multi-Instance Extensions to LDP	167
LDP Process Overview	168
Configuring LDP with CLI	171
LDP Configuration Overview	172
Basic LDP Configuration	173

Common Configuration Tasks	174
Enabling LDP	174
Configuring Graceful-Restart Helper Parameters	175
Applying Export and Import Policies	176
Targeted Session Parameters	178
Interface Parameters	180
Peer Parameters	181
LDP Signaling and Services	182
LDP Configuration Management Tasks	183
Disabling LDP	183
Modifying Targeted Session Parameters	184
Modifying Interface Parameters	185
LDP Command Reference	187
Command Hierarchies	187
LDP Commands	187
Show Commands	188
Clear Commands	188
Debug Commands	189
Standards and Protocol Support	223
Index	227

Table of Contents

LIST OF TABLES

Preface

Table 1: Configuration Process15

MPLS and RSVP

Table 2: Packet/Label Field Description19

Label Distribution Protocol

LIST OF FIGURES

MPLS and RSVP

Figure 1:	Label Placement	19
Figure 2:	Label Packet Placement	20
Figure 3:	Bypass Tunnel Nodes	25
Figure 4:	FRR Node-Protection Example	27
Figure 5:	Establishing LSPs	29
Figure 6:	LSP Using RSVP Path Set Up	30
Figure 7:	MPLS and RSVP Configuration and Implementation Flow	41

Label Distribution Protocol

Figure 8:	Subsystem Interrelationships	156
Figure 9:	LDP over RSVP Application	161
Figure 10:	LDP over RSVP Application Variant	163
Figure 11:	LDP Configuration and Implementation	169

Preface

About This Guide

Note: This guide is not applicable for 7210 SAS-M devices configured in Access uplink mode.

This guide describes the services and protocol support provided by the 7210 SAS Series 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 the 7210 SAS M-Series 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

The 7210 SAS M, X OS documentation set is composed of the following books:

- 7210 SAS M, X OS Basic System Configuration Guide
This guide describes basic system configurations and operations.
- 7210 SAS M, X OS System Management Guide
This guide describes system security and access configurations as well as event logging and accounting logs.
- 7210 SAS M, X OS Interface Configuration Guide
This guide describes card, Media Dependent Adapter (MDA), and port provisioning.
- 7210 SAS M, X OS OS Router Configuration Guide
This guide describes logical IP routing interfaces and associated attributes such as an IP address, port, link aggregation group (LAG) as well as IP and MAC-based filtering.
- 7210 SAS M, X OS OS MPLS Guide
This guide describes how to configure Multiprotocol Label Switching (MPLS), Resource Reservation Protocol (RSVP), and Label Distribution Protocol (LDP).
- 7210 SAS M OS Services Guide
This guide describes how to configure service parameters such as customer information and user services.
- 7210 SAS M, X OS OAM and Diagnostic Guide
This guide describes how to configure features such as service mirroring and Operations, Administration and Management (OAM) tools.
- 7210 SAS M Quality of Service Guide
This guide describes how to configure Quality of Service (QoS) policy management.

Technical Support

If you purchased a service agreement for your 7210 SAS M-series 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 Alcatel-Lucent service agreement, contact your welcome center:

Web: http://www1.alcatel-lucent.com/comps/pages/carrier_support.jhtml

GETTING STARTED

In This Chapter

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

Alcatel-Lucent 7210 SAS M 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 18 RSVP on page 29 Label Distribution Protocol on page 153
Reference	List of IEEE, IETF, and other proprietary entities.	Standards and Protocol Support on page 223

MPLS and RSVP

In This Chapter

This chapter provides information to configure MPLS and RSVP.

- [MPLS on page 18](#)
 - [MPLS Label Stack on page 19](#)
 - [Label Switching Routers on page 22](#)
- [RSVP on page 29](#)
 - [Using RSVP for MPLS on page 31](#)
 - [Reservation Styles on page 33](#)
- [Traffic Engineering on page 36](#)
- [Advanced MPLS/RSVP Features on page 38](#)
 - [TE Graceful Shutdown on page 39](#)
- [MPLS Service Usage on page 40](#)
- [MPLS/RSVP Configuration Process Overview on page 41](#)
- [Configuration Notes on page 42](#)

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) 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).




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 22](#)) 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 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 RSVP, so a value is reserved.
- Values 4-15 are reserved for future use.

7210 SAS M, X OS 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 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 32768 through 131,071 are dynamically assigned for both MPLS and services.
- Label values 131,072 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 7210 SAS M supports only RSVP-TE for setting up LSPs. 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 29](#).

- 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 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

The 7210 SAS supports Manual bypass tunnels, on implementation of the Manual bypass feature a LSP can be pre-configured from a PLR which is used exclusively for bypass protection. If a path message for a new LSP requests for bypass protection, the node checks if a manual bypass tunnel satisfying the path constraints exists. If a tunnel is found, it is selected. If no such tunnel exists by default, the 7210 SAS dynamically signals a bypass LSP.

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 at the PLR node. 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.

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 LSP's 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.

Configuration Guidelines

Implicit NULL must be enabled for use of Manual Bypass or Dynamic Bypass (FRR facility) if the 7210 is used as a egress LER and/or is a Merge Point.

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 5 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 5: Establishing LSPs

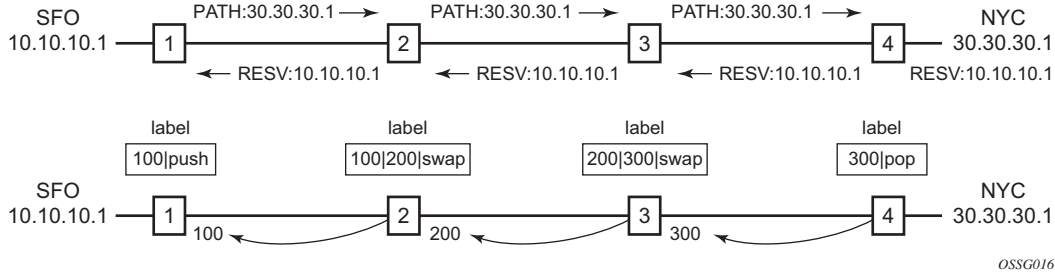


Figure 6: LSP Using RSVP Path Set Up

Figure 6 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 31](#)
 - [MD5 Authentication of RSVP Interface on page 32](#)
-

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 7210 SAS 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:

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 7210 SAS M 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 7210 SAS 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 7210 SAS 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 7210 SAS node does not attempt to send summary refresh messages.

Configuring Implicit Null

The implicit null label option allows a 7210 SAS 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 FEC signaling by the LDP control protocol.

The router can be configured to signal the implicit null label value over all RSVP interfaces and for all RSVP LSPs which have this node as the egress LER. In addition, the egress LER can be configured to receive MPLS packet with the implicit null label on a static LSP.

The following CLI command is used to configure the router:

```
config>router>ldp>implicit-null-label
```

Note: RSVP must be shutdown before changing the implicit null configuration option.

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.

Maintenance of TE links and Nodes

Graceful shutdown is used to maintain selective links and nodes in a TE network. Prior to a shutdown, headend LER nodes are notified of the imminent shutdown of the links or nodes for the purpose of maintenance, so the head-end nodes can move the paths of the LSPs away from the affected resources. Modified TE parameters for the affected links are flooded so all other nodes in the network avoid them in the CSPF calculations.

When the maintenance is over, the operator disables graceful shutdown, which reinstates and floods the user-configured TE parameters. The restored links are available for LSP path establishment.

Advanced MPLS/RSVP Features

- [TE Graceful Shutdown on page 39](#)

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 draft 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.

MPLS Service Usage

The 7210 SAS M, X routers enable service providers to deliver virtual private networks (VPNs) and Internet access using MPLS tunnels, with Ethernet interfaces.

MPLS/RSVP Configuration Process Overview

Figure 7 displays the process to configure MPLS and RSVP parameters.

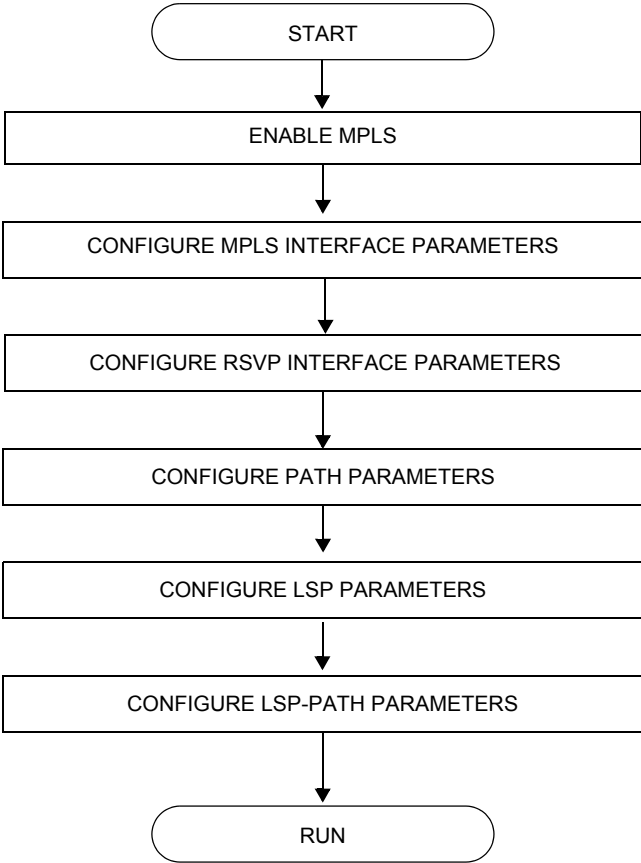


Figure 7: MPLS and RSVP Configuration and Implementation Flow

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 44](#)
 - [LSPs on page 44](#)
 - [Paths on page 44](#)
 - [Router Interface on page 45](#)
 - [Choosing the Signaling Protocol on page 45](#)
- [Basic MPLS Configuration on page 46](#)
- [Common Configuration Tasks on page 47](#)
 - [Configuring MPLS Components on page 48](#)
 - [Configuring Global MPLS Parameters on page 48](#)
 - [Configuring an MPLS Interface on page 49](#)
 - [Configuring MPLS Paths on page 50](#)
 - [Configuring an MPLS LSP on page 51](#)
- [Configuring RSVP Parameters on page 55](#)
 - [Configuring RSVP Message Pacing Parameters on page 56](#)
- [MPLS Configuration Management Tasks on page 58](#)
- [RSVP Configuration Management Tasks on page 63](#)

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 44](#)
 - [Paths on page 44](#)
 - [Router Interface on page 45](#)
-

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 7210 SAS M, X OS automatically manages label values. Labels that are automatically assigned have values ranging from 1,024 through 1,048,575 (see [Label Values on page 21](#)).

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

If only static label switched paths are used in your configurations, then you must manually define the paths through the MPLS network. Label mappings and actions configured at each hop must be specified. You do not need to enable RSVP if you are configuring static LSPs.

If dynamic LSP signaling is implemented in your network, then RSVP must be specified. Enable signaling protocols only on the links where the functionality is required.

In order to implement MPLS, the following entities must be enabled:

- MPLS must be enabled on all routers that are part of an LSP.
- RSVP must be enabled on the same routers.

When MPLS is enabled and either RSVP is also enabled, MPLS uses RSVP to set up the configured LSPs. For example, when you configure an LSP with both MPLS and RSVP running, RSVP initiates a session for the LSP. RSVP uses the local router as the RSVP session sender and the LSP destination as the RSVP session receiver. When the RSVP session is created, the LSP is set up on the path created by the session. If the session is not successfully created, RSVP notifies MPLS; MPLS can then either initiate backup paths or retry the initial path.

Basic MPLS Configuration

This section provides information to configure MPLS and configuration examples of common configuration tasks. To enable MPLS on 7210 SAS M, X Series routers, you must configure at least one MPLS interface. The other MPLS configuration parameters are optional. This follows 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)
- LDP

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 48](#)
 - [Configuring an MPLS Interface on page 49](#)
 - [Configuring MPLS Paths on page 50](#)
 - [Configuring an MPLS LSP on page 51](#)
 - [Configuring a Static LSP on page 52](#)
 - [Configuring RSVP Parameters on page 55](#)
 - [Configuring RSVP Message Pacing Parameters on page 56](#)
-

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
                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.

Listed below are the steps to configure the manual bypass tunnels:

1. Configure the option to disable the dynamic bypass tunnels on the 7210 SAS node B (if required). The CLI for this configuration is:
config>router>mpls>dynamic-bypass [disable | enable]
 The dynamic bypass tunnels are enabled by default.
2. Configure an LSP on node B, such as B-E-F-C which is 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:

Common Configuration Tasks

```
A:7210 SAS>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:7210 SAS >config>router>mpls>path#
```

3. Configure an LSP from A to D and indicate fast-reroute bypass protection, select the facility as "FRR method". (Config>router>mpls>lsp>fast-reroute facility).

Observe if the following criterions apply:

- If the LSP passes through B
- A bypass is requested
- The next hop is C
- A manually configured bypass-only tunnel exists from B to C (excluding link B to C)

Result: Node B uses the manually configured bypass-only tunnel from B to C.

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

Enable TE graceful shutdown on the maintenance interface using the **config>router>rsvp>interface>graceful-shutdown** command.

Disable graceful shutdown by executing the **no** form of the command at the RSVP interface level or at the RSVP level. This restores the user-configured TE parameters of the maintenance links, and the 7210 SAS maintenance node floods them.

MPLS Configuration Management Tasks

This section discusses the following MPLS configuration management tasks:

- [Modifying MPLS Parameters on page 58](#)
 - [Modifying MPLS Path Parameters on page 60](#)
 - [Modifying MPLS Static LSP Parameters on page 61](#)
 - [Deleting an MPLS Interface on page 62](#)
-

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 51](#).

```
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 50](#).

```
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 52](#).

```
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 63](#)
 - [Modifying RSVP Message Pacing Parameters on page 64](#)
 - [Deleting an Interface from RSVP on page 64](#)
-

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 55](#).

```
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 and then 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 62](#) for information on deleting an MPLS interface.

MPLS/RSVP Command Reference

Command Hierarchies

- [MPLS Commands on page 65](#)
- [MPLS Path Commands on page 68](#)
- [LSP Commands on page 66](#)
- [RSVP Commands on page 69](#)
- [Show Commands on page 70](#)
- [Clear Commands on page 71](#)
- [Debug Commands on page 72](#)

MPLS Commands

```

config
  — router
    — [no] mpls
      — admin-group group-name group-value
      — no admin-group group-name
      — dynamic-bypass [enable | disable]
      — [no] frr-object
      — hold-timer seconds
      — no hold-timer
      — [no] interface ip-int-name
        — [no] admin-group group-name [group-name...(up to 5 max)]
        — te-metric metric
        — no te-metric
      — resignal-timer minutes
      — no resignal-timer
      — [no] shutdown
      — [no] static-lsp lsp-name
        — no push label
        — push label nexthop ip-address
        — [no] shutdown
        — to ip-address
      — [no] static-lsp-fast-retry seconds

```

LSP Commands

```

config
  — router
    — [no] mpls
      — [no] lsp lsp-name [bypass-only | p2mp-lsp]
        — [no] adaptive
        — [no] adspec
        — bgp-transport-tunnel include | exclude
        — [no] cspf [use-te-metric]
        — [no] exclude group-name [group-name...(up to 5 max)]
        — fast-reroute frr-method
        — no fast-reroute
          — bandwidth rate-in-mbps
          — no bandwidth
          — hop-limit number
          — no hop-limit
          — [no] node-protect
        — from ip-address
        — hop-limit number
        — no hop-limit
        — [no] include group-name [group-name...(up to 5 max)]
        — [no] ldp-over-rsvp [include | exclude]
        — metric metric
        — [no] primary path-name
          — [no] adaptive
          — bandwidth rate-in-mbps
          — no bandwidth
          — [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] record
          — [no] record-label
          — [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
          — [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
          — [no] record
          — [no] record-label
          — [no] shutdown
          — [no] standby

```

- [no] **shutdown**
- **to** *ip-address*
- **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
      — [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 (for 7210 SAS M in Network Mode)
        — [no] graceful-shutdown
        — hello-interval milli-seconds
        — no hello-interval
        — [no] refresh-reduction
          — [no] reliable-delivery
        — [no] shutdown
        — subscription percentage
        — no subscription
      — keep-multiplier number
      — no keep-multiplier
      — [no] msg-pacing
        — max-burst number
        — no max-burst
        — period milli-seconds
        — no period
      — rapid-retransmit-time hundred-milliseconds
      — no rapid-retransmit-time
      — rapid-retry-limit number
      — no rapid-retry-limit
      — refresh-time seconds
      — no refresh-time
      — [no] shutdown

```

Show Commands

```

show
  — router
    — mpls
      — admin-group group-name
      — bypass-tunnel [to ip-address] [protected-lsp name] [dynamic | manual] [detail]
      — interface [ip-int-name|ip-address] [label-map label]
      — interface [ip-int-name|ip-address]
      — label start-label [end-label | in-use | label-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]
      — path [path-name] [lsp-binding]
      — 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
      — p2mp-lsp using-template [lsp-template template-name] [detail]
      — static-lsp [lsp-name]
      — static-lsp {transit | terminate}
      — static-lsp count
      — status

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

- **perform**
 - **router**
 - **mpls**
 - **csf** **to** *ip-addr* [**from** *ip-addr*] [**bandwidth** *bandwidth*] [**include-bitmap** *bitmap*] [**exclude-bitmap** *bitmap*] [**hop-limit** *limit*] [**exclude-address** *excl-addr* [*excl-addr*...(upto 8 max)]] [**use-te-metric**] [**skip-interface** *interface-name*]
 - **resignal** {**lsp** *lsp-name* **path** *path-name* | **delay** *minutes*}
-

Clear Commands

- clear**
 - **router**
 - **mpls**
 - **interface** [*ip-int-name*]
 - **lsp** *lsp-name*
 - **rsvp**
 - **interface** [*ip-int-name*] [*statistics*]
 - **statistics**

Debug Commands

```

debug
  — router
    — mpls [lsp lsp-name] [sender source-address] [endpoint endpoint-address] [tunnel-id tunnel-id] [lsp-id lsp-id]
    — 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
          — auth
          — no auth
          — misc [detail]
          — no misc
          — nbr [detail]
          — no nbr
          — path [detail]
          — no path
          — resv [detail]
          — no resv
          — rr
          — no rr
        — [no] packet
          — all [detail]
          — no all
          — ack
          — bundle [detail]
          — no bundle
          — 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**
- **srefresh** [detail]
- **no srefresh**

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 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>

admin-group

Syntax	admin-group <i>group-name</i> <i>group-value</i> no admin-group <i>group-name</i>
Context	config>router>mpls
Description	<p>This command is used to define administrative groups or link coloring for an interface. The admin group names can signify link colors, such as red, yellow, or green. MPLS interfaces advertise the link colors the support. CSPF uses the information when paths are computed for constraint-based LSPs. CSPF must be enabled in order for admin groups to be relevant.</p> <p>Network resources (links) based on zones, geographic location, link location, etc., can be classified using admin groups. MPLS interfaces must be explicitly assigned to an admin group.</p> <p>Admin groups must be defined in the config>router>mpls context before they can be assigned to an MPLS interface. The IGP communicates the information throughout the area.</p> <p>Up to 32 group names can be defined in the config>router>mpls context. The admin-group names must be identical across all routers in a single domain.</p> <p>The no form of this command deletes the administrative group. All configuration information associated with this LSP is lost.</p>
Default	none
Parameters	<p><i>group-name</i> — Specify the name of the administrative group within a virtual router instance.</p> <p><i>group-value</i> — Specify the group value associated with this administrative group. This value is unique within a virtual router instance.</p>
Values	0 — 31

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. Note :Implicit NULL must be enabled for use of Manual Bypass or Dynamic Bypass (FRR facility) if the 7210 is used as a egress LER and/or is a Merge Point.
Default	enable Note :Implicit NULL must be enabled for use of Manual Bypass or Dynamic Bypass (FRR facility) if the 7210 is used as a egress LER and/or is a Merge Point.

frr-object

Syntax	[no] frr-object
Context	config>router>mpls
Description	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.
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. The no form of the command disables the hold-timer.
Default	1 second
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

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

static-lsp

Syntax	[no] static-lsp lsp-name
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 setup 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 lsp-name command generates 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.

push

Syntax	no push <i>label</i> push <i>label</i> nexthop <i>ip-address</i>
Context	config>router>mpls>static-lsp
Description	<p>This command specifies the label to be pushed on the label stack and the next hop IP address for the static LSP.</p> <p>The no form of this command removes the association of the label to push for the static LSP.</p>

Parameters *label* — 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 7750 SR 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 *ip-address* — 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 staticLSP.

Default shutdown

to

Syntax **to** *ip-address*

Context config>router>mpls>static-lsp

Description This command specifies the system IP address of the egress router for the static LSP. This command is required while creating an LSP. For LSPs that are used as transport tunnels for services, the **to** IP address *must* 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 *ip-address* — The system IP address of the egress router.

Default none

static-lsp-fast-retry

Syntax	static-lsp-fast-retry seconds [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

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)]
Context	config>router>mpls>interface
Description	<p>This command defines admin groups that this interface supports.</p> <p>This information is advertised as part of OSPF and IS-IS to help CSPF compute constrained LSPs that must include or exclude certain admin groups. An MPLS interface is assumed to belong to all the admin groups unless the 'admin-group' command is issued under the interface config. Once an 'admin-group' command is issued the interface is assumed to belong to only the specifically listed groups for that command.</p> <p>Each single operation of the admin-group command allows a maximum of 5 groups to be specified at a time. However, a maximum of 32 groups can be specified per interface through multiple operations.</p>
Default	no admin-group
Parameters	<p><i>group-name</i> — Name of the group. The group names should be the same across all routers in the MPLS domain.</p>

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

LSP Commands

lsp

Syntax	[no] lsp <i>lsp-name</i> [bypass-only p2mp-lsp]
Context	config>router>mpls
Description	<p>This command creates an LSP that is signaled dynamically by the 7210 SAS M.</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. All other statements under the LSP hierarchy are optional. Note that the maximum number of static configurable LSPs is 4.</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 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 7210 selects it. If no manual bypass tunnel is found, the 7210 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> and <i>to</i>.</p>

adaptive

Syntax	[no] adaptive
Context	config>router>mpls>lsp
Description	<p>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.</p>
Default	adaptive

adspec

Syntax	[no] adspec
Context	config>router>mpls>lsp
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.
Parameters	

bgp-transport-tunnel

Syntax	bgp-transport-tunnel <i>include</i> <i>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>

cspf

Syntax	[no] cspf [<i>use-te-metric</i>]
Context	config>router>mpls>lsp
Description	<p>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.</p> <p>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.</p>

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	<i>use-te-metric</i> — Specifies to use the use of the TE metric for the purpose of the LSP path computation by CSPF.

exclude

Syntax	[no] exclude <i>group-name</i> [<i>group-name</i> ...(up to 5 max)]
Context	config>router>mpls>lsp
Description	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. Use the no form of the command to remove the exclude command.
Default	no exclude
Parameters	<i>group-name</i> — Specify the existing group-name to be excluded when an LSP is set up.

fast-reroute

Syntax	fast-reroute <i>frr-method</i> no fast-reroute
Context	config>router>mpls>lsp
Description	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. When fast-reroute is enabled, each node along the path of the LSP tries to establish a detour LSP as follows: <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. 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. • 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. 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.

If an LSP is configured with **fast-reroute** *frt-method* specified but does not enable CSPF, then neither global revertive nor local revertive will be available for the LSP to recover.

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.

The **no** form of **fast-reroute hop-limit** command reverts to the default value.

Default **no fast-reroute** — When fast-reroute is specified, the default fast-reroute method is one-to-one.

Parameters **Values** **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 established.

bandwidth

Syntax **bandwidth** *rate-in-mbps*
no bandwidth

Context config>router>mpls>lsp>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 *rate-in-mbps* — Specifies the amount of bandwidth in Mbps to be reserved for the LSP path.

hop-limit

Syntax **hop-limit** *limit*
no hop-limit

Context config>router>mpls>lsp>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 *limit* — Specify the maximum number of hops.

Values 0 — 255

node-protect

Syntax	[no] node-protect
Context	config>router>mpls>lsp>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
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
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>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</p>

MPLS Commands

the new value is equal to or greater than the current number hops of the established LSP, then the LSP is not affected.

The **no** form of this command returns the parameter to the default value.

Default 255

Parameters *number* — The number of hops the LSP can traverse, expressed as an integer.

Values 2 — 255

include

Syntax [**no**] **include** *group-name* [*group-name...*(up to 5max)]

Context config>router>mpls>lsp
config>router>mpls>lsp>primary
config>router>mpls>lsp>secondary

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 *group-name* — Specifies admin groups to be included when an LSP is set up.

ldp-over-rsvp

Syntax [**no**] **ldp-over-rsvp** [*include* | *exclude*]

Context config>router>mpls>lsp

Description This command specifies if this LSP will be included in LDP over RSVP. The no form of the command reverts to default operation.

Default no 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.

metric

Syntax **metric** *metric*

Context config>router>mpls>lsp

Description This command specifies the metric for this LSP which is used to select an LSP among a set of LSPs which are destined to the same egress router. The LSP with the lowest metric will be selected.

In LDP-over-RSVP, LDP performs a lookup in the Routing Table Manager (RTM) which provides the next hop to the destination PE and the advertising router (ABR or destination PE itself). 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). This lookup returns the best RSVP LSP to use to forward packets for an LDP FEC learned through the targeted LDP session. The lookup returns the LSP with the lowest metric. If multiple LSPs have the same metric, then the result of the lookup is to select the first one available in the TTM.

Default 1

Parameters *metric* — Specifies the metric for this LSP which is used to select an LSP among a set of LSPs which are destined to the same egress router.

Values 1 — 65535

to

Syntax **to** *ip-address*

Context config>router>mpls>lsp

Description This command specifies the system IP address 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.

Default No default

Parameters *ip-address* — The system IP address of the egress router.

vprn-auto-bind

Syntax **vprn-auto-bind** [**include** | **exclude**]

Context config>router>mpls>lsp

Description 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.

By configuring the command vprn-auto-bind exclude, the associated LSP will not be used by the auto-bind feature within VPRN services.

The **no** form of the command resets the flag backto the default value.

Default include

Parameters **include** — Allows an associated LSPto be used by auto-bin for vprn services

exclude — Disables the use of the associated LSP to be used with the auto-bind feature for VPRN services.

retry-limit

Syntax	retry-limit <i>number</i> no retry-limit
Context	config>router>mpls>lsp
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>The no form of this command revert the parameter to the default value.</p>
Default	0 (no limit, retries forever)
Parameters	<p><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.</p> <p>Values 0 — 10000</p>

retry-timer

Syntax	retry-timer <i>seconds</i> no retry-timer
Context	config>router>mpls>lsp
Description	<p>This command configures the time, in seconds, for LSP re-establishment attempts after it has failed.</p> <p>The no form of this command reverts to the default value.</p>
Default	30
Parameters	<p><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.</p> <p>Values 1 — 600</p>

rsvp-resv-style

Syntax	rsvp-resv-style [<i>se</i> <i>ff</i>]
Context	config>router>mpls>lsp
Description	<p>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.</p>
Default	<i>se</i>

- Parameters**
- ff* — 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.
 - se* — 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.

shutdown

Syntax	[no] shutdown
Context	config>router>mpls>lsp
Description	<p>This command disables the existing LSP including the primary and any standby secondary paths.</p> <p>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

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	<i>path-name</i> — 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

bandwidth

Syntax	bandwidth rate-in-mbps no bandwidth
Context	config>router>mpls>lsp>primary config>router>mpls>lsp>secondary
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 group-name [group-name...(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.

MPLS Commands

Use the **no** form of the command to remove the exclude command.

Default no exclude

Parameters *group-name* — Specifies the existing group-name to be excluded when an LSP is set up.

hop-limit

Syntax **hop-limit** *number*
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 *number* — The number of hops the LSP can traverse, expressed as an integer.

Values 2 — 255

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 remains an 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.

Values 1–255

record

Syntax	[no] record
Context	config>router>mpls>lsp>primary config>router>mpls>lsp>secondary
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
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

standby

Syntax	[no] standby
Context	config>router>mpls>lsp>secondary
Description	<p>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.</p> <p>The no form of this command specifies that the secondary LSP is signaled when the primary path LSP fails.</p>
Default	none

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	<p><i>lsp-name</i> — Name that identifies the LSP.</p> <p>Values Up to 32 alphanumeric characters.</p>

push

Syntax	push <i>label</i> nexthop <i>ip-address</i> no push [<i>label</i>]
Context	config>router>mpls>static-lsp
Description	<p>This command specifies the label to be pushed on the label stack and the next hop IP address for the static LSP.</p> <p>The no form of this command removes the association of the label to push for the static LSP.</p>
Parameters	<p>label — Specifies the use of the implicit label value for the push operation.</p> <p><i>label</i> — The label to push on the label stack. Label values 16 through 1,048,575 are defined as follows:</p> <ul style="list-style-type: none"> Label values 16 through 31 are 7210 SAS M 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. <p>Values 16 — 1048575</p> <p>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.</p>

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.
Parameters	<i>ip-address</i> — The system IP address of the egress router.
Default	none

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
	<p>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.</p> <p>The RDM model is defined using the following equations:</p> $\text{SUM (Reserved (CTc))} \leq \text{BCb,}$ <p>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.</p> $\text{BC0} = \text{Max-Reservable-Bandwidth, so that}$ $\text{SUM (Reserved(CTc))} \leq \text{Max-Reservable-Bandwidth,}$ <p>where the SUM is across all values of c in the range $0 \leq c \leq (\text{MaxCT} - 1)$.</p>

graceful-shutdown

Syntax	[no] graceful-shutdown
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>

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.

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).

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.

A node does not take any action on the paths of the following originating LSPs after receiving the PathErr message:

- 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

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	<p>This command enables the refresh reduction capabilities over all bypass tunnels originating on this 7210 SAS PLR node or terminating on this 7210 SAS Merge Point (MP) node.</p> <p>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.</p> <p>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.</p>
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>limit</i> no rapid-retry-limit
Context	config>router>rsvp
Description	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. The no form of this command reverts to the default value.
Default	3
Parameters	<i>limit</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

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 7210 SAS M 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 7210 SAS M 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.</p>

Interface Commands

The 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.

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 *authentication-key* — 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.

hash-key — 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 (“ ”).

This is useful when a user must configure the parameter, but for security purposes, the actual unencrypted key value is not provided.

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.

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.

bfd-enable

Syntax **[no] bfd-enable**

Context config>router>rsvp>interface

Description 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.

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.

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.

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

Description This command enables the use of the implicit null label for all LSPs.

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.

Interface Commands

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

The **no** form of this command returns the RSVP interface to use the RSVP level configuration value.

Default disable

Parameters **enable** — This parameter enables the implicit null label.
disable — This parameter disables the implicit null label.

refresh-reduction

Syntax **[no] refresh-reduction**

Context config>router>rsvp>interface

Description This command enables the use of the RSVP overhead refresh reduction capabilities on this RSVP interface.

The 7210 SAS node accepts bundle RSVP messages from its peer over the interface, performs reliable RSVP message delivery to its peer, and utilizes summary refresh messages to refresh the path and resv states. Reliable message delivery must be explicitly enabled by the user after refresh reduction is enabled.

The other two capabilities are immediately enabled.

A bundle message reduces the overall message handling load, it consists of a bundle header followed by one or more bundle sub-messages. A bundle sub-message is any RSVP message other than a bundle message. A 7210 node only processes the bundled RSVP messages received and does not generate them.

When reliable message delivery is supported by both the node and its peer over the RSVP interface, a RSVP message is sent with a message_id object. A message_id object can be added to any RSVP message or it can be a sub-message of a bundled message.

If a node 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 in the next RSVP message it sends to the node. Alternatively, an ACK message can also be used to send the message_ack object. In both cases, more than one message_ack object can be included in the same message.

The 7210 supports only the use of ACK messages to send message_ack object, but it can also process the received message_ack objects piggy-backed to hop-by-hop RSVP messages, such as path and resv.

The 7210 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 to handle unacknowledged message_id objects. A RSVP message with the same message_id is re-transmitted at an interval of $2 * \text{rapid-retransmit-time}$. The rapid-retransmit-time is referred to as the rapid retransmission interval and it should be lesser than the regular refresh interval configured in the config>router>rsvp>refresh-time context.

Rapid retry limit indicates the maximum number of retransmissions allowed for unacknowledged RSVP messages. The node stops the retransmission of unacknowledged RSVP messages when:

- The updated backoff interval exceeds the regular refresh interval.
- The number of retransmissions reaches the value of the **rapid-retry-limit** parameter, whichever comes first.

These two parameters can be configured on a system in the **config>router>rsvp** context.

Refresh summary consists of sending a summary refresh messages containing message_id list objects.

The fields of the message_id list object are populated with the values from the message_identifier field in the message_id object of a previously sent individual path or resv message. The summary refresh message is sent per refresh regular interval. The interval is 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. If any message_identifier field does not match, the node sends a message_id_nack object to the originator of the message.

The above capabilities are collectively referred to as “Refresh Overhead Reduction Extensions”. When refresh-reduction is enabled on an RSVP interface, the node sets a “refresh-reduction-capable” bit in the flag field of the common RSVP header. If both peers on a RSVP interface set the “refresh-reduction-capable” bit, all the refresh overhead reduction extensions can be implemented. The node monitors the bit in all the RSVP messages received from the peer. The router stops sending summary refresh messages once the bit is cleared, the node does not send summary refresh messages if the bit is not set by the peer.

A node (with refresh reduction, reliable message delivery enabled) attempts to perform reliable message delivery even if the “refresh-reduction-capable” bit is not set by the peer. If the peer does not support the message_id object, it returns an error message “unknown object class”. The node retransmits the RSVP message without the message_id object and adopts the same message handling method for all future messages sent to the peer.

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 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 7210 SAS 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.

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.

Interface Commands

The **no** form of the command reverts to the default value.

Default no reliable-delivery

subscription

Syntax **subscription** *percentage*
no subscription

Context config>router>rsvp>interface

Description 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.

When the **subscription** is set to zero, no new sessions are permitted on this interface. If the *percentage* is exceeded, the reservation is rejected and a log message is generated.

The **no** form of this command reverts the *percentage* to the default value.

Default 100

Parameters *percentage* — The percentage of the interface's bandwidth that RSVP allows to be used for reservations.

Values 0 — 1000

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#
```

bypass-tunnel

- Syntax** `bypass-tunnel [to ip-address] [protected-lsp [lsp-name]] [dynamic | manual] [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.
- Output** **MPLS Bypass Tunnel Output Fields** — The following table describes MPLS bypass tunnel output fields.

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-A>show>router>mpls# bypass-tunnel
=====
MPLS Bypass Tunnels
=====
Legend :  m - Manual      d - Dynamic
=====
To           State  Out I/F      Out Label    Reserved  Protected  Type
           BW (Kbps)  LSP Count
-----
10.10.36.3   Up    lag-1:10     131066       0         2         d
10.10.23.2   Up    lag-1:10     130454       0         4         d
10.10.46.4   Up    lag-2        130592       0         4         d
10.10.36.6   Up    lag-2        130591       0         2         d
-----
Bypass Tunnels : 4
=====
*A:Dut-A>show>router>mpls#
```

interface

Syntax `interface [ip-int-name | ip-address] [label-map label]`
`interface [ip-int-name | ip-address]`

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

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.
Opr	Specifies the operational state of the interface.
Te-metric	Specifies the traffic engineering metric used on the interface.
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.

```
*A:Dut-A# show router mpls interface
```

```
=====
```

```
MPLS Interfaces
```

```
=====
```

Interface	Port-id	Adm	Opr
system	system	Up	Up

```
-----
```

Show Commands

```

    Admin Groups          None
to_D                    lag-1:10          Up          Up
    Admin Groups          CLOCKWISE
to_B                    lag-2           Up          Up
    Admin Groups          ANTI_CLOCKWISE
-----
Interfaces : 3
=====
*A:Dut-A#
*A:Dut-A# show router mpls interface to_B
=====
MPLS Interface : to_B
=====
Interface                Port-id          Adm          Opr
-----
to_B                    lag-2           Up          Up
    Admin Groups          ANTI_CLOCKWISE
-----
Interfaces : 1
=====
*A:Dut-A#

```

label

Syntax `label start-label [end-label | in-use | label-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: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:Dut-A# show router mpls label-range
=====
Label Ranges
=====
Label Type      Start Label    End Label      Aging          Total Available
-----
static-lsp      32             1023           -              992
static-svc      2048           18431          -              16384
dynamic         32768          131071         0              102400
=====
*A:Dut-A#

```

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.

Label	Description (Continued)
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	<p>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.</p> <p>disabled – There is no detour LSP from each node on the primary path.</p>
ADSPEC	<p>enabled – The LSP will include advertising data (ADSPEC) objects in RSVP messages.</p> <p>disabled – The LSP will not include advertising data (ADSPEC) objects in RSVP messages.</p>
Primary	The preferred path for the LSP.
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	se – 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.

Label	Description (Continued)
	<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.
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.
LastResignalAt-tempt	Displays the system up time when the last attempt to resignal this LSP was made.
VprnAutoBind	Displays the status on VPRN auto-bind feature as enabled or disabled.

Sample Output

```
*A:SRU4>config>router>mpls# show router mpls lsp "to_110_20_1_1_cspf"
=====
MPLS LSPs (Originating)
=====
LSP Name                               To                               Fastfail   Adm   Opr
                                   Config
-----
to_110_20_1_1_cspf                     110.20.1.1                       No         Up    Up
-----
LSPs : 1
=====
*A:SRU4>config>router>mpls#

*A:Dut-A# show router mpls lsp transit detail
=====
MPLS LSPs (Transit) (Detail)
=====
LSP D_B_1::D_B_1
-----
From           : 10.20.1.4           To           : 10.20.1.2
State          : Up
In Interface   : lag-1:10            In Label     : 130668
Out Interface  : lag-2              Out Label    : 131065
Previous Hop   : 10.10.14.4         Next Hop     : 10.10.12.2
Reserved BW    : 0 Kbps
-----
*A:Dut-A#

*=====
*A:7210-SAS>show>router>mpls# lsp A detail
=====
MPLS LSPs (Originating) (Detail)
=====
```

```

-----
Type : Originating
-----
LSP Name      : A                      LSP Tunnel ID : 1
From          : 2.2.2.2                To             : 100.100.100.100
Adm State     : Up                     Oper State     : Down
LSP Up Time   : 0d 00:00:00            LSP Down Time : 0d 00:05:42
Transitions   : 2                     Path Changes   : 2
Retry Limit   : 0                     Retry Timer    : 30 sec
Signaling     : RSVP                   Resv. Style    : SE
Hop Limit     : 255                   Negotiated MTU : 0
Adaptive      : Enabled                 ClassType      : 0
FastReroute   : Disabled               Oper FR        : Disabled
CSPF          : Disabled               ADSPEC        : Disabled
Metric        : 0
Include Grps  :                        Exclude Grps   :
None                                                  None
Type          : RegularLsp             Least Fill     : Disabled
LdpOverRsvp   : Enabled                VprnAutoBind  : Enabled
Oper Metric   : 65535

Primary       : A                      Down Time      : 0d 00:05:42
Bandwidth     : 0 Mbps
-----

```

```
*A:7210-SAS>show>router>mpls# lsp 2 detail
```

```
*A:Dut-A# show router mpls lsp A_D_15 path detail
```

```
=====
MPLS LSP A_D_15 Path (Detail)
=====
```

```
Legend :
```

```

@ - Detour Available          # - Detour In Use
b - Bandwidth Protected      n - Node Protected
=====

```

```
-----
LSP A_D_15 Path A_D_15
-----
```

```

LSP Name      : A_D_15                 Path LSP ID   : 19002
From          : 10.20.1.1              To            : 10.20.1.4
Adm State     : Up                     Oper State    : Up
Path Name     : A_D_15                 Path Type     : Primary
Path Admin    : Up                     Path Oper     : Up
OutInterface  : lag-1:10                Out Label    : 130607
Path Up Time  : 0d 00:19:18            Path Dn Time  : 0d 00:00:00
Retry Limit   : 0                     Retry Timer   : 10 sec
RetryAttempt  : 0                     Next Retry   *: 0 sec
Bandwidth     : 10 Mbps                 Oper Bandwi  *: 10 Mbps
Hop Limit     : 255
Record Route  : Record                 Record Label : Record
Oper MTU      : 9194                   Negotiated  *: 9194
Adaptive      : Enabled                 MBB State    : N/A
Include Grps  :                        Exclude Grps  :
None                                                  None
Path Trans    : 2                       CSPF Queries : 34
Failure Code  : noError                 Failure Node  : n/a
ExplicitHops  :
    10.20.1.4
Actual Hops   :
    10.10.14.1(10.20.1.1) @
    -> 10.10.14.4(10.20.1.4)
Record Label  : N/A
Record Label  : 130607

```

Show Commands

```
ComputedHops:
  10.10.14.1      -> 10.10.14.4

Detour Stat*: Standby
Detour Avoi*: 10.10.14.4
Detour Acti*: n/a
In Interface: n/a
Out Interfa*: lag-2
Next Hop      : 10.10.12.2
Explicit Ho*:
  10.10.12.1    -> 10.10.12.2    -> 10.10.23.3    -> 10.10.36.6
-> 10.10.46.4

Detour Type : Originate
Detour Orig*: 10.20.1.1
Detour Up T*: 0d 00:18:36
In Label    : n/a
Out Label   : 130975
```

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#

```

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.

Show Commands

<u>Label</u>	<u>Description (Continued)</u>
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#
```

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.

```
*A:Dut-A# 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

LSP Counts      Originate      Transit          Terminate
-----
Static LSPs     12              6                12
Dynamic LSPs    96              122              95
```

```
Detour LSPs          48          64          32
=====
*A:Dat-A#
```

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 IP address and the number of packets sent and received on an interface-basis.

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.
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.
Active Resvs	The total number of active RSVP sessions that have reserved bandwidth.

Label	Description (Continued)
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.
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.
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.

Show RSVP Commands

Label	Description (Continued)
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.
Hellos	Total RSVP RESV HELLO REQ messages received on the interface.
Bfd Enabled	Yes — BFD is enabled on the RSVP interface No — BFD is disabled on the RSVP interface.

Sample Output

```
*A:7210-SAS>show>router>rsvp# interface detail
```

```
=====
RSVP Interfaces (Detailed)
=====
```

```
-----
Interface : system
-----
```

```
Interface      : system
Port ID       : system
Admin State   : Up
Active Sessions : 0
Total Sessions : 0
Subscription  : 100 %
Total BW      : 0 Mbps
Hello Interval : 3000 ms
Authentication : Disabled
Auth Rx Seq Num : n/a
Auth Tx Seq Num : n/a
Refresh Reduc. : Disabled
Bfd Enabled   : No
Oper State    : Up
Active Resvs  : 0
Port Speed    : 0 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 : 0
BC1 : 0
BC2 : 0
BC3 : 0
BC4 : 0
BC5 : 0
BC6 : 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

```

No Neighbors.

Interface : ip-10.10.12.3

```

Interface       : ip-10.10.12.3
Port ID         : 1/1/9
Admin State     : Up                Oper State      : Up
Active Sessions : 1                Active Resvs    : 0
Total Sessions  : 1
Subscription    : 100 %            Port Speed      : 1000 Mbps
Total BW        : 1000 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.  : Disabled         Reliable Deli.  : Disabled
Bfd Enabled     : No                Graceful Shut.  : 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             : 1000000          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     : 1000000
TE1-> Resv. Bw   : 0                Unresv. Bw     : 1000000
TE2-> Resv. Bw   : 0                Unresv. Bw     : 1000000
TE3-> Resv. Bw   : 0                Unresv. Bw     : 1000000
TE4-> Resv. Bw   : 0                Unresv. Bw     : 1000000
TE5-> Resv. Bw   : 0                Unresv. Bw     : 1000000
TE6-> Resv. Bw   : 0                Unresv. Bw     : 1000000
TE7-> Resv. Bw   : 0                Unresv. Bw     : 1000000

```

Neighbors : 10.10.12.2

Interface : ip-10.10.4.3

```

Interface       : ip-10.10.4.3
Port ID         : 1/1/8
Admin State     : Up                Oper State      : Up
Active Sessions : 1                Active Resvs    : 0
Total Sessions  : 1
Subscription    : 100 %            Port Speed      : 1000 Mbps
Total BW        : 1000 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

```

Show RSVP Commands

```
Refresh Reduc.      : Disabled          Reliable Deli.      : Disabled
Bfd Enabled         : No                Graceful Shut.     : 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                : 1000000            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        : 1000000
TE1-> Resv. Bw    : 0                  Unresv. Bw        : 1000000
TE2-> Resv. Bw    : 0                  Unresv. Bw        : 1000000
TE3-> Resv. Bw    : 0                  Unresv. Bw        : 1000000
TE4-> Resv. Bw    : 0                  Unresv. Bw        : 1000000
TE5-> Resv. Bw    : 0                  Unresv. Bw        : 1000000
TE6-> Resv. Bw    : 0                  Unresv. Bw        : 1000000
TE7-> Resv. Bw    : 0                  Unresv. Bw        : 1000000
Neighbors          : 10.10.4.2
```

Interface : ip-10.10.2.3

```
Interface          : ip-10.10.2.3
Port ID            : 1/1/4
Admin State        : Up                Oper State         : Down
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.     : Disabled          Reliable Deli.     : Disabled
Bfd Enabled        : No                Graceful Shut.     : 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
No Neighbors.
=====
```

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.

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, manual-bypass
- 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.

Show RSVP Commands

Label	Description (Continued)
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#
```

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: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.

Label	Description (Continued)
	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.

Sample Output

```
*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
=====
*A:SRU4>show>router>rsvp#
```

Tools Commands

cspf

Syntax	cspf to <i>ip-addr</i> [from <i>ip-addr</i>] [bandwidth <i>bandwidth</i>] [include-bitmap <i>bitmap</i>] [exclude-bitmap <i>bitmap</i>] [hop-limit <i>limit</i>] [exclude-address <i>excl-addr</i> [<i>excl-addr...</i> (up to 8 max)]] [use-te-metric] [skip-interface <i>interface-name</i>]
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>skip-interface <i>interface-name</i> — Specifies an interface name that should be skipped during setup.</p>

resignal

Syntax	resignal { lsp <i>lsp-name path path-name</i> delay <i>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	<p><i>lsp-name</i> — Specifies an existing LSP name to resignal.</p> <p><i>path-name</i> — Specifies an existing path name to resignal.</p> <p>delay <i>minutes</i> — Configures the global timer or all LSPs to resignal.</p>

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

Clear Commands

fec-egress-statistics

Syntax	fec-egress-statistics [<i>ip-prefix/mask</i>]
Context	clear>router>ldp
Description	This command resets or clears LDP FEC egress statistics.
Parameters	<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 sub-net address.</p> <p>Values 0 — 32</p>

interface

Syntax	interface <i>ip-int-name</i>
Context	clear>router>mpls
Description	This command resets or clears statistics for MPLS interfaces.
Parameters	<i>ip-int-name</i> — The name of an existing IP interface. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

lsp

Syntax	lsp <i>lsp-name</i>
Context	clear>router>mpls
Description	This command resets and restarts an LSP.
Parameters	<i>lsp-name</i> — The name of the LSP to clear up to 64 characters in length.

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.

Clear Commands

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>] no mpls
Context	debug>router
Description	This command enables and configures debugging for MPLS.
Parameters	<p>lsp <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>sender <i>source-address</i> — The system IP address of the sender.</p> <p>endpoint <i>endpoint-address</i> — The far-end system IP address.</p> <p>tunnel-id <i>tunnel-id</i> — The MPLS SDP ID.</p> <p>Values 0 — 4294967295</p> <p>lsp-id <i>lsp-id</i> — The LSP ID.</p> <p>Values 1 — 65535</p> <p>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.</p>

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.

Debug Commands

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.

auth

Syntax	auth no auth
Context	debug>router>rsvp>event
Description	This command debugs authentication events. The no form of the command disables the debugging.
Parameters	detail — Displays detailed information about authentication 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 debugs MPLS IOM events. The no form of the command disables the debugging.

Parameters **detail** — Displays detailed information about MPLS IOM events.

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.

Debug Commands

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.

rr

Syntax	rr no rr
Context	debug>router>rsvp>event
Description	This command debugs refresh reduction events. The no form of the command disables the debugging.
Parameters	detail — Displays detailed information about refresh reduction events.

packet

Syntax	[no] packet
Context	debug>router>rsvp>
Description	This command enters the syntax to debug packets.

Debug Commands

ack

Syntax	ack [detail] no ack
Context	debug>router>rsvp>packet
Description	This command debugs ack packets. The no form of the command disables the debugging.
Parameters	detail — Displays detailed information about RSVP-TE ack packets.

bundle

Syntax	bundle [detail] no bundle
Context	debug>router>rsvp>packet
Description	This command debugs bundle packets. The no form of the command disables the debugging.
Parameters	detail — Displays detailed information about RSVP-TE bundle 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.

Debug Commands

resvrr

Syntax	resvrr [detail] no resvrr
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.

resvtear

Syntax	resvtear [detail] no resvtear
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.

srefresh

Syntax	srefresh [detail] no srefresh
Context	debug>router>rsvp>packet
Description	This command debugs srefresh packets. The no form of the command disables the debugging.
Parameters	detail — Displays detailed information about RSVP-TE srefresh packets.

Label Distribution Protocol

In This Chapter

This chapter provides information to enable Label Distribution Protocol (LDP).

Topics in this chapter include:

- [Label Distribution Protocol on page 154](#)
 - [LDP and MPLS on page 154](#)
 - [LDP Architecture on page 155](#)
 - [Subsystem Interrelationships on page 156](#)
 - [Execution Flow on page 158](#)
 - [Label Exchange on page 159](#)
 -
- [LDP Process Overview on page 168](#)

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.

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 (DUS).

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 7210 SAS M 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, DUS, 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 8](#). 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.

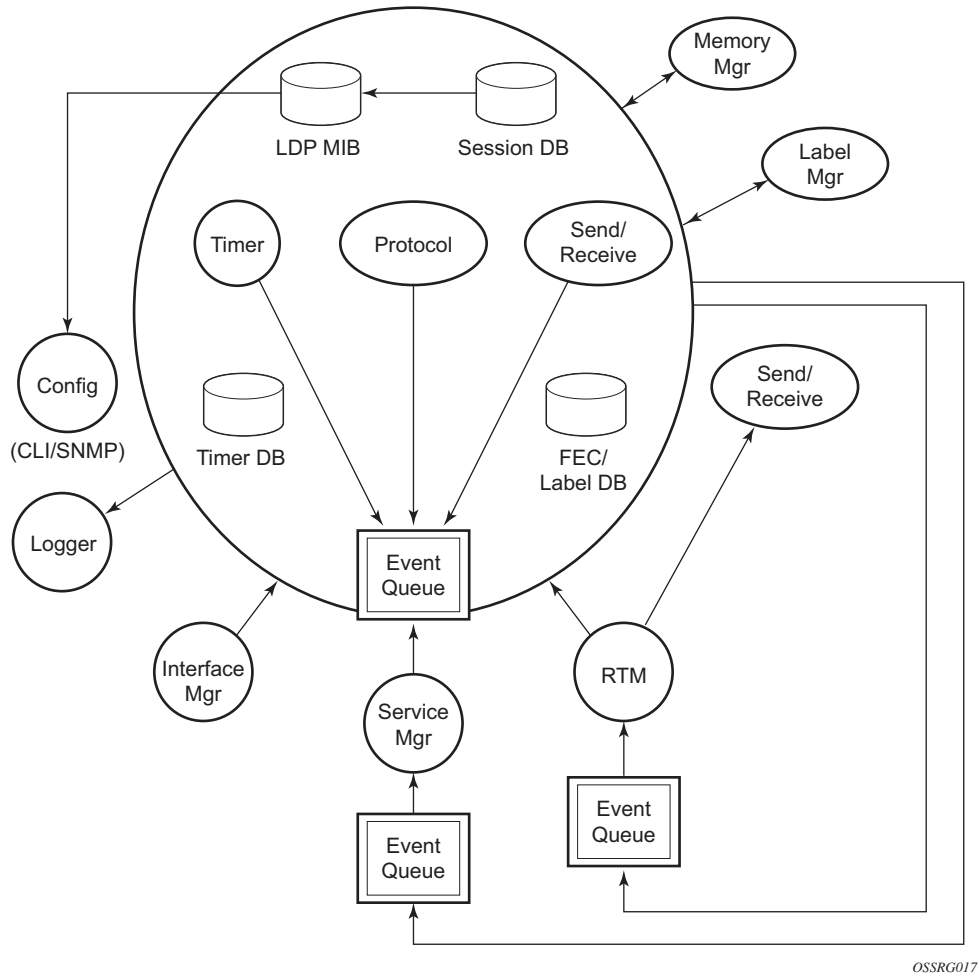


Figure 8: Subsystem Interrelationships

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 7210 SAS M uses a label range from 28672 (28K) to 131071 (128K-1) to allocate all dynamic labels, including RSVP allocated labels and VC labels.

LDP Configuration

The 7210 SAS M 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 7210 SAS M OS is limited to service-related signaling. Therefore, the configurable parameters are restricted to system-wide parameters, such as hello and keepalive timeouts.

Initialization

MPLS must be enabled when LDP is initialized. 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.

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.

LDP Filters

Both inbound and outbound LDP label binding filtering is supported.

Inbound filtering (import policy) allows configuration of a policy to control the label bindings an LSR accepts from its peers. Label bindings can be filtered based on:

- Neighbor: Match on bindings received from the specified peer.
- Prefix-list: Match on bindings with the specified prefix/prefixes.

Note: The default import behavior is to accept all FECs received from peers. The LDP export policy can be used to explicitly add FECs (or non-LDP routes) for label propagation and does not filter out or stop propagation of any FEC received from neighbors.

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.

Note: The LDP export policy will not filter out FECs. It is only used to explicitly add FECs (or non-LDP routes) for label propagation.

The default export behavior originates label bindings for system address and propagate all FECs received.

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.

NOTE:

- Use of implicit NULL MPLS label must be enabled with use of LDPoRSVP. Use the command `configure>router>rsvp> implicit-null-label` and `configure> router> ldp> implicit-null-label` to enable use of Implicit NULL MPLS labels.
 - Only FRR one-to-one is supported when LDPoRSVP is used. FRR facility is not supported. This is not blocked in CLI, but operators need to ensure it when configuring the nodes.
-

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 9: LDP over RSVP Application

The network displayed in [Figure 9](#) 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 via 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 9](#) 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 9](#), 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 10](#).

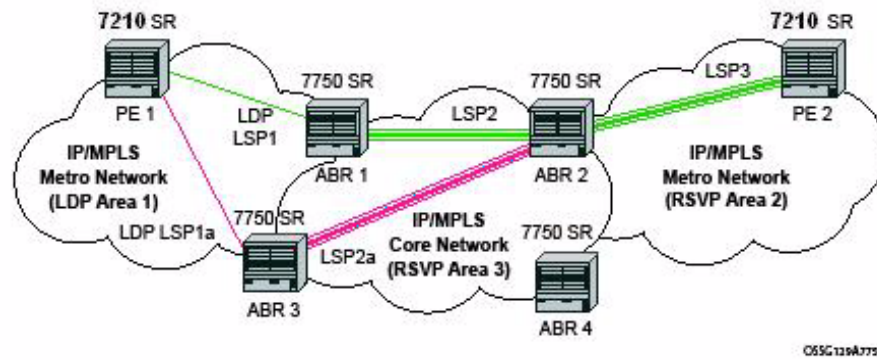


Figure 10: LDP over RSVP Application Variant

Signaling and Operation

- [LDP Label Distribution and FEC Resolution on page 164](#)
 - [Default FEC Resolution Procedure on page 165](#)
-

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. 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 166](#)
 - [ABR Protection on page 166](#)
-

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.
-

NOTE: Only FRR one-to-one is supported with LDP-over-RSVP with use of implicit NULL label. In other words, implicit NULL label must be enabled to use FRR one-to-one. FRR facility cannot be used. The software does not make any checks to enforce these restrictions. Operators must ensure this by network design and configuration.

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.

Configuring Implicit Null Label

The implicit null label option allows a 7210 SAS 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 FEC signaling by the LDP control protocol.

The user can configure to signal 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.

Multi-Area and Multi-Instance Extensions to LDP

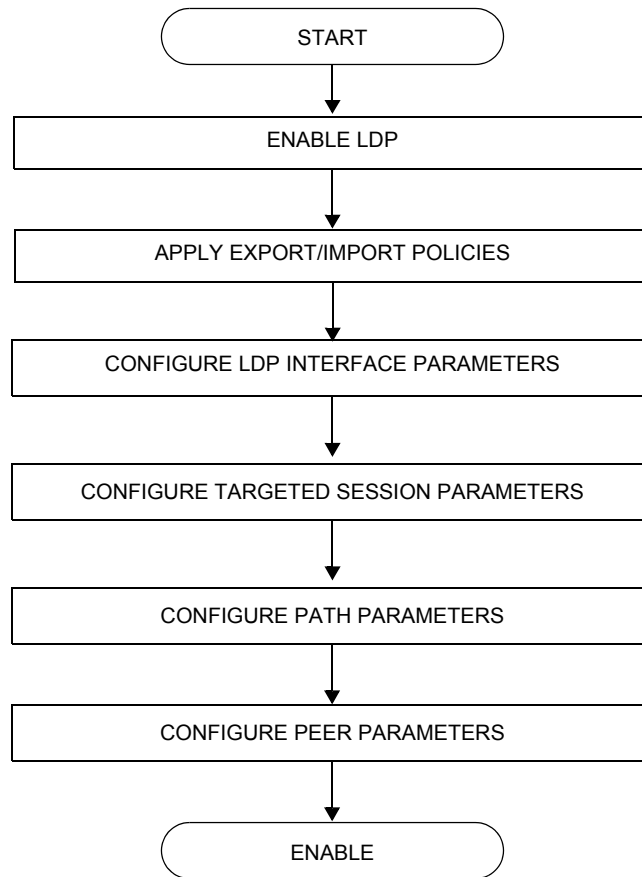
To extend LDP across multiple areas of an IGP instance or across multiple IGP instances, the current standard LDP implementation based on RFC 3036 requires that all the /32 prefixes of PEs be leaked between the areas or instances. This is because an exact match of the prefix in the routing table has to install the prefix binding in the LDP Forwarding Information Base (FIB).

The 7210 SAS performs this function by default, except in cases when the 7210 SAS is configured as Area Border Router (ABR). In this scenario, the convergence of IGP on routers increases 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). The ABR is 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 Process Overview

Figure 11 displays the process to provision basic LDP parameters.



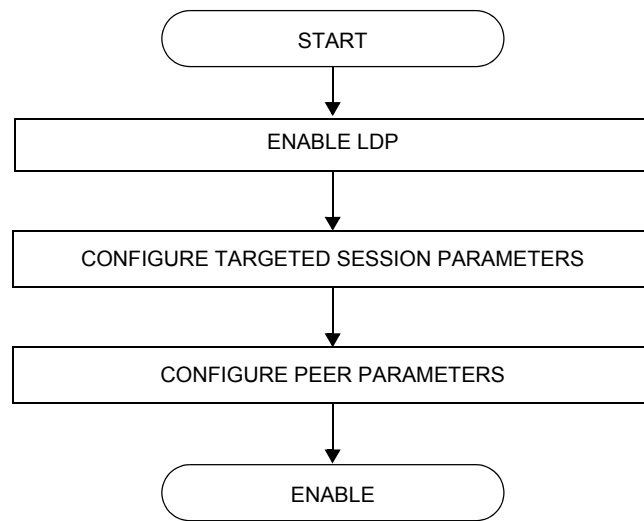


Figure 11: LDP Configuration and Implementation

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 172](#)
- [Basic LDP Configuration on page 173](#)
- [Common Configuration Tasks on page 174](#)
- [LDP Configuration Management Tasks on page 183](#)

LDP Configuration Overview

When the 7210 SAS M OS 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*.

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.

```
A:ALU_SIM11>config>router>ldp# info
-----
      aggregate-prefix-match
        prefix-exclude "sample"
      exit
    graceful-restart
    exit
  peer-parameters
    peer 1.1.1.1
      ttl-security 1
    exit
  exit
interface-parameters
  interface "a"
  exit
exit
targeted-session
exit
-----
A:ALU_SIM11>config>router>ldp#
```

Common Configuration Tasks

This section provides information to configure:

- [Enabling LDP on page 174](#)
- [Targeted Session Parameters on page 178](#)
- [Peer Parameters on page 181](#)

Enabling LDP

LDP must be enabled in order for the protocol to be active. MPLS must also be enabled. MPLS is enabled in the `config>router>mpls` context.

Use the following syntax to enable LDP on a 7210 SAS M OS router:

CLI Syntax: `ldp`

Example: `config>router# ldp`

The following displays the enabled LDP configuration.

```
A:ALU_SIM11>config>router>ldp# info
-----
      aggregate-prefix-match
        prefix-exclude "sample"
      exit
      graceful-restart
      exit
      peer-parameters
        peer 1.1.1.1
          ttl-security 1
        exit
      exit
      interface-parameters
        interface "a"
        exit
      exit
      targeted-session
      exit
-----
A:ALU_SIM11>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. Alcatel-Lucent'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
                [no] maximum-recovery-time interval
                [no] neighbor-liveness-time interval
```

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)]
A:ALU_SIM11>config>router>ldp# info
-----
aggregate-prefix-match
  prefix-exclude "sample"
exit
graceful-restart
exit
peer-parameters
  peer 1.1.1.1
    ttl-security 1
  exit
exit
interface-parameters
  interface "a"
  exit
exit
targeted-session

```


exit

Targeted Session Parameters

Use the following syntax to specify **targeted-session** parameters:

```
CLI Syntax: config>router# ldp
                targeted-session
                  disable-targeted-session
                  hello timeout factor
                  keepalive timeout factor
                  peer ip-address
                    hello timeout factor
                    keepalive timeout factor
                  no shutdown
```

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:ALU_SIM11>config>router>ldp# info
-----
aggregate-prefix-match
  prefix-exclude "sample"
exit
graceful-restart
exit
peer-parameters
  peer 1.1.1.1
    ttl-security 1
  exit
exit
interface-parameters
  interface "a"
  exit
exit
targeted-session
exit
-----
```

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]
```

```
A:ALA-1>config>router>ldp# info
-----
peer-parameters
  peer 10.10.10.104
    authentication-key "3WErEDozxyQ" hash
  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's far-end point. The exchange of LDP hellos trigger session establishment. The SDP's signaling default enables **tldp**. The service SDP uses the targeted-session parameters configured in the **config>router>ldp>targeted-session** context.

The 7210 SAS M supports only Targeted LDP (TLDP).

The following example displays the command syntax usage to configure enable LDP on an MPLS SDP:

CLI Syntax: `config>service>sdp#
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
lsp A_D_1
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#
```

LDP Configuration Management Tasks

This section discusses the following LDP configuration management tasks:

- [Disabling LDP on page 183](#)
 - [Modifying Targeted Session Parameters on page 184](#)
-

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-established. 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

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 **interface-parameter** command reverts modified values back to the defaults.

The following output displays the default values:

```
A:ALU_SIM11>config>router>ldp>targ-session# info detail
-----
                no disable-targeted-session
                hello 45 3
                keepalive 40 4
-----
A:ALU_SIM11>config>router>ldp>targ-session#
```

LDP Command Reference

Command Hierarchies

- [LDP Commands on page 187](#)
- [Show Commands on page 188](#)
- [Clear Commands on page 188](#)
- [Debug Commands on page 189](#)

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
      — export policy-name [policy-name...(up to 5 max)]
      — no export
      — [no] export-tunnel-table policy-name
      — [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...(u pto 5 max)]
      — interface-parameters
        — hello timeout factor
        — no hello
        — [no] interface ip-int-name
          — hello timeout factor
          — no hello
          — keepalive timeout factor
          — no keepalive
          — [no] shutdown
          — transport-address {system | interface}
        — keepalive timeout factor
        — no keepalive
        — transport-address {system | interface}
      — label-withdrawal-delay seconds
      — peer-parameters
        — peer ip-address
        — no peer [ip-address]
          — auth-keychain name
          — authentication-key [authentication-key | hash-key] [hash | hash2]
          — no authentication-key
          — ttl-security min-ttl-value

```

- [no] **ttl-security**
- [no] **shutdown**
- **targeted-session**
 - [no] **disable-targeted-session**
 - **hello** *timeout factor*
 - **no hello**
 - **keepalive** *timeout factor*
 - **no keepalive**
 - **peer** *ip-address*
 - **no peer** *ip-address*
 - **hello** *timeout factor*
 - **no hello**
 - **keepalive** *timeout factor*
 - **no keepalive**
 - [no] **shutdown**
- **tunnel-down-damp-time** *seconds*
- **no tunnel-down-damp-time**

Show Commands

- show**
 - **router**
 - **ldp**
 - **auth-keychain** [*keychain*]
 - **bindings**[**fec-type** *fec-type* [**detail**]] [**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*]
 - **bindings service-id** *service-id* [**detail**]
 - **bindings vc-type** *vc-type* [{**vc-id** *vc-id* | **agi** *agi*} [**session** *ip-addr[:label-space]*]]
 - **discovery** [{**peer** [*ip-address*]} | {**interface** [*ip-int-name*]}] [**state** *state*] [**detail**]
 - **interface** [*ip-int-name* | *ip-address*] [**detail**]
 - **parameters**
 - **peer** [*ip-address*] [**detail**]
 - **peer-parameters** *peer-ip-address*
 - **session** [*ip-addr[:label-space]*] [**detail** | **statistics** [*packet-type*]]
 - **status**

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
```

LDP Configuration Commands

Generic Commands

ldp

Syntax	<code>[no] ldp</code>
Context	<code>config>router</code>
Default	This command creates the context to configure an LDP parameters. LDP is not enabled by default and must be explicitly enabled (no shutdown). To suspend the LDP protocol, use the shutdown command. Configuration parameters are not affected. 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.
Default	none (LDP must be explicitly enabled)

shutdown

Syntax	<code>[no] shutdown</code>
Context	<code>config>router>ldp</code> <code>config>router>ldp>target-session>peer</code> <code>config>router>ldp>interface</code> <code>config>router>ldp>aggregate-prefix-match</code>
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

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	<p>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.</p> <p>The no form of this command removes all policies from the configuration.</p>
Default	no prefix-exclude.

export

Syntax	export <i>policy-name</i> [<i>policy-name</i> ...upto 5 max] no export
Context	config>router>ldp
Description	<p>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.</p> <p>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.</p> <p>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 five policy names can be specified.</p> <p>The no form of the command removes all policies from the configuration.</p>
Default	no export — No export route policies specified.
Parameters	<p><i>policy-name</i> — 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.</p> <p>The specified name(s) must already be defined.</p>

export-tunnel-table

Syntax	[no] export-tunnel-table <i>policy-name</i>
Context	config>router>ldp
Description	<p>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.</p> <p>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.</p> <p>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.</p> <p>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:</p> <ul style="list-style-type: none"> • 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 is 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.

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.

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

import

Syntax	import <i>policy-name</i> [<i>policy-name</i> ...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	<i>policy-name</i> — 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 <i>seconds</i>
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 transport tunnels (IPv4 prefix FECs) and is not applicable to pseudowires (service FECs).

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

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 IOM NHLFE has been de-programmed.

If the FEC gets resolved before the tunnel-down-damp-time timer expires, then LDP programs the IOM 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.

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.

Values 0 — 20

keepalive

Syntax **keepalive** *timeout factor*
no keepalive

Context config>router>ldp>interface-parameters
config>router>ldp>targ-session
config>router>ldp>targ-session>peer
config>router>ldp>if-params>if

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 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 3 — 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

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.

hello

Syntax **hello** *timeout factor*
no hello

Context config>router>ldp>interface-parameters
 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.

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 targeted-session level sets the **hello timeout** and the **hello factor** to the default values.

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 3 — 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

interface

Syntax [**no**] **interface** *ip-int-name*

Context config>router>ldp>if-params

Description This command enables LDP on the specified IP interface.

The **no** form of the command deletes the LDP interface and all configuration information associated with the LDP interface.

The LDP interface must be disabled using the **shutdown** command before it can be deleted.

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.

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> <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.</p>
Default	system — The system IP address is used.
Parameters	<p>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.</p> <p>system — The system IP address is used to set up the LDP session between neighbors.</p>

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>

tll-security

Syntax	tll-security <i>min-ttl-value</i> no tll-security
Context	config>router>ldp>peer-parameters>peer
Description	<p>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.</p> <p>The no form of the command disables TTL security.</p>
Default	no tll-security
Parameters	<p><i>min-ttl-value</i> — Specify the minimum TTL value for an incoming packet.</p> <p>Values 1 — 255</p>

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

disable-targeted-session

Syntax	[no] disable-targeted-session
Context	config>router>ldp>targ-session
Description	<p>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.</p> <p>The no form of the command enables the set up of any targeted sessions.</p>
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.

Show LDP Commands

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	bindings [<i>fec-type fec-type</i> [<i>detail</i>]] [<i>session ip-addr[:label-space]</i>] bindings <i>label-type start-label</i> [<i>end-label</i>] bindings { <i>prefix ip-prefix/mask</i> [<i>detail</i>]} [<i>session ip-addr[:label-space]</i>] bindings active [<i>prefix ip-prefix/mask</i>] bindings service-id <i>service-id</i> [<i>detail</i>] bindings vc-type <i>vc-type</i> [{ <i>vc-id vc-id</i> <i>agi agi</i> }] [<i>session ip-addr[:label-space]</i>]
Context	show>router>ldp
Description	This command displays the contents of the label information base.
Parameters	detail — Displays detailed information. <i>label-space</i> — Specifies the label space identifier that the router is advertising on the interface. Values 0 — 65535 <i>start-label</i> — Specifies a label value to begin the display. Values 16 — 1048575 <i>end-label</i> — Specifies a label value to end the display. Values 17 — 1048575

Show Commands

vc-type — Specifies the VC type to display.

Values ethernet , vlan , mirror

vc-id — Specifies the VC ID to display.

Values 1 — 4294967295

service-id — Specifies the service ID number to display.

Values 1 — 2147483647

Output **LDP Bindings Output** — The following table describes the LDP bindings fields.

Label	Description
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)
Type	The service type exchanging labels. 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.
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-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

```

```

LDP Prefix Bindings

```

```

Prefix          Peer          IngLbl  EgrLbl  EgrIntf/LspId  EgrNextHop
-----

```

```

No Matching Entries Found

```

```

LDP Service Bindings

```

```

Type  VCId      SvcId      SDPIId  Peer          IngLbl  EgrLbl  LMTU  RMTU
-----
V-Vlan 201      201        1210    10.20.1.2    130604U 131036 1500 1500
V-Vlan 201      201        1410    10.20.1.4    130728U 131016S 1500 1500
V-Vlan 202      202        1210    10.20.1.2    130603U 131035 1500 1500
V-Vlan 202      202        1410    10.20.1.4    130727U 131003S 1500 1500
V-Vlan 203      203        1210    10.20.1.2    130602U 131034 1500 1500
...

```

```

*A:Dut-A#

```

```

*A:Dut-A# show router ldp bindings ingress-label 32768 131071

```

```

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

```

```

LDP Prefix Bindings

```

```

Prefix          Peer          IngLbl  EgrLbl  EgrIntf          EgrNextHop
-----

```

```

No Matching Entries Found

```

```

LDP Service Bindings

```

```

Type  VCId      SvcId      SDPIId  Peer          IngLbl  EgrLbl  LMTU  RMTU
-----
V-Vlan 201      201        1210    10.20.1.2    130604U 131036 1500 1500
V-Vlan 201      201        1410    10.20.1.4    130728U 131016S 1500 1500
V-Vlan 202      202        1210    10.20.1.2    130603U 131035 1500 1500
V-Vlan 202      202        1410    10.20.1.4    130727U 131003S 1500 1500
V-Vlan 203      203        1210    10.20.1.2    130602U 131034 1500 1500
...

```

```

*A:Dut-A#

```

```

*A:ALU_SIM11>show>router>ldp# bindings

```

```

LDP LSR ID: 1.1.1.2

```

```

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

```

Show Commands

```
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
                Peer          IngLbl    EgrLbl  EgrIntf/  EgrNextHop
                Peer          IngLbl    EgrLbl  LspId
-----
1.1.1.1/32      1.1.1.1      --        131071  1/1/1     6.6.6.1
1.1.1.2/32      1.1.1.1      131071U   --      --        --
-----
No. of Prefix Bindings: 2
=====
LDP Service FEC 128 Bindings
=====
Type  VCId      SvcId      SDFId  Peer          IngLbl  EgrLbl  LMTU  RMTU
-----
No Matching Entries Found
=====
LDP Service FEC 129 Bindings
=====
AGI          SvcId      SDFId  SAII  Peer          TAII
Type          SvcId      SDFId  Peer          IngLbl  EgrLbl  LMTU  RMTU
-----
No Matching Entries Found
=====
*A:ALU_SIM11>show>router>ldp#
```

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.
Local IP Address	The IP address used locally for the LDP session.

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:ALU_SIM11>show>router>ldp# interface
=====
LDP Interfaces
=====
Interface                Adm Opr  Hello Hold  KA    KA    Transport
                        Factor Time  Factor Timeout Address
-----
a                        Up  Up    3     15   3     30   System
-----
No. of Interfaces: 1
=====
```



```
*A:ALU_SIM11>show>router>ldp# interface detail
*A:ALU_SIM11>show>router>ldp#
=====
LDP Interfaces (Detail)
=====
Interface "a"
-----
Admin State      : Up                Oper State      : Up
Hold Time       : 15                Hello Factor   : 3
Keepalive Timeout : 30             Keepalive Factor : 3
Transport Addr  : System           Last Modified  : 07/06/2010 10:36:59
Active Adjacencies : 1
Tunneling       : Disabled
Lsp Name        : None
=====
*A:ALU_SIM11>show>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 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	Enabled — Authentication using MD5 message based digest protocol is enabled. Disabled — No authentication is used.
Passive-Mode	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.
Targeted-Sessions	true — Targeted sessions are enabled. false — Targeted sessions are disabled.

Sample Output

Show Commands

```
*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.

Label	Description (Continued)
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.

Sample Output

```
*A:SRU4>config>router>ldp# show router ldp peer
=====
LDP Peers
=====
```

Peer	Adm	Opr	Hello Factor	Hold Time	KA Factor	KA Timeout	Passive Mode	Auto 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

Show Commands

```
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
BFDD 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
BFDD 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
BFDD 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-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.

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

```
*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

```
-----
No. of Sessions: 1
=====
```

```
*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 that have been sent to the peer associated with this FEC.
Serv FECs Recv	The number of labels that have been received from the peer associated with this FEC.

Label	Description (Continued)
Attempted Sessions	The total number of attempted sessions for this LDP instance.
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.

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
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.

Parameters *ip-address* — The IP address of the LDP peer.

event

Syntax [no] event

Context 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>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>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.

Show Commands

hello

Syntax	hello [detail] no hello
Context	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.

Standards and Protocol Support

Standards Compliance

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 802.3 10BaseT
IEEE 802.3ad Link Aggregation
IEEE 802.3ae 10Gbps Ethernet
IEEE 802.3ah Ethernet OAM
IEEE 802.3u 100BaseTX
IEEE 802.3z 1000BaseSX/LX ITU-T Y.1731 OAM functions and mechanisms for Ethernet based networks draft-ietf-disman-alarm-mib-04.txt IANA-IFType-MIB
IEEE8023-LAG-MIB ITU-T G.8032 Ethernet Ring Protection Switching (version 2)

Protocol Support

BGP

RFC 1397 BGP Default Route Advertisement
RFC 1772 Application of BGP in the Internet
RFC 1997 BGP Communities Attribute
RFC 2385 Protection of BGP Sessions via MD5
RFC 2439 BGP Route Flap Dampening
RFC 2547 bis BGP/MPLS VPNs draft-ietf-idr-rfc2858bis-09.txt.
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 4760 Multi-protocol Extensions for BGP
RFC 4893 BGP Support for Four-octet AS Number Space

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)
RFC 5287 Control Protocol Extensions for the Setup of Time-Division Multiplexing (TDM) Pseudowires in MPLS Networks

DHCP

RFC 2131 Dynamic Host Configuration Protocol (REV)

DIFFERENTIATED SERVICES

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 2697 A Single Rate Three Color Marker
RFC 2698 A Two Rate Three Color Marker
RFC 4115 A Differentiated Service Two-Rate, Three-Color Marker with Efficient Handling of in-Profile Traffic

IPv6

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 2740 OSPF for IPv6
RFC 3587 IPv6 Global Unicast Address Format
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 5095 Deprecation of Type 0 Routing Headers in IPv6
draft-ietf-isis-ipv6-05
draft-ietf-isis-wg-multi-topology-xx.txt

IS-IS

RFC 1142 OSI IS-IS Intra-domain Routing Protocol (ISO 10589)
RFC 1195 Use of OSI IS-IS for routing in TCP/IP & dual environments
RFC 2763 Dynamic Hostname Exchange for IS-IS
RFC 2966 Domain-wide Prefix Distribution with Two-Level IS-IS
RFC 2973 IS-IS Mesh Groups
RFC 3373 Three-Way Handshake for Intermediate System to Intermediate System (IS-IS) Point-to-Point Adjacencies
RFC 3567 Intermediate System to Intermediate System (ISIS) Cryptographic Authentication
RFC 3719 Recommendations for Interoperable Networks using IS-IS
RFC 3784 Intermediate System to Intermediate System (IS-IS) Extensions for Traffic Engineering (TE)
RFC 3787 Recommendations for Interoperable IP Networks
RFC 3847 Restart Signaling for IS-IS – GR helper

Standards and Protocols

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

MPLS - General

RFC 3031 MPLS Architecture
RFC 3032 MPLS Label Stack Encoding
RFC 4379 Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures
RFC 4182 Removing a Restriction on the use of MPLS Explicit NULL

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) [Only in 7210 SAS-M access-uplink mode]

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 1907 SNMPv2-MIB
RFC 2011 IP-MIB
RFC 2012 TCP-MIB
RFC 2013 UDP-MIB
RFC 2096 IP-FORWARD-MIB
RFC 2138 RADIUS
RFC 2206 RSVP-MIB
RFC 2571 SNMP-FRAMEWORKMIB
RFC 2572 SNMP-MPD-MIB

RFC 2573 SNMP-TARGET-&-NOTIFICATION-MIB
RFC 2574 SNMP-USER-BASEDSMMIB
RFC 2575 SNMP-VIEW-BASEDACM-MIB
RFC 2576 SNMP-COMMUNITY-MIB
RFC 2665 EtherLike-MIB
RFC 2819 RMON-MIB
RFC 2863 IF-MIB
RFC 2864 INVERTED-STACK-MIB
RFC 3014 NOTIFICATION-LOGMIB
RFC 3164 Syslog
RFC 3273 HCRMON-MI
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
draft-ietf-mpls-lsr-mib-06.txt
draft-ietf-mpls-te-mib-04.txt
draft-ietf-mpls-ldp-mib-07.txt

OSPF

RFC 1765 OSPF Database Overflow
RFC 2328 OSPF Version 2
RFC 2370 Opaque LSA Support
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

MPLS - RSVP-TE

RFC 2430 A Provider Architecture DiffServ & TE
RFC 2702 Requirements for Traffic Engineering over MPLS
RFC2747 RSVP Cryptographic Authentication

RFC3097 RSVP Cryptographic Authentication
RFC 3209 Extensions to RSVP for Tunnels
RFC 4090 Fast reroute Extensions to RSVP-TE for LSP Tunnels
RFC 5817 Graceful Shutdown in MPLS and GMPLS Traffic Engineering Networks

PSEUDO-WIRE

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 4448 Encapsulation Methods for Transport of Ethernet over MPLS Networks (draft-ietf-pwe3-ethernet-encap-11.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
RFC6073, Segmented Pseudowire (draft-ietf-pwe3-segmented-pw-18.txt)
draft-ietf-l2vpn-vpws-iw-oam-02.txt
OAM Procedures for VPWS Interworking
draft-ietf-pwe3-oam-msg-map-14.txt, Pseudowire (PW) OAM Message Mapping
Pseudowire Preferential Forwarding Status bit definition
draft-pwe3-redundancy-02.txt
Pseudowire (PW) Redundancy

RADIUS

RFC 2865 Remote Authentication Dial In User Service
RFC 2866 RADIUS Accounting

SSH

draft-ietf-secsh-architecture.txt SSH Protocol Architecture
 draft-ietf-secsh-userauth.txt SSH Authentication Protocol
 draft-ietf-secsh-transport.txt SSH Transport Layer Protocol
 draft-ietf-secsh-connection.txt SSH Connection Protocol
 draft-ietf-secsh-newmodes.txt SSH Transport Layer Encryption Modes

TACACS+

draft-grant-tacacs-02.txt

TCP/IP

RFC 768 UDP
 RFC 1350 The TFTP Protocol
 RFC 791 IP
 RFC 792 ICMP
 RFC 793 TCP
 RFC 826 ARP
 RFC 854 Telnet
 RFC 1519 CIDR
 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

Timing

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.

ITU-T G.8264 Telecommunication Standardization Section of ITU, Distribution of timing information through packet networks, issued 10/2008.

IEEE Std 1588™-2008, IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems.

VPLS

RFC 4762 Virtual Private LAN Services Using LDP (previously draft-ietf-l2vpn-vpls-ldp-08.txt)

VRRP

RFC 2787 Definitions of Managed Objects for the Virtual Router Redundancy Protocol
 RFC 3768 Virtual Router Redundancy Protocol

Proprietary MIBs

ALCATEL-IGMP-SNOOPING-MIB.mib
 TIMETRA-CAPABILITY-7210-SAS-M-V5v0.mib
 (7210 SAS-M Only)
 TIMETRA-CAPABILITY-7210-SAS-X-V5v0.mib (7210 SAS-X Only)
 TIMETRA-CHASSIS-MIB.mib
 TIMETRA-CLEAR-MIB.mib
 TIMETRA-DOT3-OAM-MIB.mib
 TIMETRA-FILTER-MIB.mib
 TIMETRA-GLOBAL-MIB.mib
 TIMETRA-IEEE8021-CFM-MIB.mib
 TIMETRA-LAG-MIB.mib
 TIMETRA-LOG-MIB.mib
 TIMETRA-MIRROR-MIB.mib
 TIMETRA-NTP-MIB.mib
 TIMETRA-OAM-TEST-MIB.mib
 TIMETRA-PORT-MIB.mib
 TIMETRA-QOS-MIB.mib
 TIMETRA-SAS-ALARM-INPUT-MIB.mib
 TIMETRA-SAS-FILTER-MIB.mib
 TIMETRA-SAS-IEEE8021-CFM-MIB.mib
 TIMETRA-SAS-IEEE8021-PAE-MIB.mib
 TIMETRA-SAS-GLOBAL-MIB.mib
 TIMETRA-SAS-LOG-MIB.mib.mib

TIMETRA-SAS-MIRROR-MIB.mib
 TIMETRA-SAS-MPOINT-MGMT-MIB.mib (Only for 7210 SAS-X)
 TIMETRA-SAS-PORT-MIB.mib
 TIMETRA-SAS-QOS-MIB.mib
 TIMETRA-SAS-SDP-MIB.mib
 TIMETRA-SAS-SYSTEM-MIB.mib
 TIMETRA-SAS-SERV-MIB.mib
 TIMETRA-SAS-VRTR-MIB.mib
 TIMETRA-SCHEDULER-MIB.mib
 TIMETRA-SECURITY-MIB.mib
 TIMETRA-SERV-MIB.mib
 TIMETRA-SYSTEM-MIB.mib
 TIMETRA-TC-MIB.mib
 TIMETRA-ISIS-MIB.mib
 TIMETRA-ROUTE-POLICY-MIB.mib
 TIMETRA-MPLS-MIB.mib
 TIMETRA-RSVP-MIB.mib
 TIMETRA-LDP-MIB.mib
 TIMETRA-VRRP-MIB.mib
 TIMETRA-VRTR-MIB.mib

INDEX

L

LDP

- overview 154
 - architecture 155
 - execution flow 158
 - label exchange 159
 - LDP and MPLS 154
 - subsystems 156
 - label manager 157
 - LDP configuration 157
 - logger 157
 - memory manager 157
 - service manager 157
- configuring 172
 - basic 173
 - command reference 187
 - enabling 174
 - interface 180
 - LDP services signaling 182
 - management tasks 183
 - peer 181
 - targeted sessions 178

M

MPLS

- overview 18
 - label stack 19, 22, 21
 - LSP types 22
 - reservation types 33
 - services 40, 45
 - traffic engineering 36
- configuring
 - admin-group 48
 - basic 46
 - command reference 65, 66
 - interface 49
 - LSPs 44, 51
 - management tasks 58
 - overview 44
 - paths 44, 50
 - router interface 45

signaling protocol 45, 52

R

RSVP

- overview 29
 - for MPLS 31
 - hellos 31
 - message pacing 33
 - traffic engineering 31
- configuring
 - command reference 69
 - message pacing 56, 63
 - RSVP parameters 55

T

T-LDP 162

