



Alcatel-Lucent 7705

SERVICE AGGREGATION ROUTER OS | RELEASE 4.0
MPLS GUIDE

Alcatel-Lucent assumes no responsibility for the accuracy of the information presented, which is subject to change without notice.

Alcatel, Lucent, Alcatel-Lucent and the Alcatel-Lucent logo are trademarks of Alcatel-Lucent. All other trademarks are the property of their respective owners.

Copyright 2010 Alcatel-Lucent.
All rights reserved.

Disclaimers

Alcatel-Lucent products are intended for commercial uses. Without the appropriate network design engineering, they must not be sold, licensed or otherwise distributed for use in any hazardous environments requiring fail-safe performance, such as in the operation of nuclear facilities, aircraft navigation or communication systems, air traffic control, direct life-support machines, or weapons systems, in which the failure of products could lead directly to death, personal injury, or severe physical or environmental damage. The customer hereby agrees that the use, sale, license or other distribution of the products for any such application without the prior written consent of Alcatel-Lucent, shall be at the customer's sole risk. The customer hereby agrees to defend and hold Alcatel-Lucent harmless from any claims for loss, cost, damage, expense or liability that may arise out of or in connection with the use, sale, license or other distribution of the products in such applications.

This document may contain information regarding the use and installation of non-Alcatel-Lucent products. Please note that this information is provided as a courtesy to assist you. While Alcatel-Lucent tries to ensure that this information accurately reflects information provided by the supplier, please refer to the materials provided with any non-Alcatel-Lucent product and contact the supplier for confirmation. Alcatel-Lucent assumes no responsibility or liability for incorrect or incomplete information provided about non-Alcatel-Lucent products.

However, this does not constitute a representation or warranty. The warranties provided for Alcatel-Lucent products, if any, are set forth in contractual documentation entered into by Alcatel-Lucent and its customers.

This document was originally written in English. If there is any conflict or inconsistency between the English version and any other version of a document, the English version shall prevail.

Table of Contents

Preface	27
Getting Started	31
Alcatel-Lucent 7705 SAR MPLS Configuration Process	31
Notes on 7705 SAR-8, 7705 SAR-18, and 7705 SAR-F	32
MPLS and RSVP-TE	35
Overview	37
MPLS	38
Traffic Engineering for MPLS	38
TE Metric and IGP Metric	39
MPLS Label Stack	39
Label Values	41
Label Edge and Label Switch Routers	43
LSP Types	43
RSVP and RSVP-TE	45
RSVP-TE Overview	45
Using RSVP-TE for MPLS	47
RSVP-TE Extensions for MPLS	47
Hello Protocol	48
MD5 Authentication of RSVP-TE Interface	49
RSVP-TE Signaling	50
General Attributes of RSVP-TE	50
Authentication	51
OAM: BFD	51
Timers	51
LSP Resignal Limit	52
RSVP-TE Message Pacing	52
RSVP-TE Overhead Refresh Reduction	53
RSVP-TE Reservation Styles	54
LSP Redundancy	55
Fast Reroute (FRR)	56
FRR Terminology	57
FRR Behavior	59
Dynamic and Manual Bypass LSPs	60
Bypass LSP Selection Rules for the PLR	60
FRR Node Protection (Facility Backup)	63
Shared Risk Link Groups	65
SRLGs for Secondary LSP Paths	65
SRLGs for FRR LSP Paths	66
Disjoint and Non-disjoint Paths	66
Enabling Disjoint Backup Paths	67
RSVP-TE Graceful Shutdown	70
MPLS Service Usage	71
Service Destination Points	71
MPLS and RSVP-TE Configuration Process Overview	72

Table of Contents

Configuration Notes	73
Reference Sources	73
Configuring MPLS and RSVP-TE with CLI	75
MPLS Configuration Overview	76
Router Interface	76
E-LSP for Differentiated Services	76
Paths	76
LSPs	77
Pseudowires	77
Signaling Protocol	78
Basic MPLS Configuration	79
Common Configuration Tasks	80
Configuring MPLS Components	80
Configuring Global MPLS Parameters	81
Configuring an MPLS Interface	82
Configuring MPLS Paths	83
Configuring an MPLS LSP	84
Configuring a Static LSP	85
Configuring Manual Bypass Tunnels	86
Configuring RSVP-TE Parameters	88
Configuring RSVP-TE Message Pacing Parameters	89
MPLS Configuration Management Tasks	90
Deleting MPLS	90
Modifying MPLS Parameters	90
Modifying an MPLS LSP	91
Modifying MPLS Path Parameters	91
Modifying MPLS Static LSP Parameters	92
Deleting an MPLS Interface	93
RSVP-TE Configuration Management Tasks	94
Modifying RSVP-TE Parameters	94
Modifying RSVP-TE Message Pacing Parameters	95
Deleting an Interface from RSVP-TE	95
MPLS and RSVP-TE Command Reference	97
Command Hierarchies	97
Command Descriptions	104
Configuration Commands (MPLS)	105
Configuration Commands (RSVP-TE)	136
Show Commands (MPLS)	149
Show Commands (RSVP)	172
Clear Commands	182
Debug Commands	184
Label Distribution Protocol	195
Label Distribution Protocol	196
LDP and MPLS	196
LDP Architecture	197
LDP Subsystem Interrelationships	198
Memory Manager and LDP	198
Label Manager	198
LDP Configuration	199

Logger	200
Service Manager	200
Execution Flow	200
Initialization	200
Session Lifetime	200
Label Exchange	202
Other Reasons for Label Actions	202
Cleanup	202
LDP Filters	202
ECMP Support for LDP	203
Label Operations	204
Graceful Restart Helper	205
LDP Process Overview	206
Configuration Notes	207
Reference Sources	207
Configuring LDP with CLI	209
LDP Configuration Overview	210
Basic LDP Configuration	211
Common Configuration Tasks	212
Enabling LDP	212
Configuring Graceful Restart Helper Parameters	213
Applying Import Policies	214
Configuring Interface Parameters	215
Specifying Targeted Session Parameters	216
Specifying Peer Parameters	217
Enabling LDP Signaling and Services	218
LDP Configuration Management Tasks	219
Disabling LDP	219
Modifying Targeted Session Parameters	219
Modifying Interface Parameters	220
LDP Command Reference	221
Command Hierarchies	221
Command Descriptions	225
Configuration Commands	226
Show Commands	239
Clear Commands	262
Debug Commands	264
Standards and Protocol Support	269

List of Tables

Getting Started	31
Table 1: Configuration Process	31
Table 2: 7705 SAR-8, 7705 SAR-18, and 7705 SAR-F Comparison	32
MPLS and RSVP-TE	35
Table 3: Packet/Label Field Description	40
Table 4: Ingress Label Values (Pop Labels)	42
Table 5: Egress Label Values (Push Labels)	42
Table 6: FRR Terminology	57
Table 7: Disabled and Enabled Options for Bypass-Only	87
Table 8: Show Router MPLS Admin-Group Output Fields	149
Table 9: Show Router MPLS Bypass-Tunnel Output Fields	150
Table 10: Show Router MPLS Interface Output Fields	152
Table 11: Show Router MPLS Label Output Fields	154
Table 12: Show Router MPLS Label Range Output Fields	154
Table 13: Show Router MPLS LSP Output Fields	157
Table 14: Show Router MPLS LSP Detail Output Fields	158
Table 15: Show Router MPLS LSP Path Detail Output Fields	161
Table 16: Show Router MPLS LSP Path MBB Output Fields	164
Table 17: Show Router MPLS Path Output Fields	166
Table 18: Show Router MPLS SRLG Group Output Fields	167
Table 19: Show Router MPLS Static LSP Output Fields	169
Table 20: Show Router MPLS Status Output Fields	170
Table 21: Show Router RSVP-TE Interface Output Fields	173
Table 22: Show Router RSVP-TE Interface Detail Output Fields	174
Table 23: Show Router RSVP-TE Interface Statistics Output Fields	176
Table 24: Show Router RSVP-TE Neighbor Output Fields	177
Table 25: Show Router RSVP-TE Session Output Fields	179
Table 26: Show Router RSVP-TE Statistics Output Fields	180
Table 27: Show Router RSVP-TE Status Output Fields	181
Label Distribution Protocol	195
Table 28: Hello Timeout Factor Default Values	230
Table 29: Keepalive Timeout Factor Default Values	232
Table 30: LDP Bindings Output Fields	242
Table 31: LDP Discovery Output Fields	248
Table 32: LDP Interface Output Fields	250
Table 33: LDP Parameters Output Fields	252

List of Tables

Table 34:	LDP Peer Output Fields	255
Table 35:	LDP Peer-Parameter Output Fields	257
Table 36:	LDP Session Output Fields	258
Table 37:	LDP Status Output Fields	260

List of Figures

MPLS and RSVP-TE	35
Figure 1: Label Structure	40
Figure 2: Label Packet Placement	40
Figure 3: Establishing LSPs	46
Figure 4: LSP Using RSVP-TE Path Setup	46
Figure 5: Bypass Tunnel Node Example	60
Figure 6: FRR Node-Protection Example	63
Figure 7: Disjoint Primary and Secondary LSPs	68
Figure 8: Disjoint FRR Bypass LSPs	69
Figure 9: MPLS and RSVP-TE Configuration and Implementation Flow	72
Figure 10: Manual Bypass Tunnels	86
Label Distribution Protocol	195
Figure 11: LDP Subsystem Interrelationships	199
Figure 12: LDP Configuration and Implementation	206

List of Figures

List of Acronyms

Acronym	Expansion
2G	second generation wireless telephone technology
3DES	triple DES (data encryption standard)
3G	third generation mobile telephone technology
5620 SAM	5620 Service Aware Manager
7705 SAR	7705 Service Aggregation Router
7710 SR	7710 Service Router
7750 SR	7750 Service Router
9500 MPR	9500 Microwave Packet Radio
ABR	available bit rate area border router
AC	alternating current attachment circuit
ACK	acknowledge
ACL	access control list
ACR	adaptive clock recovery
ADP	automatic discovery protocol
AFI	authority and format identifier
AIS	alarm indication signal
ANSI	American National Standards Institute
Apipe	ATM VLL
APS	automatic protection switching
ARP	address resolution protocol
A/S	active/standby
AS	autonomous system

Acronym	Expansion
ASAP	any service, any port
ASBR	autonomous system boundary router
ASN	autonomous system number
ATM	asynchronous transfer mode
ATM PVC	ATM permanent virtual circuit
B3ZS	bipolar with three-zero substitution
Batt A	battery A
B-bit	beginning bit (first packet of a fragment)
Bellcore	Bell Communications Research
BFD	bidirectional forwarding detection
BGP	border gateway protocol
BITS	building integrated timing supply
BMCA	best master clock algorithm
BMU	<p>broadcast, multicast, and unknown traffic</p> <p>Traffic that is not unicast. Any nature of multipoint traffic:</p> <ul style="list-style-type: none"> • broadcast (that is, all 1s as the destination IP to represent all destinations within the subnet) • multicast (that is, traffic typically identified by the destination address, uses special destination address); for IP, the destination must be 224.0.0.0 to 239.255.255.255 • unknown (that is, the destination is typically a valid unicast address but the destination port/interface is not yet known; therefore, traffic needs to be forwarded to all destinations; unknown traffic is treated as broadcast)
BOF	boot options file
BPDU	bridge protocol data unit
BRAS	Broadband Remote Access Server
BSC	Base Station Controller
BSTA	Broadband Service Termination Architecture

Acronym	Expansion
BTS	base transceiver station
CAS	channel associated signaling
CBN	common bonding networks
CBS	committed buffer space
CC	control channel continuity check
CCM	continuity check message
CE	customer edge circuit emulation
CEM	circuit emulation
CES	circuit emulation services
CESoPSN	circuit emulation services over packet switched network
CFM	connectivity fault management
CIDR	classless inter-domain routing
CIR	committed information rate
CLI	command line interface
CLP	cell loss priority
CoS	class of service
CPE	customer premises equipment
Cpipe	circuit emulation (or TDM) VLL
CPM	Control and Processing Module (CPM is used instead of CSM when referring to CSM filtering to align with CLI syntax used with other SR products). CSM management ports are referred to as CPM management ports in the CLI.
CPU	central processing unit
CRC	cyclic redundancy check
CRON	a time-based scheduling service (from chronos = time)

Acronym	Expansion
CSM	Control and Switching Module
CSNP	complete sequence number PDU
CSPF	constrained shortest path first
C-TAG	customer VLAN tag
CV	connection verification customer VLAN (tag)
CW	control word
DC	direct current
DC-C	DC return - common
DCE	data communications equipment
DC-I	DC return - isolated
DCO	digitally controlled oscillator
DDoS	distributed DoS
DES	data encryption standard
DF	do not fragment
DHB	decimal, hexadecimal, or binary
DHCP	dynamic host configuration protocol
DHCPv6	dynamic host configuration protocol for IPv6
DIS	designated intermediate system
DM	delay measurement
DNS	domain name server
DoS	denial of service
dot1p	IEEE 802.1p bits, found in Ethernet or VLAN ingress packet headers and used to map traffic to up to eight forwarding classes
dot1q	IEEE 802.1q encapsulation for Ethernet interfaces
DPI	deep packet inspection

Acronym	Expansion
DPLL	digital phase locked loop
DSCP	differentiated services code point
DSL	digital subscriber line
DSLAM	digital subscriber line access multiplexer
DTE	data termination equipment
DU	downstream unsolicited
DUID	DHCP unique identifier
DV	delay variation
e911	enhanced 911 service
EAP	Extensible Authentication Protocol
EAPOL	EAP over LAN
E-bit	ending bit (last packet of a fragment)
ECMP	equal cost multi-path
EFM	Ethernet in the first mile
EGP	exterior gateway protocol
EIA/TIA-232	Electronic Industries Alliance/Telecommunications Industry Association Standard 232 (also known as RS-232)
ELER	egress label edge router
E&M	ear and mouth earth and magneto exchange and multiplexer
Epipe	Ethernet VLL
EPL	Ethernet private line
ERO	explicit route object
ESD	electrostatic discharge
ESMC	Ethernet synchronization message channel
ETE	end-to-end

Acronym	Expansion
ETH-CFM	Ethernet connectivity fault management (IEEE 802.1ag)
EVDO	evolution - data optimized
EVPL	Ethernet virtual private link
EXP bits	experimental bits (currently known as TC)
FC	forwarding class
FCS	frame check sequence
FDB	forwarding database
FDL	facilities data link
FEAC	far-end alarm and control
FEC	forwarding equivalence class
FF	fixed filter
FIB	forwarding information base
FIFO	first in, first out
FNG	fault notification generator
FOM	figure of merit
FRR	fast reroute
FTN	FEC-to-NHLFE
FTP	file transfer protocol
GFP	generic framing procedure
GigE	Gigabit Ethernet
GRE	generic routing encapsulation
GSM	Global System for Mobile Communications (2G)
HCM	high capacity multiplexing
HDB3	high density bipolar of order 3
HEC	header error control
HMAC	hash message authentication code

Acronym	Expansion
HSDPA	high-speed downlink packet access
HSPA	high-speed packet access
HVPLS	hierarchical virtual private line service
IANA	internet assigned numbers authority
IBN	isolated bonding networks
ICMP	Internet control message protocol
ICMPv6	Internet control message protocol for IPv6
ICP	IMA control protocol cells
IEEE	Institute of Electrical and Electronics Engineers
IEEE 1588v2	Institute of Electrical and Electronics Engineers standard 1588-2008
IES	Internet Enhanced Service
IETF	Internet Engineering Task Force
IGP	interior gateway protocol
ILER	ingress label edge router
ILM	incoming label map
IMA	inverse multiplexing over ATM
IOM	input/output module
IP	Internet Protocol
IPCP	Internet Protocol Control Protocol
IPIP	IP in IP
Ipipe	IP interworking VLL
IPoATM	IP over ATM
IS-IS	Intermediate System-to-Intermediate System
IS-IS-TE	IS-IS-traffic engineering (extensions)
ISO	International Organization for Standardization

Acronym	Expansion
LB	loopback
lbf-in	pound force inch
LBM	loopback message
LBO	line buildout
LBR	loopback reply
LCP	link control protocol
LDP	label distribution protocol
LER	label edge router
LFIB	label forwarding information base
LIB	label information base
LLDP	link layer discovery protocol
LLDPDU	link layer discovery protocol data unit
LLF	link loss forwarding
LLID	loopback location ID
LM	loss measurement
LSA	link-state advertisement
LSDB	link-state database
LSP	label switched path link-state PDU (for IS-IS)
LSR	label switch router link-state request
LSU	link-state update
LT	linktrace
LTE	line termination equipment
LTM	linktrace message
LTN	LSP ID to NHLFE

Acronym	Expansion
LTR	linktrace reply
MA	maintenance association
MAC	media access control
MA-ID	maintenance association identifier
MBB	make-before-break
MBS	maximum buffer space maximum burst size media buffer space
MBSP	mobile backhaul service provider
MC-MLPPP	multi-class multilink point-to-point protocol
MD	maintenance domain
MD5	message digest version 5 (algorithm)
MDA	media dependent adapter
MDDDB	multidrop data bridge
MDL	maintenance data link
ME	maintenance entity
MED	multi-exit discriminator
MEF	Metro Ethernet Forum
MEG	maintenance entity group
MEG-ID	maintenance entity group identifier
MEN	Metro Ethernet network
MEP	maintenance association end point
MFC	multi-field classification
MHF	MIP half function
MIB	management information base
MIP	maintenance association intermediate point

Acronym	Expansion
MIR	minimum information rate
MLPPP	multilink point-to-point protocol
MP	merge point multilink protocol
MP-BGP	multiprotocol border gateway protocol
MPLS	multiprotocol label switching
MPR	see 9500 MPR
MRRU	maximum received reconstructed unit
MRU	maximum receive unit
MSDU	MAC Service Data Unit
MS-PW	multi-segment pseudowire
MTIE	maximum time interval error
MTSO	mobile trunk switching office
MTU	maximum transmission unit multi-tenant unit
M-VPLS	management virtual private line service
MW	microwave
N·m	newton meter
NBMA	non-broadcast multiple access (network)
NE	network element
NET	network entity title
NHLFE	next hop label forwarding entry
NHOP	next-hop
NLRI	network layer reachability information
NNHOP	next next-hop
NNI	network-to-network interface

Acronym	Expansion
Node B	similar to BTS but used in 3G networks — term is used in UMTS (3G systems) while BTS is used in GSM (2G systems)
NSAP	network service access point
NSSA	not-so-stubby area
NTP	network time protocol
OAM	operations, administration, and maintenance
OAMPDU	OAM protocol data units
OC3	optical carrier, level 3
ORF	outbound route filtering
OS	operating system
OSI	Open Systems Interconnection (reference model)
OSINLCP	OSI Network Layer Control Protocol
OSPF	Open Shortest Path First
OSPF-TE	OSPF-traffic engineering (extensions)
OSS	operations support system
OSSP	Organization Specific Slow Protocol
OTP	one time password
PADI	PPPoE active discovery initiation
PADR	PPPoE active discovery request
PAE	port authentication entities
PCP	priority point code
PDU	protocol data units
PDV	packet delay variation
PDVT	packet delay variation tolerance
PE	provider edge router
PHB	per-hop behavior

Acronym	Expansion
PHY	physical layer
PID	protocol ID
PIR	peak information rate
PLCP	Physical Layer Convergence Protocol
PLR	point of local repair
POP	point of presence
POS	packet over SONET
PPP	point-to-point protocol
PPPoE	point-to-point protocol over Ethernet
PRC	primary reference clock
PSN	packet switched network
PSNP	partial sequence number PDU
PTP	precision time protocol performance transparency protocol
PVC	permanent virtual circuit
PVCC	permanent virtual channel connection
PW	pseudowire
PWE	pseudowire emulation
PWE3	pseudowire emulation edge-to-edge
QL	quality level
QoS	quality of service
RADIUS	Remote Authentication Dial In User Service
RAN	Radio Access Network
RBS	robbed bit signaling
RD	route distinguisher
RDI	remote defect indication

Acronym	Expansion
RED	random early discard
RESV	reservation
RIB	routing information base
RJ-45	registered jack 45
RNC	Radio Network Controller
RRO	record route object
RS-232	Recommended Standard 232 (also known as EIA/TIA-232)
RSHG	residential split horizon group
RSTP	Rapid Spanning Tree Protocol
RSVP-TE	resource reservation protocol - traffic engineering
RT	receive/transmit
RTM	routing table manager
RTN	battery return
RTP	real-time protocol
R&TTE	Radio and Telecommunications Terminal Equipment
RTU	remote terminal unit
RU	rack unit
SAA	service assurance agent
SAP	service access point
SAR-8	7705 Service Aggregation Router - 8-slot chassis
SAR-18	7705 Service Aggregation Router - 18-slot chassis
SAR-F	7705 Service Aggregation Router - fixed form-factor chassis
SAToP	structure-agnostic TDM over packet
SCADA	surveillance, control and data acquisition
SCP	secure copy
SD	signal degrade

Acronym	Expansion
SDH	synchronous digital hierarchy
SDI	serial data interface
SDP	service destination point
SE	shared explicit
SF	signal fail
SFP	small form-factor pluggable (transceiver)
SGT	self-generated traffic
SHA-1	secure hash algorithm
SHG	split horizon group
SIR	sustained information rate
SLA	Service Level Agreement
SNMP	Simple Network Management Protocol
SNPA	subnetwork point of attachment
SNTP	simple network time protocol
SONET	synchronous optical networking
S-PE	switching provider edge router
SPF	shortest path first
SPT	shortest path tree
SR	service router (includes 7710 SR, 7750 SR)
SRLG	shared risk link group
SSH	secure shell
SSM	synchronization status messaging
SSU	system synchronization unit
S-TAG	service VLAN tag
STM1	synchronous transport module, level 1
SVC	switched virtual circuit

Acronym	Expansion
SYN	synchronize
TACACS+	Terminal Access Controller Access-Control System Plus
TC	traffic class (formerly known as EXP bits)
TCP	transmission control protocol
TDEV	time deviation
TDM	time division multiplexing
TE	traffic engineering
TFTP	trivial file transfer protocol
TLDP	targeted LDP
TLV	type length value
ToS	type of service
T-PE	terminating provider edge router
TPID	tag protocol identifier
TPMR	two-port MAC relay
TTL	time to live
TTM	tunnel table manager
U-APS	unidirectional automatic protection switching
UBR	unspecified bit rate
UDP	user datagram protocol
UMTS	Universal Mobile Telecommunications System (3G)
UNI	user-to-network interface
V.35	V-series Recommendation 35
VC	virtual circuit
VCC	virtual channel connection
VCCV	virtual circuit connectivity verification
VCI	virtual circuit identifier

Acronym	Expansion
VID	VLAN ID
VLAN	virtual LAN
VLL	virtual leased line
VoIP	voice over IP
Vp	peak voltage
VP	virtual path
VPC	virtual path connection
VPI	virtual path identifier
VPLS	virtual private LAN service
VPN	virtual private network
VPRN	virtual private routed network
VRF	virtual routing and forwarding table
VSE	vendor-specific extension
VSO	vendor-specific option
WCDMA	wideband code division multiple access (transmission protocol used in UMTS networks)
WRED	weighted random early discard
WTR	wait to restore

About This Guide

This guide describes the services and protocol support provided by the Alcatel-Lucent 7705 Service Aggregation Router and presents examples to configure and implement MPLS and LDP protocols.

This guide 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 guide is intended for network administrators who are responsible for configuring the 7705 SAR routers. It is assumed that the network administrators have an understanding of networking principles and configurations. Protocols and concepts described in this guide include the following:

- Multiprotocol Label Switching (MPLS)
- Resource Reservation Protocol for Traffic Engineering (RSVP-TE)
- Label Distribution Protocol (LDP)

List of Technical Publications

The 7705 SAR OS documentation set is composed of the following guides:

- 7705 SAR OS Basic System Configuration Guide
This guide describes basic system configurations and operations.
- 7705 SAR OS System Management Guide
This guide describes system security and access configurations as well as event logging and accounting logs.
- 7705 SAR OS Interface Configuration Guide
This guide describes card and port provisioning.
- 7705 SAR OS Router Configuration Guide
This guide describes logical IP routing interfaces, IP-based filtering, and routing policies.
- 7705 SAR OS MPLS Guide
This guide describes how to configure Multiprotocol Label Switching (MPLS), Resource Reservation Protocol for Traffic Engineering (RSVP-TE), and Label Distribution Protocol (LDP).
- 7705 SAR OS Services Guide
This guide describes how to configure service parameters such as service access points (SAPs), service destination points (SDPs), customer information, and user services.
- 7705 SAR OS Quality of Service Guide
This guide describes how to configure Quality of Service (QoS) policy management.
- 7705 SAR OS Routing Protocols Guide
This guide provides an overview of dynamic routing concepts and describes how to configure them.
- 7705 SAR OS OAM and Diagnostics Guide
This guide provides information on Operations, Administration and Maintenance (OAM) tools.

Multiple PDF File Search

You can use Adobe Reader, Release 6.0 or later, to search multiple PDF files for a term. Adobe Reader displays the results in a display panel. The results are grouped by PDF file. You can expand the entry for each file.



Note: The PDF files in which you search must be in the same folder.

To search multiple PDF files for a term:

Step 1. Open Adobe Reader.

Step 2. Choose Edit – Search from the Adobe Reader main menu. The Search panel appears.

Step 3. Enter the term to search for.

Step 4. Select the All PDF Documents in radio button.

Step 5. Choose the folder in which to search using the drop-down menu.

Step 6. Select the following criteria if required:

- Whole words only
- Case-Sensitive
- Include Bookmarks
- Include Comments

Step 7. Click on the Search button.

Adobe Reader displays the search results. You can expand the entries for each file by clicking on the + symbol.

Step 8. Click on a search result to go directly to that location in the selected file.

Technical Support

If you purchased a service agreement for your 7705 SAR 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, check this link for instructions to contact Support personnel:

Web: <http://support.alcatel-lucent.com>

Getting Started

In This Chapter

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

Alcatel-Lucent 7705 SAR MPLS Configuration Process

[Table 1](#) lists the tasks necessary to configure MPLS application 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	MLPS	MPLS on page 38
	RSVP-TE	RSVP and RSVP-TE on page 45
	LDP	Label Distribution Protocol on page 195
Reference	List of IEEE, IETF, and other proprietary entities	Standards and Protocol Support on page 269

Notes on 7705 SAR-8, 7705 SAR-18, and 7705 SAR-F

The 7705 SAR-8, 7705 SAR-18, and 7705 SAR-F run the same operating system software. The main difference between the products is their hardware platforms.

The 7705 SAR-8 is an 8-slot chassis that supports 2 CSMs, a Fan module, and 6 adapter cards. The 7705 SAR-18 chassis has 18 slots; in Release 4.0, it supports 2 CSMs, a Fan module, an Alarm module, and 12 adapter cards.

The 7705 SAR-F chassis has a fixed hardware configuration. The 7705 SAR-F replaces the CSM, Fan module, and the 16-port T1/E1 ASAP Adapter card and 8-port Ethernet Adapter card with an all-in-one unit that provides comparable functional blocks, as detailed in [Table 2](#).

The fixed configuration of the 7705 SAR-F means that provisioning the router at the “card slot” and “type” levels is preset and is not user-configurable. Operators begin configurations at the port level.



Note: Unless stated otherwise, references to the terms “Adapter card” and “CSM” throughout the 7705 SAR OS documentation set include the equivalent functional blocks on the 7705 SAR-F.

Table 2: 7705 SAR-8, 7705 SAR-18, and 7705 SAR-F Comparison

7705 SAR-8, 7705 SAR-18	7705 SAR-F	Notes
CSM	Control and switching functions	The control and switching functions include the console and management interfaces, the alarm and fan functions, the synchronization interfaces, system LEDs, and so on.
Fan module	Integrated with the control and switching functions	

Table 2: 7705 SAR-8, 7705 SAR-18, and 7705 SAR-F Comparison (Continued)

7705 SAR-8, 7705 SAR-18	7705 SAR-F	Notes
16-port T1/E1 ASAP Adapter card	16 individual T1/E1 ports on the faceplate	<p>The T1/E1 ports on the 7705 SAR-F are equivalent to the T1/E1 ports on the 16-port T1/E1 ASAP Adapter card, version 1, except that the 16 T1/E1 ports on the 7705 SAR-F support multiple synchronization sources to support two timing references. The 16-port T1/E1 ASAP Adapter card, version 2, also supports two timing references.</p> <p>On the 7705 SAR-8 and 7705 SAR-18, the CLI indicates the MDA type for the 16-port T1/E1 ASAP Adapter card as <code>a16-chds1</code> for version 1 and <code>a16-chds1v2</code> for version 2.</p> <p>On the 7705 SAR-F, the CLI indicates the MDA type for the 7705 SAR-F ports as <code>i16-chds1</code>.</p>
8-port Ethernet Adapter card	8 individual Ethernet ports on the faceplate	<p>The –48 VDC versions of the 7705 SAR-8 support two versions of the 8-port Ethernet Adapter card, with version 2 having additional support for Synchronous Ethernet. The +24 VDC version of the 7705 SAR-8 supports only version 2 of the 8-port Ethernet Adapter card.</p> <p>The 7705 SAR-18 supports only version 2 of the card.</p> <p>The Ethernet ports on the 7705 SAR-F are functionally equivalent to the Ethernet ports on version 2 of the 8-port Ethernet Adapter card and support multiple synchronization sources to support two timing references.</p> <p>On the 7705 SAR-8, the CLI indicates the MDA type for the 8-port Ethernet Adapter card as <code>a8-eth</code> or <code>a8-ethv2</code>. On the 7705 SAR-18, the CLI indicates the MDA type as <code>a8-ethv2</code>. On the 7705 SAR-F, the CLI indicates the MDA type for the 7705 SAR-F Ethernet ports as <code>i8-eth</code>.</p>
Requires user configuration at card (IOM) and MDA (adapter card) levels	Configuration at card (IOM) and MDA (adapter card) levels is preset and users cannot change these types	

MPLS and RSVP-TE

In This Chapter

This chapter provides information required to configure Multiprotocol Label Switching (MPLS) and Resource Reservation Protocol for Traffic Engineering (RSVP-TE) for the 7705 SAR. For information on dynamic LSPs with LDP, refer to the chapter [Label Distribution Protocol](#).

Topics in this chapter include:

- [Overview on page 37](#)
- [MPLS on page 38](#)
 - [Traffic Engineering for MPLS on page 38](#)
 - [MPLS Label Stack on page 39](#)
 - [Label Edge and Label Switch Routers on page 43](#)
 - [LSP Types on page 43](#)
- [RSVP and RSVP-TE on page 45](#)
 - [RSVP-TE Overview on page 45](#)
- [RSVP-TE Signaling on page 50](#)
 - [General Attributes of RSVP-TE on page 50](#)
- [LSP Redundancy on page 55](#)
- [Fast Reroute \(FRR\) on page 56](#)
 - [FRR Terminology on page 57](#)
 - [FRR Behavior on page 59](#)
 - [Dynamic and Manual Bypass LSPs on page 60](#)
- [Shared Risk Link Groups on page 65](#)
 - [SRLGs for Secondary LSP Paths on page 65](#)
 - [SRLGs for FRR LSP Paths on page 66](#)
 - [Disjoint and Non-disjoint Paths on page 66](#)
 - [Enabling Disjoint Backup Paths on page 67](#)
- [RSVP-TE Graceful Shutdown on page 70](#)

- [MPLS Service Usage on page 71](#)
 - [Service Destination Points on page 71](#)
- [MPLS and RSVP-TE Configuration Process Overview on page 72](#)
- [Configuration Notes on page 73](#)
 - [Reference Sources on page 73](#)
- [Configuring MPLS and RSVP-TE with CLI on page 75](#)
- [MPLS and RSVP-TE Command Reference on page 97](#)

Overview

The 7705 SAR provides MPLS technology using static LSPs, RSVP-TE for traffic-engineered signaled routing of LSPs, and LDP for non-traffic-engineered signaled routing of LSPs. A network operator may choose to use any combination of static LSPs, RSVP-TE, and LDP to establish paths for services. Furthermore, the 7705 SAR can be used as an ingress and egress Label Edge Router (ILER and ELER), and as a transit router. A transit router is also referred to as a Label Switch Router (LSR). Consider RSVP-TE and LDP as the Layer 2.5 protocols.

OSPF and IS-IS are the interior gateway protocols with traffic engineering extensions (IGP-TE) available to the 7705 SAR. These are the Layer 3 protocols. Typically, one or the other of these gateway protocols will be in use in the network. Whichever protocol is the chosen gateway protocol, it must be working in order for LDP or RSVP-TE to function. These Layer 3 protocols identify the next hop, which is information needed by the Layer 2.5 protocols (LDP or RSVP-TE) in order to assign labels.

In addition, the 7705 SAR provides link and node redundancy protection through LSP redundancy and Fast Reroute (FRR) features.

The LSP redundancy and FRR features have the ability to take shared risk link groups (SRLGs) into consideration when the Constrained Shortest Path First (CSPF) algorithm is used to determine an alternate LSP. The selection of a route is determined by the IGP-TE protocol. The added constraints imposed by SRLGs and CSPF will ensure that the redundant route selected will be unique from the principal route (route being protected); that is, it will use physical equipment that is different from the equipment that carries the principal route. CSPF will constrain the alternate route to be the shortest possible alternative route. Note that there may be more than one alternative route.

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 independent of any routing protocol but is considered multiprotocol because it works with protocols such as IP, ATM, Ethernet, and circuit emulation.

Traffic Engineering for MPLS

Without traffic engineering (TE), routers route traffic according to the Shortest Path First (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 less-congested path 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 (bidirectional) 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 (from next hop to next hop) 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 using SPF to reach the destination IP address.

TE Metric and IGP Metric

When the TE metric is selected for an LSP, the shortest path computation will select an LSP path based on the TE metric constraints instead of the IGP metric (for OSPF and IS-IS), which is the default metric. The user configures the TE metric under the `router>mpls>interface` context and the IGP metric under the `router>ospf>area>interface` context (for OSPF) and the `router>isis>interface>level` context (for IS-IS). 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 the IGP protocols. For more information on the OSPF and IS-IS routing protocols, refer to the 7705 SAR OS Routing Protocols Guide.

Typically, the TE metric is used to allow Constrained Shortest Path First (CSPF) to represent a dual TE topology for the purpose of computing LSP paths, where one TE topology is based on the RSVP-TE database and the other is based on the IGP-TE database.

An LSP dedicated to 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 amount of delay, or combined delay and jitter, 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 to 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 value as required.

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

MPLS Label Stack

Routers that support MPLS are known as Label Edge Routers (LERs) and Label Switch Routers (LSRs). MPLS requires a set of procedures to enhance network layer packets with label stacks, which turns them into labeled packets. In order to initiate, transmit, or terminate a labeled packet on a particular data link, an LER or 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 (that is, remove the top label), or swap the label and push one or more labels onto 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 other labels may have been above it in the past or that 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](#) shows the structure of a label and [Table 3](#) describes the fields. [Figure 2](#) shows the label placement in a packet.

Figure 1: Label Structure

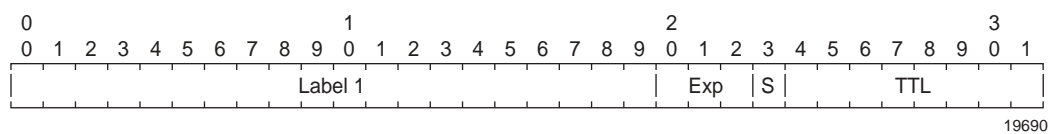
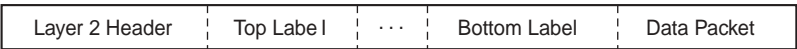


Table 3: Packet/Label Field Description

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 time-to-live 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.

Label Values

The 7705 SAR uses RSVP-TE and LDP protocols for label forwarding. For packet-based services such as VLL, the 7705 SAR uses T-LDP for signaling PW labels between peer nodes.

Packets traveling along an LSP are identified by the packet label, which is the 20-bit, unsigned integer (see [Label Edge and Label Switch Routers](#)). The range is 0 through 1 048 575. Label values 0 to 15 are reserved and are defined below:

- A value of 0 represents the IPv4 Explicit NULL label. This label value is legal only at the bottom of the label stack if the label stack is immediately followed by an IPv4 header, in which case the packet forwarding is based on the IPv4 header. If the IPv4 Explicit NULL label is not at the bottom of the label stack, then the packet forwarding is based on the subsequent label.
- 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 be 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. An LER advertises this when it is requesting penultimate hop popping and expecting unlabeled packets. Thus, the label value 3 should never appear in the label stack.
- Values 4 through 15 are reserved for future use.

[Table 4](#) lists the label ranges available for use by ingress labels (pop labels).

Table 4: Ingress Label Values (Pop Labels)

Label Values	Description
16 through 31	Reserved for future use
32 through 1023	Available for static outer LSP tunnel label assignment
1024 through 2047	Reserved for future use
2048 through 18 431	Statically assigned for services (inner pseudowire label)
32 768 through 131 071	Dynamically assigned for both MPLS and services
131 072 through 1 048 575	Reserved for future use

[Table 5](#) lists the label ranges available for use by egress labels (push labels).

Table 5: Egress Label Values (Push Labels)

Label Values	Description
16 through 1 048 575	Can be used for static LSP tunnel and static PW labels
16 through 1 048 575	Can be dynamically assigned for both MPLS tunnel labels and PW labels

Label Edge and Label Switch Routers

A 7705 SAR performs different functions based on its position in an LSP—ingress, egress, or transit—as described in the following list:

- ingress Label Edge Router (ILER) — The router at the beginning of an LSP is the ILER. The ingress router encapsulates packets with an MPLS header and forwards the packets to the next router along the path. An LSP can only have one ingress router.
- Label Switching Router (LSR) — An LSR can be any intermediate router in the LSP between the ingress and egress routers, swapping the incoming label with the outgoing MPLS label and forwarding the MPLS packets it receives to the next router in the LSP. An LSP can have 0 to 253 transit routers.
- egress Label Edge Router (ELER) — The router at the end of an LSP is the 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. An LSP can have only one egress router. The ingress and egress routers in an LSP cannot be the same router.

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

Constrained-path LSPs are signaled and are confined to one Interior Gateway Protocol (IGP) area. These LSPs cannot cross an autonomous system (AS) boundary.

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

LSP Types

The following LSP types are supported:

- static LSPs — a static LSP specifies a static path. All routers that the LSP traverses must be configured manually with labels. No RSVP-TE or LDP signaling is required. Static LSPs are discussed in this chapter.
- signaled LSPs — LSPs are set up using the RSVP-TE or LDP signaling protocol. 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. RSVP-TE is discussed in this chapter, and LDP is discussed in [Label Distribution Protocol](#).

There are two types of signaled LSP:

- 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 other routers (loose). Thus, you can control how the path is set up. Explicit-path LSPs are similar to static LSPs but require less configuration. See [RSVP and RSVP-TE](#). Note that an explicit path that has not specified any hops will follow the IGP route.
- constrained-path LSPs — for 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 that satisfies the constraints for the LSP. In turn, CSPF relies on the topology database provided by an extended IGP such as OSPF or IS-IS.

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

If Fast Reroute (FRR) is configured, the ingress router signals the downstream routers so that each downstream router can preconfigure a detour route for the LSP that will be used if there is a failure on the original LSP. If a downstream router does not support FRR, the request is ignored and the router continues to support the original LSP. This can cause some of the detour routes to fail, but the original LSP is not impacted. For more information on FRR, see [Fast Reroute \(FRR\)](#).

No bandwidth is reserved for the reroute 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 establishing the backup LSP. The following warning message is displayed:

“The fast reroute bandwidth command is not supported in this release.”

RSVP and RSVP-TE

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 operational state to provide the requested service. In general, RSVP requests result in resources reserved in each node along the data path.

The Resource Reservation Protocol for Traffic Engineering (RSVP-TE) is an extended version of RSVP for MPLS. RSVP-TE uses traffic engineering extensions to support automatic signaling of LSPs. MPLS uses RSVP-TE to set up traffic-engineered LSPs. See [RSVP-TE Extensions for MPLS](#) for more information.

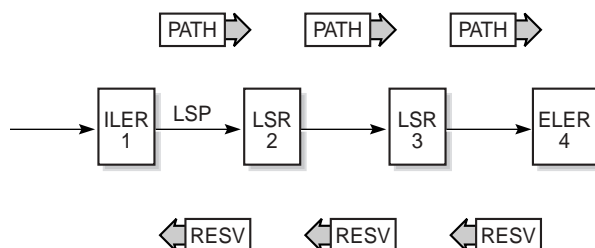
RSVP-TE Overview

RSVP-TE requests resources for simplex (unidirectional) flows. Therefore, RSVP-TE 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-TE is a signaling protocol, not a routing protocol. RSVP-TE operates with unicast and multicast routing protocols. Routing protocols determine where packets are forwarded. RSVP-TE consults local routing tables to relay RSVP-TE messages.

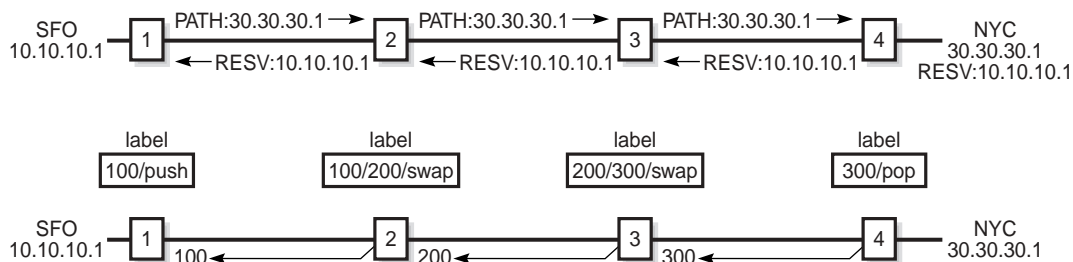
RSVP-TE uses two message types to set up LSPs, PATH and RESV. [Figure 3](#) depicts the process to establish an LSP.

- The sender (the ingress LER (ILER)) sends PATH messages toward the receiver, (the egress LER (ELER)) to indicate the forwarding equivalence class (FEC) for which label bindings are desired. PATH messages are used to signal and request the 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 3: Establishing LSPs

20120

Figure 4 displays an example of an LSP path set up using RSVP-TE. The ingress label edge router (ILER 1) transmits an RSVP-TE 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.

Figure 4: LSP Using RSVP-TE Path Setup

20121

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

- explicit route object (ERO) — when the ERO is present, the RSVP-TE 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
- session attribute object — controls the path setup priority, holding priority, and local rerouting features

Upon receiving a PATH message containing a label request object, the ELER transmits an 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-TE for MPLS

Hosts and routers that support both MPLS and RSVP-TE can associate labels with RSVP-TE flows. When MPLS and RSVP-TE 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 Forwarding Equivalence Class (FEC) that defines the RSVP-TE flow.

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

RSVP-TE Extensions for MPLS

The RSVP-TE extensions enable MPLS to support the creation of explicitly routed LSPs, with or without resource reservation. Several of the features enabled by these extensions were implemented to meet the requirements for traffic engineering over MPLS, which enables the creation of traffic trunks with specific characteristics. None of the TE extensions result in backward compatibility problems with traditional RSVP implementations.

To run properly, the traffic engineering capabilities of RSVP-TE require an underlying TE-enabled IGP routing protocol. The 7705 SAR supports OSPF and IS-IS with TE extensions.

Routing protocols make it possible to advertise the constraints imposed over various links in the network. For example, in order for the nodes in a network to choose the best link for signaling a tunnel, the capacity of a particular link and the amount of reservable capacity must be advertised by the IGP. RSVP-TE makes use of these constraints to request the setup of a path or LSP that traverses only those links that are part of an administrative group (admin groups are described in the following list). Thus, both RSVP-TE and the IGP-TE (that is, OSPF-TE or IS-IS-TE for the 7705 SAR) must be enabled and running simultaneously.

The following TE capabilities are supported:

- **hop limit** — the hop limit is the maximum number of LSR nodes that a given LSP can traverse, including the ingress and the egress LER nodes. Typically, the hop limit is used to control the maximum delay time for mission-critical traffic such as voice traffic.

The hop limit applies to the primary LSP, any backup LSPs, and LSPs configured to be used in Fast Reroute (FRR) situations.

- **admin groups** — administrative groups provide a way to define which LSR nodes should be included or excluded while signaling an LSP. For example, it might be desirable to avoid some nodes or links that are known to be utilized heavily from being included in the path of an LSP, or to include a specific LSR node to ensure that a newly signaled RSVP-TE tunnel traverses that LSR node.

Administrative groups apply to both primary and secondary LSPs. They are defined under the `config>router>mpls` context, and are applied at the MPLS interface level, as well as at the LSP and the primary and secondary LSP levels through `include` and `exclude` commands.

- **bandwidth** — the bandwidth capability (supported by RSVP-TE), is similar to the Connection Admission Control (CAC) function in ATM. During the establishment phase of RSVP-TE, the LSP PATH message contains the bandwidth reservation request. If the requested capacity is available, the RESV message confirms the reservation request. The amount of reserved bandwidth stated in the request is deducted from the amount of reservable bandwidth for each link over which the LSP traverses.

The bandwidth capability applies to both primary and secondary LSPs, and LSPs configured to be used in Fast Reroute (FRR) situations.

Hello Protocol

The Hello protocol detects the loss of a neighbor node (node failure detection) or the reset of a neighbor's RSVP-TE 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 LERs. 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-TE 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-TE Interface

When enabled on an RSVP-TE 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
- the key used with the authentication algorithm
- the lifetime of the key. A key is a user-generated key using third-party software or hardware. The value is entered as a static string into the CLI configuration of the RSVP interface. The key will continue to be valid until it is removed from that RSVP interface.
- the source address of the sending system
- the 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 that also contains a Flags field, a Key Identifier field, and a Sequence Number field. The RSVP sender complies with 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.

If a point of local repair (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 merge point (MP) node, then the message is discarded since there is no security association with the next-next-hop MP node.

The 7705 SAR MD5 implementation does not support the authentication challenge procedures in RFC 2747.

RSVP-TE Signaling

RSVP-TE-based signaling provides a means to establish tunnels dynamically.

RSVP-TE uses the Downstream on Demand (DOD) label distribution mode, sending PATH messages from the ingress LER node to the egress LER, and RESV messages in the reverse direction. DOD label distribution is a router's response to an explicit request from another router for label binding information. The DOD mode is in contrast to LDP on the 7705 SAR, which uses the Downstream Unsolicited (DU) label distribution mode for both PWs and LSPs. A router in DU mode will distribute label bindings to another router that has not explicitly requested the label bindings.

RSVP-TE signaling is supported when the 7705 SAR is deployed as an LER and as an LSR. When used as an LER, the 7705 SAR uses RSVP-TE signaling to set up constrained paths because only the LER knows all the constraints imposed on the LSP. When used as an LSR, the 7705 SAR uses RSVP-TE to interpret the RSVP-TE messages (including all the constraints).

With RSVP-TE, users can choose which services and PWs may use a particular LSP. One-to-one or many-to-one scenarios for binding PWs to RSVP-TE LSPs is supported, which is similar to binding PWs to static LSPs. Furthermore, each RSVP-TE LSP can be configured with its own set of attributes and constraints.

General Attributes of RSVP-TE

The following general attributes of RSVP-TE on the 7705 SAR are supported:

- [Authentication](#)
- [OAM: BFD](#)
- [Timers](#)
- [LSP Resignal Limit](#)
- [RSVP-TE Message Pacing](#)
- [RSVP-TE Overhead Refresh Reduction](#)
- [RSVP-TE Reservation Styles](#)

Authentication

In order to ensure the integrity of a peer router, authentication for RSVP-TE is supported. It can be enabled on a per-link basis and is bidirectional. Hence, both of the nodes must either enable authentication or disable it on a per-peer or per-link basis. The MD5-based authentication algorithm is implemented and sequence numbers are used to keep track of messages.

OAM: BFD

Bidirectional Forwarding Detection (BFD) is supported on the 7705 SAR. In the case of BFD for RSVP-TE, an RSVP-TE enabled link is registered with the BFD state machine, and if a failure occurs the RSVP-TE interface is taken out of service. The BFD implementation on the 7705 SAR works on a hop-by-hop basis, and if BFD detects a link failure, only the two directly connected MPLS nodes are aware of that failure. If the node that detects the link failure is an LSR node, it generates PATH-ERR messages to the originators (the LER nodes) of the failing LSPs. If FRR is configured, the detecting node takes corrective action itself. See [LSP Redundancy](#) and [Fast Reroute \(FRR\)](#) for more information on these topics.

Timers

The following timers are implemented to ensure the successful operation of RSVP-TE:

- hold-timer — the hold timer defines the amount of time before an LSP is brought up and is in service, which provides protection against unreliable nodes and links
- resignal-timer — the resignal timer is used in conjunction with the route optimization process, especially after a reroute has occurred. If the newly computed path for an LSP has a better metric than the currently recorded hop list, then an attempt is made to resignal that LSP, and if the attempt is successful, then a make-before-break switchover occurs. If the attempt to resignal an LSP fails, the LSP continues to use the existing path and another resignal attempt is made the next time the timer expires.

When the resignal timer expires, a trap and syslog message are generated.

- retry-timer — the retry timer defines a period of time before a resignal attempt is made after an LSP failure. This delay time protects network resources against excessive signaling overhead.

LSP Resignal Limit

When an LSP fails, an LER node tries to resignal it. The following limit can be configured:

- **retry-limit** — the retry limit defines the number of ressignaling attempts in order to conserve the resources of the nodes in the network. There could be a serious loss of capacity due to a link failure where an infinite number of retries generate unnecessary message overhead.

RSVP-TE Message Pacing

RSVP-TE message pacing provides a means to limit the overwhelming number of RSVP-TE signaling messages that can occur in large MPLS networks during node failures. RSVP-TE message pacing allows the messages to be sent in timed intervals.

To protect nodes from receiving too many messages, the following message pacing parameters can be configured:

- **msg-pacing** — message pacing can be enabled or disabled
- **max-burst** — maximum burst defines the number of RSVP-TE messages that can be sent in the specified period of time
- **period** — period defines the interval of time used in conjunction with the max-burst parameter to send message pacing RSVP-TE messages

Message pacing needs to be enabled on all the nodes in a network to ensure the efficient operation of tier-1 nodes. Message pacing affects the number of RSVP-TE messages that a particular node can generate, not the number of messages it can receive. Thus, each node must be paced at a rate that allows the most loaded MPLS nodes to keep up with the number of messages they receive.



Note: Typically, a tier-1 node is an aggregator of tier-2 node transmissions, which is an aggregator of tier-3 node transmissions. Tier-1 nodes are often installed at an MTSO, while tier-3 nodes are often installed at cell sites.

RSVP-TE Overhead Refresh Reduction

RFC 2961, *RSVP Refresh Overhead Reduction Extensions*, defines enhancements to the RSVP-TE signaling protocol that reduce refresh overhead, which are in addition to the message pacing function.

These extensions are:

- **RSVP-TE message bundling** — RSVP-TE message bundling reduces the total number of RSVP-TE messages by aggregating the status information of multiple LSPs into a single RSVP-TE PDU. The 7705 SAR supports the receipt and processing of bundled RSVP-TE messages but not the transmission of bundled messages as specified in RFC 2961, section 3.3.
- **reliable message delivery** — reliable message delivery extends RSVP-TE to support MESSAGE_ACK. Each RSVP-TE PDU has a unique message-id for sequence tracking purposes. When an RSVP-TE message arrives, the recipient acknowledges the reception of the specific message-id (this is similar to TCP ACK messages). Lost PDUs can be detected and re-sent with this method, which helps reduce the refresh rate because there are two endpoints tracking the received/lost messages.
- **summary refresh** — the summary refresh capability uses a single message-id list to replace many individual refresh messages and sends negative ACKs (NACKs) for any message-id that cannot be matched (verified). The summary refresh capability reduces the number of message exchanges and message processing between peers. It does not reduce the amount of soft state stored in the node. The term soft state refers to the control state in hosts and routers that will expire if not refreshed within a specified amount of time (see RFC 2205 for information on soft state).

These capabilities can be enabled on a per-RSVP-TE interface basis and are referred to collectively as “refresh overhead reduction extensions”. When `refresh-reduction` is enabled on a 7705 SAR RSVP-TE interface, the node indicates this to its peer by setting a `refresh-reduction-capable` bit in the flags field of the common RSVP-TE header. If both peers of an RSVP-TE interface set this bit, all three of the capabilities listed above can be used. The node monitors the setting of this bit in received RSVP-TE messages from the peer on the interface. If the bit is cleared, the node stops sending summary refresh messages. If a peer did not set the `refresh-reduction-capable` bit, a 7705 SAR node does not attempt to send summary refresh messages.

Also, reliable delivery of RSVP-TE messages over the RSVP-TE interface can be enabled using the `reliable-delivery` option.

RSVP-TE Reservation Styles

LSPs can be signaled with explicit reservation styles for the reservation of resources, such as bandwidth. A reservation style describes a set of attributes for a reservation, including the sharing attributes and sender selection attributes. The style information is part of the LSP configuration. The 7705 SAR OS supports two reservation styles:

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

If the FRR option is enabled for the LSP and the facility FRR method is selected at the head-end node, only the SE reservation style is allowed. Furthermore, if a 7705 SAR 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 backup method supports both FF and SE styles.

LSP Redundancy

Each primary LSP can be protected by up to two secondary LSPs. When the LER detects a primary LSP failure, it signals its secondary LSPs, if any have been configured, and automatically switches to the first one that is available. LSP redundancy supports shared risk link groups (SRLG). See [Shared Risk Link Groups](#) for more information on SRLG.

LSP redundancy differs from the Fast Reroute (FRR) feature in that LSP redundancy is controlled by the LER that initiated the LSP, whereas FRR uses the node that detects the failure to take recovery action. This means that LSP redundancy takes longer to reroute traffic than FRR because failure messages need to traverse multiple hops to reach the LER and activate LSP redundancy, whereas an FRR-configured node responds immediately to bypass the failed node or link. See [Fast Reroute \(FRR\)](#) for more information on FRR.

The following parameters can be configured for primary and secondary LSPs:

- bandwidth — the amount of bandwidth needed for the secondary LSP can be reserved and can be any value; it does not need to be identical to the value reserved by the primary LSP. Bandwidth reservation can be set to 0, which is equivalent to reserving no bandwidth.
- inclusion and exclusion of nodes — by including or excluding certain nodes, you can ensure that the primary and secondary LSPs do not traverse the same nodes and therefore ensure successful recovery. Each secondary LSP can have its own list of included and excluded nodes.
- hop limit — the hop limit is the maximum number of LSR nodes that a secondary LSP can traverse, including the ingress and egress LER nodes
- standby (secondary LSPs only) — when a secondary LSP is configured for standby mode, it is signaled immediately and is ready to take over traffic the moment the LER learns of a primary LSP failure. This mode is also called hot-standby mode.

When a secondary LSP is not in standby mode, then it is only signaled when the primary LSP fails. If there is more than one secondary LSP, they are all signaled at the same time (upon detection of a primary LSP failure) and the first one to come up is used.

Fast Reroute (FRR)

FRR is a mechanism to protect against RSVP-TE signaled LSP failures by reacting to these failures as soon as possible. FRR is set up from the ILER, which signals the transit routers to pre-compute their backup LSPs. FRR creates a pre-computed backup LSP from each node in the LSP path. If a link or LSP between two routers fails, traffic is rerouted immediately onto the pre-computed backup LSP.



Note: In order for FRR to work, CSPF must be enabled.

The 7705 SAR supports FRR facility backup and one-to-one backup.

Facility backup mode allows FRR to be enabled on an aggregate basis and protects a whole node or a whole link, regardless of the number of LSPs using that link. In other words, facility backup mode creates a common bypass tunnel to protect all LSP-paths traversing a common facility path. It provides flexibility, faster provisioning, and faster convergence times compared with one-to-one backup or LSP redundancy. One-to-one backup allows FRR to be enabled on a per-LSP basis.

With both methods, MPLS switches build many possible detour routes on the nodes between the ingress and egress nodes of an LSP. The facility backup method creates a detour route between two nodes, called a bypass tunnel, which is a single tunnel that follows the primary LSP path except where the link or node has failed. Traffic then switches to the bypass tunnel. The bypass tunnel merges with the original LSP path at the merge point (MP) as soon as possible. The one-to-one backup method creates a detour route, called a detour LSP, for each LSP that needs to be rerouted. Unlike the bypass tunnel, the detour LSP takes the best path to the termination point, and does not merge with the original LSP as soon as possible. The detour LSPs of a one-to-one backup LSP can merge at a detour merge point (DMP), which can either be at the termination point or at a point along the primary LSP.

One of the major differences between facility and one-to-one backup is the scalability offered by the protection method. In facility backup mode, all LSPs of the same type are rerouted over the bypass tunnel. Hence, they are all protected against the failure of a node or link in the network. In facility backup mode, each LSR along the path verifies that it has a bypass tunnel available to meet its requirements; otherwise, if it can, it signals a new bypass tunnel based on the requirements. If a new LSP is configured for FRR facility backup, the existing backup tunnels are scanned and if any one of them can be used for recovery, it is preferred. If there are no common links, then a new bypass tunnel will be signaled, assuming that the LSP requirements can be met. One-to-one backup mode uses similar reroute and protection methods except a detour route is applied on a per-LSP basis.

The 7705 SAR uses CSPF to calculate the explicit route and dynamically signal the FRR LSP.

With facility backup mode, routers check the contents of the Record Route Object (RRO) in the received RESV message to determine the bypass tunnel endpoint in the FRR facility. For link protection, the router uses the RRO to check the IP address of the next-hop router attached to the far end of the link along with the label allocation information and to build the bypass tunnel. This label is preserved until the LSP is merged at the MP. For node protection, the router uses the RRO to determine the next-next-hop router and the label it is expecting. The collection of RRO information is enabled through the `record` and `record-label` options.

If, after this process, another LSP requests FRR using the facility backup method, then the router checks and compares its session object to the existing session object(s) and if there is a match, the router binds that LSP to the same bypass tunnel. If there is no match, another bypass is created.

FRR Terminology

[Table 6](#) provides definitions of terms used for FRR.

Table 6: FRR Terminology

Term	Definition
Backup path	The LSP that is responsible for backing up a protected LSP. A backup path can be a backup tunnel (facility backup) or a detour LSP (one-to-one backup).
Backup tunnel	The LSP that is used to back up one of the many LSPs in FRR facility (many-to-one) backup
Bypass tunnel	An LSP that is used to protect a set of LSPs passing over a common facility in FRR facility backup. A bypass tunnel can be configured manually or dynamically (see Dynamic and Manual Bypass LSPs).
CSPF	Constraint-based shortest path first
Detour route	Any alternate route that protects the primary path, such as a secondary path, FRR bypass tunnel, or FRR detour LSP. Note that the term “detour route” should not be confused with the term “detour LSP”. Detour route is a general term that refers to any alternate route, while detour LSP is a specific term that applies to one-to-one backup.

Table 6: FRR Terminology (Continued)

Term	Definition
Detour LSP	The LSP that is used to reroute traffic around a failure in FRR one-to-one backup. Note that the term “detour LSP” should not be confused with the term “detour route”. Detour route is a general term that refers to any alternate route, while detour LSP is a specific term that applies to one-to-one backup.
DMP	Detour merge point In the case of one-to-one backup, this is an LSR where multiple detours converge. Only one detour is signaled beyond that LSR.
Disjoint	See SRLG disjoint
Facility backup	A local repair method in which a single bypass tunnel is used to protect one or more LSPs that traverse the PLR, the resource being protected, and the Merge Point (in that order). Facility backup is distinct from a one-to-one backup tunnel, which has one backup path per protected path.
MP	Merge point The LSR where one or more backup tunnels rejoin the path of the protected LSP downstream of the potential failure. The same LSR may be both an MP and a PLR simultaneously.
NHOP bypass tunnel	Next-hop bypass tunnel A backup tunnel that bypasses a single link of the protected LSP
NNHOP bypass tunnel	Next-next-hop bypass tunnel A backup tunnel that bypasses a single node of the protected LSP
One-to-one backup	A local repair method in which a backup LSP is separately created for each protected LSP at a PLR
PLR	Point of local repair The head-end router of a backup tunnel or a detour LSP, where the term local repair refers to techniques used to repair an LSP tunnel quickly when a node or link along an LSP path fails
Primary path	An LSP that uses the routers specified by the path defined by the <code>primary path-name</code> command
Protected LSP	An LSP is protected at a given hop if it has one or more associated backup tunnels originating at that hop

Table 6: FRR Terminology (Continued)

Term	Definition
Reroutable LSP	Any LSP for which the head-end router requests local protection
Secondary path	An LSP that protects a primary path that uses LSP redundancy protection rather than FRR protection
SRLG disjoint	A path is considered to be SRLG disjoint from a given link or node if the path does not use any links or nodes that belong to the same SRLG as the given link or node

FRR Behavior

The FRR MPLS facility backup method and one-to-one backup method are configured on the ingress LER (ILER) by using the `fast-reroute` command.

The behavior of an LSP at an ILER with both FRR and a standby LSP path configured is as follows.

- When a downstream detour route (alternative path) becomes active at a Point of Local Repair (PLR):
The ILER switches to the standby LSP path as soon as it is notified of the reroute. If the primary LSP path is subsequently repaired at the PLR, the LSP switches back to the primary path. If the standby path goes down, the LSP is switched back to the primary path, even though the primary path is still on the detour route at the PLR.
- If the primary path goes down at the ILER while the LSP is on the standby path, the detour route at the ILER is torn down and, for one-to-one backup detour routes, a “path tear” is sent for the detour route. In other words, the detour route at the ILER does not protect the standby LSP. If and when the primary LSP is again successfully resignaled, the ILER detour route will be restarted.
- When the primary LSP fails at the ILER:
The LSP switches to the detour route. If the primary path undergoes a global revertive recovery, the LSP switches back to the primary path. If the LSP is on the detour route and the detour route fails, the LSP is switched to the standby path.
- Administrative groups are not taken into account when creating the detour routes for LSPs.

Dynamic and Manual Bypass LSPs

Users can disable dynamic bypass creation on a per-node basis using the `config>router>mpls>dynamic-bypass` command. Disabling dynamic bypass means that manual bypass is enabled. Dynamic bypass is enabled by default.

Dynamic bypass tunnels are implemented as per RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*. When an LSP is signaled and the Local Protection flag in the Session_attribute object is set, or the FRR object in the PATH message indicates that facility backup is desired, the PLR establishes a bypass tunnel to provide node and link protection. If there exists a bypass LSP that merges with the protected LSP at a downstream node, and if this LSP satisfies the constraints in the FRR object, then this bypass tunnel is selected and used. The `frr-object` command specifies whether facility backup is signaled in the FRR object.

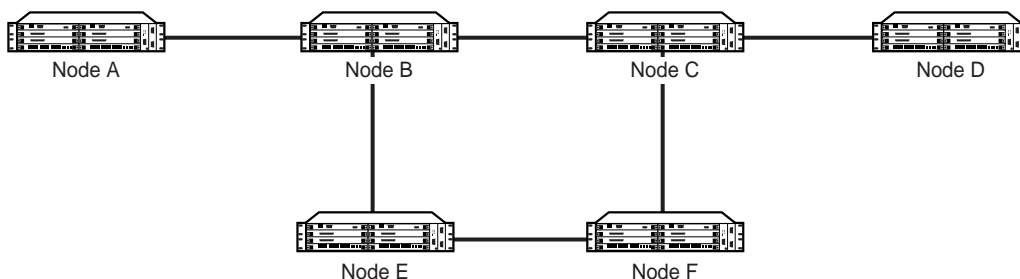
The manual bypass feature allows an LSP to be preconfigured from a Point of Local Repair (PLR) that will be used exclusively for bypass protection. When a PATH message for a new LSP requests bypass protection, the node first checks for a manual bypass tunnel that satisfies the path constraints. If one is found, it is selected and used. If no manual bypass tunnel is found, the 7705 SAR dynamically signals a bypass LSP in the default behavior. To configure a manual bypass LSP, use the `bypass-only` option in the `config>router>mpls>lsp lsp-name [bypass-only]` command.

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

Bypass LSP Selection Rules for the PLR

[Figure 5](#) shows a sample network used to illustrate the LSP selection rules for a PLR bypass scenario.

Figure 5: Bypass Tunnel Node Example



20123

The PLR uses the following rules to select a bypass LSP from among multiple bypass LSPs (manually and dynamically created) when establishing the primary LSP path or when searching for a bypass for a protected LSP that does not have an association with a bypass tunnel.

1. The MPLS/RSVP-TE task in the PLR node checks for an existing manual bypass tunnel that satisfies the constraints. If the PATH message for the primary LSP path indicated that node protection is desired, which is the default LSP FRR setting at the head-end node, then the MPLS/RSVP-TE task searches for a node-protect bypass LSP. If the PATH message for the primary LSP path indicated that link protection is desired, then it searches for a link-protect bypass LSP.
2. If multiple manual bypass LSPs satisfying the path constraints exist, the PLR will prefer a manual bypass LSP terminating closer to the PLR over a manual bypass LSP terminating further away. If multiple manual bypass LSPs satisfying the path constraints terminate on the same downstream node, the PLR selects the one with the lowest IGP path cost, or if there is a tie, it picks the first one available.
3. If none of the manual bypass LSPs satisfy the constraints and dynamic bypass tunnels have not been disabled on the PLR node, then the MPLS/RSVP-TE task in the PLR node checks to determine if any of the already established dynamic bypass LSPs of the requested type satisfy the constraints.
4. If none of the dynamic bypass LSPs satisfy the constraints, then the MPLS/RSVP-TE 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 is desired, and no manual bypass was found after Step 1, and/or no dynamic bypass LSP was found after three attempts to perform Step 3, the MPLS/RSVP-TE task will repeat Steps 1 to 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 is 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 LSPs that requested FRR protection but are not currently associated with a bypass tunnel, the PLR node—upon reception of an RESV refresh message on the primary LSP path—repeats Steps 1 to 7.

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

If the user configures a bypass tunnel on Node B ([Figure 5](#)) and dynamic bypass tunnels have been disabled, LSPs that had been previously signaled and that were not associated with any manual bypass tunnel (for example, none existed) will be associated with the manual bypass tunnel, if it is suitable. The node checks for the availability of a suitable bypass tunnel for each of the outstanding LSPs every time an 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 that had been previously signaled over dynamic bypass tunnels will not automatically be switched to the manual bypass tunnel, even if the manual bypass tunnel is a more optimized path. The user must perform a make-before-break switchover at the head end of these LSPs. The make-before-break process is enabled using the `adaptive` option.

If the manual bypass tunnel goes into the down state on 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 the “protection available” flag in the next RESV refresh message for each of these LSPs. If it cannot find one, it will keep checking for one every time an RESV message is received for each of the remaining LSPs. When the manual bypass tunnel is back up, the LSPs that did not find a match are associated back with this tunnel and the protection available flag is set starting in the next RESV refresh message.

If the manual bypass tunnel goes into the down state on Node B and dynamic bypass tunnels have not been disabled, Node B will automatically signal a dynamic bypass tunnel to protect the LSPs if a suitable one does not exist. Similarly, if an LSP is signaled while the manual bypass tunnel 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 up, the node will not switch the protected LSPs from the dynamic bypass tunnel to the manual bypass tunnel.

FRR Node Protection (Facility Backup)

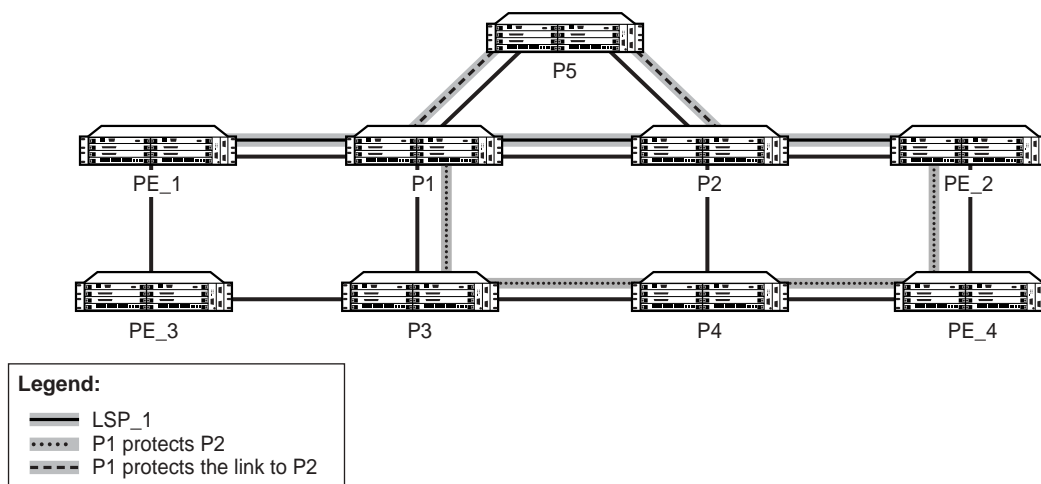
The MPLS Fast Reroute (FRR) functionality enables PLRs to be aware of the lack of node protection and lets them regularly probe for a node bypass via the `node-protect` command.

When enabled, the `node-protect` command provides node protection for the specified LSP. If node protection cannot be provided, link protection is attempted. If link protection cannot be provided, no protection is provided. When disabled via the `no` form of the command, link protection is attempted, and if link protection cannot be provided, no protection is provided.

For example, assume the following for the LSP scenario in [Figure 6](#).

1. LSP_1 is between PE_1 and PE_2 (via P1 and P2), and has CSPF, FRR facility backup, and FRR node protection enabled.
2. P1 protects P2 with bypass nodes P1 - P3 - P4 - PE_4 - PE_3.
3. If P4 fails, P1 tries to establish the bypass node three times.
4. When the bypass node creation fails (there is no bypass route), P1 will protect link P1-P2.
5. P1 protects the link to P2 through P1 - P5 - P2.
6. P4 returns online.

Figure 6: FRR Node-Protection Example



20124

Fast Reroute (FRR)

LSP_1 had requested node protection, but due to lack of an available path it could only obtain link protection. Therefore, every 60 s, the PLR for LSP_1 will search for a new path that might be able to provide node protection. Once P4 is back online and such a path is available, a new bypass tunnel will be signaled and LSP_1 will be associated with this new bypass tunnel.

Shared Risk Link Groups

A shared risk link group (SRLG) represents a set of interfaces (or links) that share the same risk of failing because they may be subjected to the same resource failures or defects. Two examples where the same risk of failure exists are fiber links that share the same conduit, and multiple wavelengths that share the same fiber.

SRLGs are supported by both LSP redundancy protection and FRR protection. SRLGs allow the user to prepare a detour route that is disjoint from the primary LSP path. See [Disjoint and Non-disjoint Paths](#).

The SRLG feature ensures that a primary and secondary LSP path, or a bypass tunnel or detour LSP path, do not share SRLGs. That is, they do not share the same sets of links that are considered to have a similar (or identical) chance of failure.

To use SRLGs, the user first creates an SRLG by assigning one or more routers to the SRLG. Then, the user links the SRLG to an MPLS interface and enables the SRLG feature on the LSP path. Note that SRLGs cannot be assigned to the system interface.

SRLGs for Secondary LSP Paths

SRLGs for secondary LSP paths apply when LSP redundancy protection is used.

When setting up the secondary path, enable the `srlg` option on the secondary path to ensure that CSPF includes the SRLG constraint in its route calculation. To make an accurate computation, CSPF requires that the primary LSP be established and in the up state (because the head-end LER needs the most current explicit route object (ERO) for the primary path, and the most current ERO is built during primary path CSPF computation). The ERO includes the list of SRLGs.

At the establishment of a secondary path with the SRLG constraint, the MPLS/RSVP-TE task queries CSPF again, which provides the list of SRLGs to be avoided. CSPF prunes all links having interfaces that belong to the same SRLGs as the interfaces included in the ERO of the primary path. If CSPF finds an eligible path, the secondary path is set up. If CSPF does not find an eligible path, MPLS/RSVP-TE keeps retrying the requests to CSPF.

SRLGs for FRR LSP Paths

When setting up the FRR bypass or detour LSP, enable the `srlg-frr` option on FRR to ensure that CSPF includes the SRLG constraint in its route calculation. CSPF prunes all links that are in the SRLG being used by the primary LSP during the calculation of the FRR path. If one or more paths are found, CSPF sets up the FRR bypass or detour LSP based on the best cost and signals the FRR LSP.

If there is no path found based on the above calculation and the `srlg-frr` command has the `strict` option set, then the FRR LSP is not set up and the MPLS/RSVP-TE task keeps trying to set up a path. If the `strict` option is not set, then the FRR LSP is set up based on the other TE constraints (that is, excluding the SRLG constraint).

Disjoint and Non-disjoint Paths

A path is considered to be SRLG disjoint from a given link (or node) if the path does not use any links (or nodes) that belong to the same SRLG as the given link (or node). Eligible disjoint paths are found by CSPF when the SRLG constraint is included in the CSPF route calculation (referred to as the strict SRLG condition).

When LSP redundancy is used, the secondary LSP is always signaled with a strict SRLG condition.

When FRR is used, the FRR bypass or detour LSP may have a strict or non-strict SRLG condition. If the `strict` option is used with the `srlg-frr` command, then the bypass LSP must be on the list of eligible paths found by the CSPF calculation that included the SRLG constraint. If the `strict` option is not used, then it is possible for the bypass or detour LSP to be non-disjoint. The non-disjoint case is supported only if the SRLG is not strict.

At the PLR, if an FRR tunnel is needed to protect a primary LSP, the priority order for selecting that FRR tunnel is as follows:

1. Manual bypass disjoint
2. Manual bypass non-disjoint (eligible only if `srlg-frr` is non-strict)
3. Dynamic bypass disjoint
4. Dynamic bypass non-disjoint (eligible only if `srlg-frr` is non-strict)

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

Enabling Disjoint Backup Paths

A typical application of the SRLG feature is to provide automatic setup of secondary LSPs or FRR bypass or detour LSPs, in order to minimize the probability that they share the same failure risks with the primary LSP path (see [Figure 7](#) and [Figure 8](#)).

[Figure 7](#) illustrates SRLG when LSP redundancy is used, where SRLG_1 contains the interfaces that define links A-B, B-C, and C-D. The primary path uses these links to connect node A to node D. In the event of a failure along the primary path, the secondary path cannot use any of the links in SRLG_1 and takes the path from node A to nodes E, F, G, H, J, and D.

[Figure 8](#) illustrates SRLG when FRR bypass is used, where SRLG_1 is the same as in [Figure 7](#). Since FRR bypass is used, the following possible reroutes may occur, depending on where the failure occurs:

- if node B fails, the bypass is from node A to nodes E, F, G, H, and C
- if node C fails, the bypass is from node B to nodes F, G, H, J, and D
- if link C-D fails, the bypass is from node C to nodes H, J, and D

The SRLG feature is supported on OSPF and IS-IS interfaces for which RSVP-TE is enabled.

The following steps describe how to enable SRLG disjoint backup paths for LSP redundancy and FRR.

LSP Redundancy for Primary/Secondary (standby) SRLG Disjoint Configuration

- Create an SRLG-group (similar to creating an admin group).
- Link the SRLG-group to MPLS interfaces.
- Configure primary and secondary LSP paths, and enable SRLG on the secondary LSP path. Note that the SRLG secondary LSP path(s) will always perform a strict CSPF query.

The setting of the `srlg-frr` command is irrelevant in this case (see the [srlg-frr](#) command).

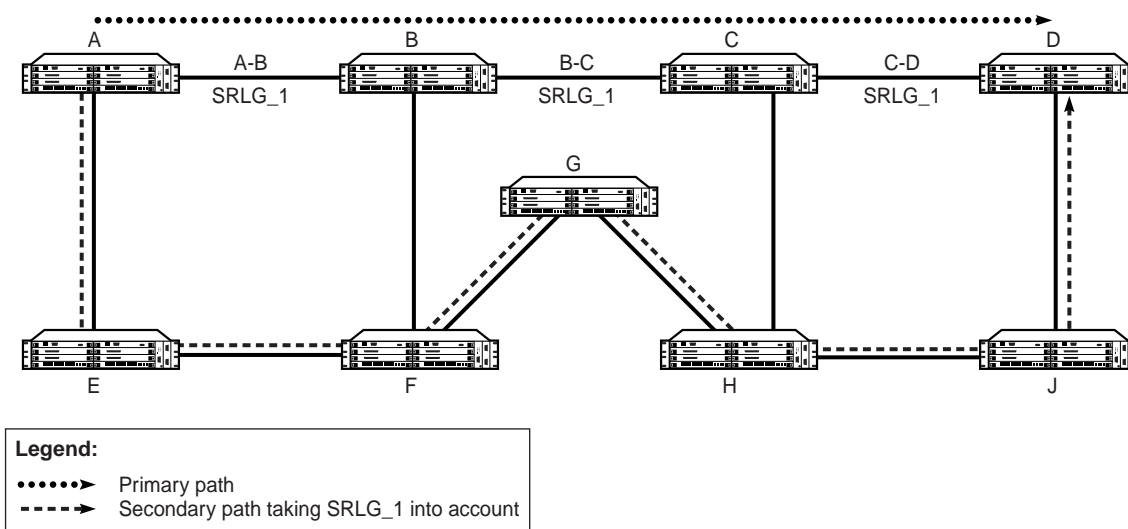
FRR Bypass Tunnel or Detour LSP SRLG Disjoint Configuration

- Create an SRLG-group (similar to creating an admin group).
- Link the SRLG-group to MPLS interfaces.
- Enable the `strict` option on the `srlg-frr` command, which is a system-wide command that forces the CSPF calculation for every LSP path to take any configured SRLG membership(s) into account.

- Configure primary FRR (facility backup or one-to-one backup) LSP path(s). Note that each PLR will create a bypass or detour LSP that will only avoid the SRLG membership(s) configured on the primary LSP path egress interface. For one-to-one backup, detour-detour merging is out of the control of the PLR. The PLR will not ensure that the FRR detour will be prohibited from merging with a colliding detour LSP. For facility backup, given that there are several bypass types to bind to, the priority rules shown in [Disjoint and Non-disjoint Paths](#) are used.

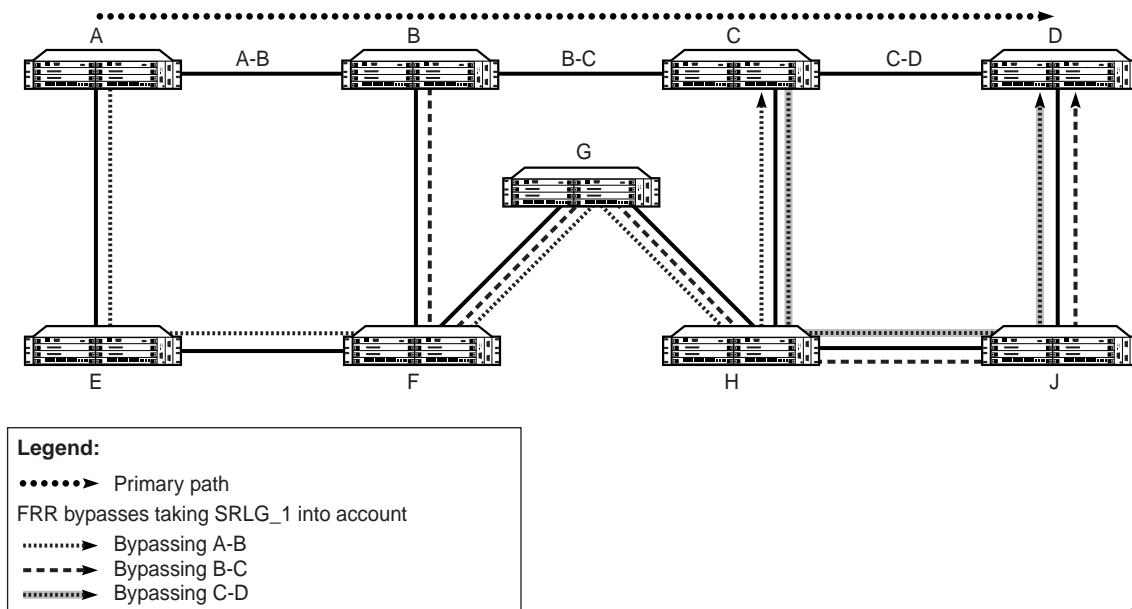
Manually configured bypasses that do not use CSPF are not considered as possible backup paths.

Figure 7: Disjoint Primary and Secondary LSPs



20482

Figure 8: Disjoint FRR Bypass LSPs



20483

RSVP-TE Graceful Shutdown

RSVP-TE graceful shutdown provides a method to reroute transit LSPs in a bulk fashion away from a node prior to maintenance of that node. A PathErr message with the error code “Local Maintenance on TE Link required Flag” (if the affected network element is a link) or the error code “Local node maintenance required” (if the affected network element is the node) is sent before the links or node are taken out of service.

When an LER receives the message, it performs a make-before-break on the LSP path to move the LSPs away from the links/nodes whose IP addresses are indicated in the PathErr message and reroute them. Affected link/node resources are flagged in the TE database so that other routers will signal LSPs using the affected resources only as a last resort.

Graceful shutdown can be enabled on a per-interface basis or on all interfaces on the node if the whole node must be taken out of service.

MPLS Service Usage

Alcatel-Lucent routers enable service providers to deliver virtual private networks (VPNs) and Internet access using Generic Routing Encapsulation (GRE), IP, and/or MPLS tunnels, with Ethernet and/or SONET/SDH interfaces.

Service Destination Points

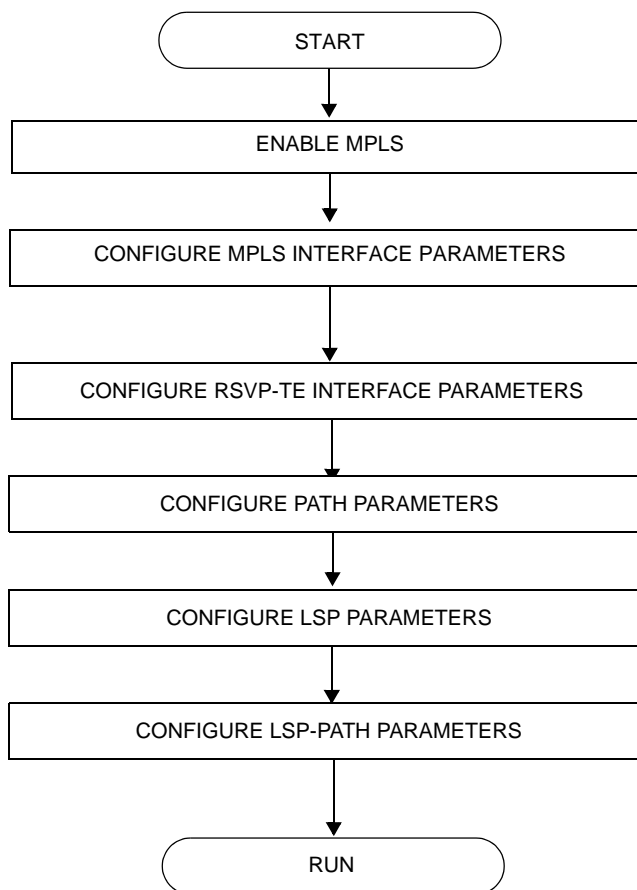
A service destination point (SDP) acts as a logical way of directing traffic from one 7705 SAR router to another through a unidirectional (one-way) service tunnel. The SDP terminates at the far-end 7705 SAR router, which directs packets to the correct service egress service access point (SAP) on that device. All services mapped to an SDP use the GRE, IP, or MPLS transport encapsulation type.

For information about service transport tunnels, refer to the 7705 SAR OS Services Guide. Service transport tunnels can support up to eight forwarding classes and can be used by multiple services. Multiple LSPs with the same destination can be used to load balance traffic.

MPLS and RSVP-TE Configuration Process Overview

Figure 9 displays the process to configure MPLS and RSVP-TE parameters.

Figure 9: MPLS and RSVP-TE Configuration and Implementation Flow



Configuration Notes

Network and system interfaces must be configured in the `config>router>interface` context before they can be specified in MPLS. Refer to the 7705 SAR OS Router Configuration Guide for interface configuration information.

This section describes MPLS and RSVP-TE 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.
- CSPF must be enabled in order for administrative groups and SRLGs to be relevant.

Reference Sources

For information on supported IETF drafts and standards, as well as standard and proprietary MIBs, refer to [Standards and Protocol Support](#).

Configuring MPLS and RSVP-TE with CLI

This section provides information to configure MPLS and RSVP-TE using the CLI.

Topics in this section include:

- [MPLS Configuration Overview on page 76](#)
- [Basic MPLS Configuration on page 79](#)
- [Common Configuration Tasks on page 80](#)
- [MPLS Configuration Management Tasks on page 90](#)
- [RSVP-TE Configuration Management Tasks on page 94](#)

MPLS Configuration Overview

MPLS enables routers to forward traffic based on a label embedded in the packet header. A router examines the label to determine the next hop for the packet, instead of router address lookups to the next node when forwarding packets.

To implement MPLS on an LSP for outer tunnel and pseudowire assignment, the following entities must be configured:

- [Router Interface](#)
→ [E-LSP for Differentiated Services](#)
- [Paths](#)
- [LSPs](#)
- [Pseudowires](#)
- [Signaling Protocol](#) (for RSVP-TE or LDP)

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.

E-LSP for Differentiated Services

An EXP-inferred LSP (E-LSP) is an LSP that can support a variety of VLLs or traffic types. Up to eight types of traffic can be multiplexed over an E-LSP.

The prioritization of mission-critical traffic is handled by the settings of the three EXP bits. The EXP bits designate the importance of a particular packet. The classification and queuing at the Provider (P) or Provider Edge (PE) nodes typically take place based on the value of the EXP bits. Refer to the 7705 SAR OS Quality of Service Guide for more information on the use of EXP bits and differentiated services on the 7705 SAR.

Paths

To configure signaled LSPs, you must first create one or more named paths on the ingress router using the `config>router>mpls>path` command. For each path, the transit routers (hops) in the path are specified.

LSPs

The 7705 SAR supports static and dynamic LSPs.

To configure MPLS-signaled (dynamic) LSPs, the LSP must run from an ingress LER to an egress LER. Configure the dynamic LSP only at the ingress router, and either configure the LSP to allow the router software to make the forwarding decisions or configure some or all routers in the LSP path statically. The LSP is set up by RSVP-TE signaling messages. The 7705 SAR OS automatically manages label values. Labels that are automatically assigned have values ranging from 1,024 through 1 048 575 (see [Label Values](#)).

A static LSP is a manually configured LSP where the next hop IP address and the outgoing label are explicitly specified.

To establish a static LSP, an LSP must be configured from an ingress LER to an egress LER. Labels must be manually assigned and the label values must be within the range of 32 to 1023 (see [Label Values](#)).

Pseudowires

To configure PW/VLL labels, the PW/VLL service must be configured. PW/VLL labels can be configured manually as statically allocated labels using any unused label within the static label range. Pseudowire/VLL labels can also be dynamically assigned by targeted LDP. Statically allocated labels and dynamically allocated labels are designated differently in the label information base.

PW/VLL labels are uniquely identified against a 7705 SAR, not against an interface or module.

As defined in RFC 3036, *LDP Specification*, and RFC 4447 *Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)*, label distribution is handled in the Downstream Unsolicited (DU) mode. Generic Label TLV is used for all setup and maintenance operations.

Signaling Protocol

For static LSPs, the path and the label mappings and actions configured at each hop must be specified manually. RSVP-TE or LDP is not required for static LSPs.

For dynamic LSPs, RSVP-TE or LDP must be turned on. See [RSVP and RSVP-TE](#) or [Label Distribution Protocol](#).

To implement dynamic pseudowire/VLL labels, entities must be enabled as follows:

- MPLS must be enabled on all routers that are part of a static LSP
- LDP must be enabled on the ingress and egress LERs

When MPLS is enabled and either RSVP-TE or LDP is also enabled, MPLS uses RSVP-TE or LDP to set up the configured LSPs. For example, when you configure an LSP with both MPLS and RSVP-TE running, RSVP-TE initiates a session to create the LSP. RSVP-TE uses the local router as the RSVP-TE session sender and the LSP destination as the RSVP-TE session receiver. Once the RSVP-TE session is created, the LSP is set up on the path created by the session. If the session is not successfully created, RSVP-TE 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 gives configuration examples of common configuration tasks. To enable MPLS on a 7705 SAR router, you must configure at least one MPLS interface. The MPLS interface is configured in the `config>router>mpls` context. The other MPLS configuration parameters are optional.

The following example displays an MPLS configuration output.

```
A:ALU-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 "to-NYC"
    hop 1 10.10.10.104 strict
    no shutdown
exit
path "secondary-path"
    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:ALU-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-TE (for RSVP-TE-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 the path 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 settings under the LSP hierarchy are optional.

Configuring MPLS Components

Use the MPLS and RSVP-TE CLI syntax shown in the following information for:

- [Configuring Global MPLS Parameters](#)
- [Configuring an MPLS Interface](#)
- [Configuring MPLS Paths](#)
- [Configuring an MPLS LSP](#)
- [Configuring a Static LSP](#)
- [Configuring Manual Bypass Tunnels](#)
- [Configuring RSVP-TE Parameters](#)
- [Configuring RSVP-TE Message Pacing Parameters](#)

Configuring Global MPLS Parameters

Admin groups can signify link colors, such as red, yellow, or green, or some other link quality. Shared risk link groups (SRLGs) are lists of interfaces that share the same risk of failure due to shared resources. MPLS interfaces advertise the admin groups and SRLGs that they support. CSPF uses the information when paths are computed for constraint-based LSPs. CSPF must be enabled in order for admin groups and SRLGs to be relevant.

To configure global MPLS parameters, enter the following commands:

CLI Syntax:

```
config>router>mpls
    admin-group group-name group-value
    dynamic-bypass [enable | disable]
    frr-object
    hold-timer seconds
    resignal-timer minutes
    srlg-frr [strict]
    srlg-group group-name value group-value
```

Example:

```
config>router# mpls
config>router>mpls# admin-group "green" 15
config>router>mpls# admin-group "red" 25
config>router>mpls# admin-group "yellow" 20
config>router>mpls# frr-object
config>router>mpls# hold-timer 3
config>router>mpls# resignal-timer 500
config>router>mpls# srlg-frr strict
config>router>mpls# srlg-group "SRLG_fiber_1" value 50
```

The following example displays a global MPLS configuration output.

```
A:ALU-1>config>router>mpls# info
-----
    admin-group "green" 15
    admin-group "red" 25
    admin-group "yellow" 20
    frr-object
    hold-timer 3
    resignal-timer 500
    srlg-frr strict
    srlg-group "SRLG_fiber_1" 50
-----
A:ALU-1>config>router>mpls# info
```

Configuring an MPLS Interface

The interface must exist in the system before it can be configured as an MPLS interface; refer to the 7705 SAR OS Router Configuration Guide for more information.

Once the MPLS protocol instance is created, the `no shutdown` command is not required since MPLS is administratively enabled upon creation. Configure the `label-map` parameters if the interface is used in a static LSP.

Use the following CLI syntax to configure an MPLS interface on a router:

CLI Syntax:

```
config>router>mpls
interface ip-int-name
    admin-group group-name [group-name...(up to 32 max)]
    label-map in-label
        pop
        swap out-label next-hop ip-address
    no shutdown
    srlg-group group-name [group-name...(up to 5 max)]
    te-metric value
    no shutdown
```

Example:

```
config>router# mpls
config>router>mpls# interface to-104
config>router>mpls>if# label-map 35
config>router>mpls>if>label-map# swap 36 next-hop
10.10.10.91
config>router>mpls>if>label-map# no shutdown
config>router>mpls>if>label-map# exit
config>router>mpls>if# srlg-group "SRLG_fiber_1"
config>router>mpls>if# no shutdown
config>router>mpls# exit
```

The following example displays the interface configuration output.

```
A:ALU-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
  srlg-group "SRLG_fiber_1"
  exit
exit
no shutdown
-----
A:ALU-1>config>router>mpls#
```

Configuring MPLS Paths

Configure an MPLS path for use by LSPs. When configuring an MPLS path, the IP address of each hop 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 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 example displays a path configuration output.

```
A:ALU-1>config>router>mpls# info
-----
      interface "system"
      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
      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
-----
A:ALU-1>config>router>mpls#
```

Configuring an MPLS LSP

Configure an LSP 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:ALU-1>config>router>mpls# 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:ALU-1>config>router>mpls#
```

Configuring a Static LSP

An LSP can be explicitly (manually) configured. The reserved range of static LSP labels is 32 to 1023. Static LSPs are configured on every node along the LSP 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 label nexthop ip-address
    no shutdown
```

Example:

```
config>router# mpls
config>router>mpls# static-lsp static-LSP
config>router>mpls>static-lsp$ to 10.10.10.124
config>router>mpls>static-lsp# push 60 nexthop 10.10.42.3
config>router>mpls>static-lsp# no shutdown
config>router>mpls>static-lsp# exit
```

The following example displays the static LSP configuration output.

```
ALU-1>config>router>mpls# info
-----
...
    static-lsp "static-LSP"
        to 10.10.10.124
        push 60 nexthop 10.10.42.3
        no shutdown
    exit
-----
```

Configuring a Fast-Retry Timer for Static LSPs

A fast-retry timer can be configured for static LSPs. When a static LSP is trying to come up, MPLS tries to resolve the ARP entry for the next hop of the LSP. This request may fail because the next hop might still be down or unavailable. In that case, MPLS starts a retry timer before making the next request. The fast-retry command allows the user to configure the retry timer so that the LSP comes up shortly after the next hop is available.

Use the following CLI syntax to configure a fast-retry timer for static LSPs:

CLI Syntax:

```
config>router>mpls
    static-lsp-fast-retry seconds
```

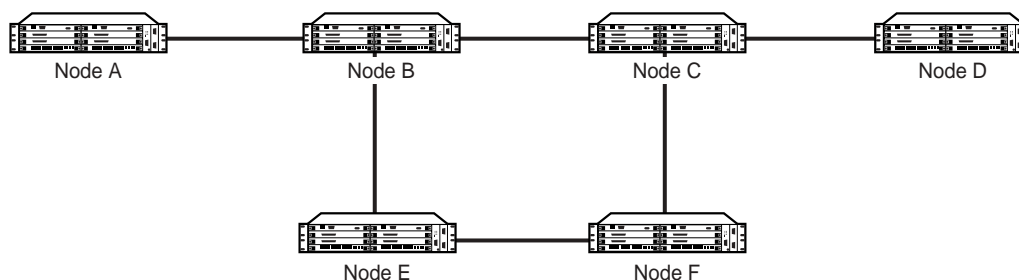
Example:

```
config>router# mpls
config>router>mpls# static-lsp-fast-retry 15
```

Configuring Manual Bypass Tunnels

Consider the following network setup in [Figure 10](#). Assume that a manual bypass tunnel must be configured on Node B.

Figure 10: Manual Bypass Tunnels



20123

Step 1. Disable dynamic bypass tunnels on Node B.

The CLI syntax for this configuration is:

```
config>router>mpls>dynamic-bypass [disable|enable]
```

By default, dynamic bypass tunnels are enabled.

Step 2. Configure an LSP on Node B, such as B-E-F-C, which will be used only as a bypass. Specify each hop in the path and assign its `strict` or `loose` option; in this case, the bypass LSP will have a strict path. Designate the LSP as a primary LSP.

The CLI syntax for this configuration is:

```
config>router>mpls>path path-name>hop hop-index ip-address
[strict | loose]
```

```
config>router>mpls>lsp lsp-name bypass-only
```

Including the **bypass-only** keyword disables some options under the LSP configuration. See [Table 7](#).

Table 7: Disabled and Enabled Options for Bypass-Only

Disabled Options	Enabled Options
<ul style="list-style-type: none"> bandwidth fast-reroute secondary 	<ul style="list-style-type: none"> adaptive adspec cspf exclude hop-limit include metric

The following example displays a bypass tunnel configuration output.

```
A:ALU-48>config>router>mpls># 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" bypass-only
      to 10.10.10.15
      primary "BEFC"
      exit
      no shutdown
  ...
-----
A:ALU-48#
```

Step 3. Configure an LSP from A to D and indicate fast-reroute bypass protection by selecting facility as the FRR method.

The CLI syntax for this configuration is:

```
config>router>mpls>lsp lsp-name>fast-reroute facility
```

If the LSP from A to D goes through Node B and bypass is requested, the next hop is Node C, and there is a manually configured bypass-only tunnel from B to C that excludes link BC (that is, path BEFC), then Node B uses the bypass-only tunnel.

Configuring RSVP-TE Parameters

RSVP-TE is used to set up LSPs. RSVP-TE 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 `config>router>rsvp` context.

Initially, interfaces are configured in the `config>router>mpls>interface` context. Only these existing (MPLS) interfaces are available to be modified in the `config>router>rsvp` context. Interfaces cannot be directly added in the `rsvp` context.

The following example displays an RSVP-TE configuration output.

```
A:ALU-1>config>router>rsvp# info
-----
      interface "system"
          no shutdown
      exit
      interface to-104
          hello-interval 4000
          no shutdown
      exit
      no shutdown
-----
A:ALU-1>config>router>rsvp#
```

Configuring RSVP-TE Message Pacing Parameters

RSVP-TE 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-TE message pacing parameters:

CLI Syntax:

```
config>router>rsvp
no shutdown
msg-pacing
    period milli-seconds
    max-burst number
```

The following example displays an RSVP-TE message pacing configuration output.

```
A:ALU-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:ALU-1>config>router>rsvp#
```

MPLS Configuration Management Tasks

This section discusses the following MPLS configuration management tasks:

- [Deleting MPLS](#)
- [Modifying MPLS Parameters](#)
- [Modifying an MPLS LSP](#)
- [Modifying MPLS Path Parameters](#)
- [Modifying MPLS Static LSP Parameters](#)
- [Deleting an MPLS Interface](#)

Deleting MPLS

The `no mpls` command typically removes an MPLS instance and all associated information. However, MPLS must be disabled (shut down) and all SDP bindings to LSPs removed before an MPLS instance can be deleted. Once MPLS is shut down, the `no mpls` command deletes the protocol instance and removes all configuration parameters for the MPLS instance.

If MPLS is not shut down first, when the `no mpls` command is executed, a warning message on the console indicates that MPLS is still administratively up.

To delete the MPLS instance:

1. Disable the MPLS instance using the `shutdown` command.
2. Remove the MPLS instance from the router using the `no mpls` command.

CLI Syntax: `config>router# no mpls`

Modifying MPLS Parameters



Note: You must shut down MPLS entities in order to modify parameters. Re-enable (no shutdown) the entity for the change to take effect.

Modifying an MPLS LSP

Some MPLS LSP parameters (such as primary and secondary), must be shut down before they can be edited or deleted from the configuration.

The following example displays an MPLS LSP configuration output. Refer to [Configuring an MPLS Interface](#).

```
A:ALU-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:ALU-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 example displays an MPLS path configuration output. Refer to [Configuring MPLS Paths](#).

```
A:ALU-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:ALU-1>config>router>mpls#
```

Modifying MPLS Static LSP Parameters

Use the `show>service>router>static-lsp` command to display a list of LSPs.

In order to modify static LSP parameters, the `config>router>mpls>static-lsp lsp-name` context must be shut down.

To modify an LSP:

1. Access the specific LSP by specifying the LSP name, and then shut it down.
2. Enter the parameter to modify and then enter the new information.

Example:

```
config>router# mpls
config>router>mpls# static-lsp "static-LSP"
config>router>mpls>static-lsp# shutdown
config>router>mpls>static-lsp# to 10.10.0.234
config>router>mpls>static-lsp# push 1023 nexthop
10.10.8.114
config>router>mpls>static-lsp# no shutdown
config>router>mpls>static-lsp# exit
```

The following example displays the static LSP configuration output.

```
ALU-1>config>router>mpls# info
-----
...
    static-lsp "static-LSP"
        to 10.10.10.234
        push 1023 nexthop 10.10.8.114
        no shutdown
    exit
    no shutdown
-----
ALU-1>config>router>mpls#
```

Deleting an MPLS Interface

To delete an interface from the MPLS configuration:

1. Administratively disable the interface using the `shutdown` command.
2. Delete the interface with the `no interface` command.

CLI Syntax: `mpls`

```

    interface ip-int-name
    shutdown
    exit
    no interface ip-int-name
  
```

Example:

```

config>router# mpls
config>router>mpls# interface to-104
config>router>mpls>if# shutdown
config>router>mpls>if# exit
config>router>mpls# no interface to-104
  
```

The following example displays the configuration output when interface “to-104” has been deleted.

```

A:ALU-1>config>router>mpls# info
-----
...
    admin-group "green" 15
    admin-group "red" 25
    admin-group "yellow" 20
    interface "system"
    exit
    no shutdown
-----
A:ALU-1>config>router>mpls#
  
```

RSVP-TE Configuration Management Tasks

This section discusses the following RSVP-TE configuration management tasks:

- [Modifying RSVP-TE Parameters](#)
- [Modifying RSVP-TE Message Pacing Parameters](#)
- [Deleting an Interface from RSVP-TE](#)

Modifying RSVP-TE Parameters

Only interfaces configured in the MPLS context can be modified in the `rsvp` context.

The `no rsvp` command deletes this RSVP-TE protocol instance and removes all configuration parameters for this RSVP-TE instance. The `shutdown` command suspends the execution and maintains the existing configuration.

The following example displays a modified RSVP-TE configuration output.

```
A:ALU-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:ALU-1>config>router>rsvp#
```

Modifying RSVP-TE Message Pacing Parameters

RSVP-TE 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 a modified RSVP-TE message pacing configuration output. Refer to [Configuring RSVP-TE Message Pacing Parameters](#).

```
A:ALU-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:ALU-1>config>router>rsvp#
```

Deleting an Interface from RSVP-TE

Interfaces cannot be deleted directly from the RSVP-TE configuration. Because an interface is created in the `mpls` context and then configured in the `rsvp` context, it can only be deleted in the `mpls` context. This removes the association from RSVP-TE.

Refer to [Deleting an MPLS Interface](#).

MPLS and RSVP-TE Command Reference

Command Hierarchies

- [MPLS Commands](#)
- [RSVP-TE Commands](#)
- [Show Commands](#)
- Tools Commands (refer to Tools section of 7705 SAR OS OAM and Diagnostics Guide)
- [Clear Commands](#)
- [Debug Commands](#)

MPLS Commands

```

config
— router [router-name]
— [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)]
— [no] label-map in-label
— [no] pop
— swap out-label nexthop ip-address
— no swap
— [no] shutdown
— [no] shutdown
— [no] srlg-group group-name [group-name...(up to 5 max)]
— te-metric value
— no te-metric
— [no] lsp lsp-name [bypass-only]
— [no] adaptive
— [no] adspec
— [no] cspf [use-te-metric]
— [no] exclude group-name [group-name...(up to 5 max)]
— [no] fast-reroute [frr-method]
— bandwidth rate-in-mbps
— no bandwidth
— hop-limit limit
— 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)]
— 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] record
    — [no] record-label
    — [no] shutdown
    — [no] srlg
    — [no] standby
— [no] shutdown
— to ip-address
— [no] path path-name
    — hop hop-index ip-address {strict | loose}
    — no hop hop-index
    — [no] shutdown
— resignal-timer minutes
— no resignal-timer
— srlg-frr [strict]
— no srlg-frr
— srlg-group group-name {value group-value}
— no srlg-group group-name
— [no] shutdown
— [no] static-lsp lsp-name
    — push label nexthop ip-address
    — no push label
    — to ip-address
    — [no] shutdown
— static-lsp-fast-retry seconds
— no static-lsp-fast-retry

```

RSVP-TE Commands

```

config
  — router
    — [no] rsvp
      — [no] graceful-shutdown
      — [no] interface ip-int-name
        — authentication-key [authentication-key | hash-key] [hash | hash2]
        — no authentication-key
        — [no] bfd-enable
        — [no] graceful-shutdown
        — hello-interval milli-seconds
        — no hello-interval
        — [no] refresh-reduction
          — [no] reliable-delivery
        — [no] shutdown
        — subscription percentage
        — no subscription
      — [no] 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-reduction-over-bypass [enable | disable]
      — refresh-time seconds
      — no refresh-time
      — [no] shutdown

```

Show Commands

```

show
  — router
    — mpls
      — admin-group group-name
      — bypass-tunnel [to ip-address] [protected-lsp [lsp-name]] [dynamic | manual]
        [detail]
      — interface [ip-int-name | ip-address] [label-map [label]]
      — interface [ip-int-name | ip-address] statistics
      — 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]
      — lsp [lsp-name] path [path-name] mbb
      — path [path-name] [lsp-binding]
      — static-lsp [lsp-name]
      — static-lsp [lsp-type]
      — static-lsp count
      — srlg-group [group-name]
      — status

show
  — router
    — rsvp
      — interface [ip-int-name | ip-address] statistics [detail]
      — neighbor [ip-address] [detail]
      — session [session-type] [from ip-address | to ip-address] lsp-name name] [status
        {up | down}] [detail]
      — statistics
      — status

```

Clear Commands

```

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

```

Debug Commands

```

debug
  — router
    — [no] mpls [lsp lsp-name] [sender source-address] [endpoint endpoint-address] [tunnel-id
      tunnel-id] [lsp-id lsp-id] [interface ip-int-name]
      — [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
      — [no] rsvp [lsp lsp-name] [sender sender-address] [endpoint endpoint-address] [tunnel-id
        tunnel-id] [lsp-id lsp-id] [interface ip-int-name]
        — [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
          — ack [detail]
          — no ack
          — all [detail]
          — no all
          — 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**

Command Descriptions

- [Configuration Commands \(MPLS\) on page 105](#)
- [Configuration Commands \(RSVP-TE\) on page 136](#)
- [Show Commands \(MPLS\) on page 149](#)
- [Show Commands \(RSVP\) on page 172](#)
- [Clear Commands on page 182](#)
- [Debug Commands on page 184](#)

Configuration Commands (MPLS)

- [Generic Commands on page 106](#)
- [Interface Commands on page 112](#)
- [Interface Label-Map Commands on page 115](#)
- [LSP Commands on page 117](#)
- [Primary and Secondary Path Commands on page 126](#)
- [LSP Path Commands on page 132](#)
- [Static LSP Commands on page 134](#)

Generic Commands

mpls

Syntax	[no] mpls
Context	config>router
Description	<p>This command creates the MPLS protocol instance and enables MPLS configuration. The MPLS protocol instance is not created by default, but once it is created, a no shutdown command is not required since MPLS is enabled automatically. The shutdown command administratively disables MPLS.</p> <p>The no form of this command deletes this MPLS protocol instance and all configuration parameters for this MPLS instance.</p> <p>MPLS must be shut down and all SDP bindings to LSPs removed before the MPLS instance can be deleted. If MPLS is not shut down, when the no mpls command is executed, a warning message on the console indicates that MPLS is still administratively up.</p>

shutdown

Syntax	[no] shutdown
Context	config>router>mpls config>router>mpls>interface config>router>mpls>if>label-map config>router>mpls>path config>router>mpls>static-lsp
Description	<p>The shutdown command administratively disables an entity. The operational state of the entity is disabled as well as the operational state of any entities contained within. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. Many objects must be shut down before they can be deleted. Many entities must be explicitly enabled using the no shutdown command.</p> <p>In the label-map context, all packets that match the specified <i>in-label</i> are dropped when the label map is shut down.</p> <p>In the path context, 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>

The **no** form of this command places the entity into an administratively enabled state. In the **mpls** and **mpls>interface** contexts, this triggers any LSPs that were previously defined under the associated context to come back up. In the **path** context, the **no** form of this command administratively enables the path and all LSPs—where the path is defined as a primary or a standby secondary path—are (re)established.

Default **mpls** — **no shutdown**
 interface — **shutdown**
 label-map — **no shutdown**
 path — **shutdown**
 static-lsp — **shutdown**

admin-group

Syntax **admin-group** *group-name group-value*
 no admin-group *group-name*

Context config>router>mpls

Description 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 they 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.

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.

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.

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.

The **no** form of this command deletes the admin group. All configuration information associated with this LSP is lost.

Default n/a

Parameters *group-name* — specifies the name of the admin group within a router instance
 group-value — specifies the group value associated with this admin group. This value is unique within a router instance.

Values 0 to 31

dynamic-bypass

Syntax	dynamic-bypass [enable disable]
Context	config>router>mpls
Description	This command disables the creation of dynamic bypass LSPs in FRR. One or more manual bypass LSPs must be configured to protect the primary LSP path at the PLR nodes.
Default	enable

frr-object

Syntax	[no] frr-object
Context	config>router>mpls
Description	This command specifies whether signaling the frr-object is on or off. 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 — by default, the value is inherited by all LSPs

hold-timer

Syntax	hold-timer <i>seconds</i> no hold-timer
Context	config>router>mpls
Description	<p>This command specifies the amount of time that the ingress node waits before programming its data plane and declaring to the service module that the LSP status is up.</p> <p>The no form of the command disables the hold-timer.</p>
Parameters	<i>seconds</i> — specifies the hold time, in seconds
Values	0 to 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, that the 7705 SAR OS software waits before attempting to resignal the LSPs.</p> <p>When the resignal timer expires, if the newly computed path for an LSP has a better metric than that for the currently recorded hop list, an attempt is made to resignal that LSP using the make-before-break (MBB) 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>When the resignal timer expires, a trap and syslog message are generated.</p> <p>The no form of the command disables timer-based LSP resignaling.</p>
Default	no resignal-timer
Parameters	<p><i>minutes</i> — specifies the time the software waits before attempting to resignal the LSPs, in minutes</p> <p>Values 30 to 10080</p>

srlg-frr

Syntax	srlg-frr [strict] no srlg-frr
Context	config>router>mpls
Description	<p>This system-wide command enables or disables the use of the shared risk link group (SRLG) constraint in the computation of an FRR bypass or detour LSP for any primary LSP path on the system. When srlg-frr is enabled, CSPF includes the SRLG constraint in the computation of an FRR bypass or detour LSP for protecting the primary LSP path.</p> <p>The strict option is a system-wide option that forces the CSPF to consider any configured SRLG membership lists in its calculation of every LSP path.</p> <p>CSPF and FRR</p>

CSPF prunes all links with interfaces that belong to the same SRLG as the interface being protected, where the interface being protected is the outgoing interface at the PLR used by the primary path. If one or more paths are found, the MPLS/RSVP-TE task selects one path based on best cost and signals the setup of the FRR bypass or detour LSP. If no path is found and the user included the **strict** option, the FRR bypass or detour LSP is not set up and the MPLS/RSVP-TE task keeps retrying the request to CSPF. If no path is found and the **strict** option is disabled, if a path exists that meets all the TE constraints except the SRLG constraint, then the FRR bypass or detour LSP is set up.

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

When the MPLS/RSVP-TE task is searching for an SRLG bypass tunnel to associate with the primary path of the protected LSP, the task does the following steps.

- First, the task checks for any configured manual bypass LSP that has CSPF enabled and that satisfies the SRLG constraints.
- The task skips any non-CSPF bypass LSP since there is no ERO returned with which to check the SRLG constraint.
- If no path is found, the task checks for an existing dynamic bypass LSP that satisfies the SRLG and other primary path constraints.
- If no bypass path is found, then the task makes a request to CSPF to try to create one.

Primary Path and FRR Behavior

Once the primary path of the LSP is set up and is operationally up, any subsequent changes to the SRLG membership of an interface that the primary path is using will not be considered by the MPLS/RSVP-TE task at the PLR for FRR bypass or detour LSP association until the next opportunity that the primary path is resigaled. The path may be resigaled due to a failure or to a make-before-break (MBB) operation. A make-before-break operation occurs as a result of a global revertive operation, a reoptimization of the LSP path (timer-based or manual), or a change by the user to any of the path constraints.

Once the FRR bypass or detour LSP is set up and is operationally up, any subsequent change to the SRLG membership of an interface that the FRR bypass or detour LSP is using will not be considered by the MPLS/RSVP-TE task at the PLR until the next opportunity that the association with the primary LSP path is rechecked. The association is rechecked if the FRR bypass or detour LSP is reoptimized. Detour routes are not reoptimized and are resigaled if the primary path is down.

The user must first shut down MPLS before enabling or disabling the **srlg-frr** option in CLI.

An RSVP-TE interface can belong to a maximum of 64 SRLGs. The user creates SRLGs using the **config>router>mpls>srlg-group** command. The user associates the SRLGs with an RSVP-TE interface using the **srlg-group** command in the **config>router>mpls>interface** context.

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

Default	no srlg-frr
Parameters	strict — specifies that the CSPF calculation for the FRR backup must include the SRLG constraint and the backup must be on the resulting list of eligible backup paths
Values	non-strict: srlg-frr strict: srlg-frr strict

srlg-group

Syntax	srlg-group <i>group-name</i> { value <i>group-value</i> } no srlg-group <i>group-name</i>
Context	config>router>mpls
Description	<p>This command is used to assign a name and a value to a shared risk link group (SRLG). An SRLG represents a set of interfaces (or links) that share the same risk of failing because they may be subjected to the same resource failures or defects.</p> <p>RSVP-TE interfaces must be explicitly assigned to an SRLG. SRLGs must be defined in the config>router>mpls context before they can be assigned to an RSVP-TE interface. Two different SRLG names cannot share the same <i>group-value</i>. Once an SRLG has been bound to an MPLS interface, its value cannot be changed until the binding is removed.</p> <p>The IGP communicates the information throughout the area using the TE link state advertisement. CSPF uses the information when paths are computed for constraint-based LSPs. CSPF must be enabled in order for SRLGs to be relevant.</p> <p>Up to 256 group names can be defined in the config>router>mpls context. SRLG names must be identical across all routers in a single domain. Up to five group names can be defined using one srlg-group command.</p> <p>The no form of this command deletes the SRLG.</p>
Default	n/a
Parameters	<p><i>group-name</i> — specifies the name of the SRLG within a router instance, up to 32 characters</p> <p><i>group-value</i> — specifies the group value associated with this SRLG; the group value is unique within a router instance</p> <p>Values 0 to 4294967295</p>

Interface Commands

interface

Syntax	[no] interface <i>ip-int-name</i>
Context	config>router>mpls
Description	<p>This command enables MPLS protocol support on an IP interface. MPLS commands are not executed on an IP interface where MPLS is not enabled.</p> <p>The no form of this command deletes all MPLS commands that are defined under the interface, such as label-map. The interface must be shut down before it can be deleted. If the interface is not shut down, the no interface <i>ip-int-name</i> command issues a warning message on the console indicating that the interface is administratively up.</p>
Default	shut down
Parameters	<i>ip-int-name</i> — identifies the network IP interface. The interface name character string cannot be in the form of an IP address. 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.

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 configuration. When 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	<i>group-name</i> — specifies the name of the group. The group names should be the same across all routers in the MPLS domain.

srlg-group

Syntax	[no] srlg-group <i>group-name</i> [<i>group-name...</i> (up to 5 max)]
Context	config>router>mpls>interface
Description	<p>This command associates an RSVP-TE interface with one or more SRLGs. An interface can belong to up to 64 SRLGs. Each operation of the srlg-group command allows a maximum of five groups to be specified at a time.</p> <p>The no form of this command deletes the association of the interface with the SRLG.</p>
Default	n/a
Parameters	<i>group-name</i> — specifies the group name of the SRLG within a router instance, up to 32 characters

shutdown

Syntax	[no] shutdown
Context	config>router>mpls>interface
Description	<p>This command disables the MPLS-related functions for the interface. The MPLS configuration information associated with this interface is retained. Shutting down the interface causes the LSPs associated with this interface to go down.</p> <p>The no form of this command administratively enables the MPLS interface. Any LSPs previously associated with this interface will attempt to come back up.</p>
Default	shutdown

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 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. The IS-IS-TE metric is encoded as sub-TLV type 18 as part of the extended IS reachability TLV. The metric value is encoded as a 24-bit unsigned integer.</p>

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. 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.

The TE metric in CSPF LSP path computation can be configured by entering the command **config>router>mpls>lsp *lsp-name*>cspf use-te-metric**.

The TE metric is only used in CSPF computations for MPLS paths and not in the regular SPF computation for IP reachability.

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

Default	no te-metric
Parameters	<i>value</i> — 1 to 16777215

Interface Label-Map Commands

label-map

Syntax	[no] label-map <i>in-label</i>
Context	config>router>mpls>interface
Description	<p>This command is used on either transit or egress LSP routers when a static LSP is defined. The static LSP on the ingress router is initiated using the config>router>mpls>static-lsp <i>lsp-name</i> command. The <i>in-label</i> is associated with a pop action or a swap action, but not both. If both actions are specified, the last action specified takes effect.</p> <p>The no form of this command deletes the static LSP configuration associated with the <i>in-label</i>.</p>
Parameters	<i>in-label</i> — specifies the incoming MPLS label on which to match
Values	32 to 1023

pop

Syntax	[no] pop
Context	config>router>mpls>if>label-map
Description	<p>This command specifies that the incoming label must be popped (removed). No label stacking is supported for a static LSP. The service header follows the top label. Once the label is popped, the packet is forwarded based on the service header.</p> <p>The no form of this command removes the pop action for the <i>in-label</i>.</p>
Default	n/a

swap

Syntax	swap <i>out-label</i> nexthop <i>ip-address</i> no swap
Context	config>router>mpls>interface>label-map
Description	<p>This command swaps the incoming label and specifies the outgoing label and next-hop IP address on an LSR for a static LSP.</p> <p>The no form of this command removes the swap action associated with the <i>in-label</i>.</p>
Default	n/a

Parameters *out-label* — specifies the label value to be swapped with the *in-label*. Label values 16 through 1048575 are defined as follows:

Label values 16 through 31 are 7705 SAR reserved

Label values 32 through 1023 are available for static assignment

Label values 1024 through 2047 are reserved for future use

Label values 2048 through 18431 are statically assigned for services

Label values 28672 through 131071 are dynamically assigned for both MPLS and services

Label values 131072 through 1048575 are reserved for future use

Values 16 to 1048575

ip-address — specifies the IP address to forward to. If an ARP entry for the next hop exists, then the static LSP will be marked operational. If an ARP entry does not exist, software will set the operational status of the static LSP to down and continue to ARP for the configured next-hop at a fixed interval.

LSP Commands

lsp

Syntax	[no] lsp <i>lsp-name</i> [bypass-only]
Context	config>router>mpls
Description	<p>This command creates an LSP that is signaled dynamically by the 7705 SAR OS.</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.</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 shut down and unbound from all SDPs before it can be deleted.</p>
Default	n/a
Parameters	<p><i>lsp-name</i> — specifies the 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 7705 SAR selects it. If no manual bypass tunnel is found, the 7705 SAR dynamically signals a bypass LSP as the default behavior. The CLI for this feature includes a command that provides the user with the option to disable dynamic bypass creation on a per-node basis.</p>

adaptive

Syntax	[no] adaptive
Context	config>router>mpls>lsp
Description	<p>This command enables the make-before-break (MBB) functionality for an LSP or LSP path. When enabled for the LSP, a make-before-break operation will be performed for the primary path and all the secondary paths of the LSP.</p>
Default	adaptive

adspec

Syntax	[no] adspec
Context	config>router>mpls>lsp
Description	When enabled, the advertised data (ADSPEC) object will be included in RSVP-TE messages.
Default	no adspec

cspf

Syntax	[no] cspf [use-te-metric]
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 LSPs that take configuration constraints into account. CSPF is also used to calculate the FRR bypass or detour LSP 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 is set up using RSVP-TE signaling from ingress to egress.</p> <p>If an LSP is configured with fast-reroute specified but does not enable CSPF, then neither global revertive nor local revertive will be available for the LSP to recover.</p>
Default	no cspf
Parameters	use-te-metric — specifies to use 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	<p>This command specifies the admin groups to be excluded when an LSP is set up in the primary or secondary contexts. Each single operation of the exclude command allows a maximum of 5 groups to be specified at a time. However, a maximum of 32 groups can be specified per LSP through multiple operations. The admin groups are defined in the config>router>mpls>admin-group context.</p> <p>Use the no form of the command to remove the exclude command.</p>
Default	no exclude
Parameters	<i>group-name</i> — specifies the existing group name to be excluded when an LSP is set up

fast-reroute

Syntax	[no] fast-reroute [<i>frr-method</i>]
Context	config>router>mpls>lsp
Description	This command creates a pre-computed protection LSP from each node in the path of the LSP. In case of a link or LSP failure between two nodes, traffic is immediately rerouted on the pre-computed protection LSP. When fast-reroute is specified, the default fast-reroute method is the facility method.

When **fast-reroute** is enabled, each node along the path of the LSP tries to establish a protection LSP as follows.

- Each upstream node sets up a protection LSP that avoids only the immediate downstream node, and merges back onto the actual path of the LSP as soon as possible.
- If it is not possible to set up a protection LSP that avoids the immediate downstream node, a protection LSP can be set up to the downstream node on a different interface.
- The protection LSP may take one or more hops (see [hop-limit](#)) before merging back onto 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-TE to set up their protection LSP. TE must be enabled for fast reroute to work.

Note that CSPF must be enabled for fast reroute to work. 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.

The one-to-one fast reroute method creates a separate detour LSP for each backed-up LSP.

The facility fast reroute method, sometimes called many-to-one, takes advantage of the MPLS label stack. Instead of creating a separate LSP for every backed-up LSP, a single LSP is created that serves to back up a set of LSPs. This LSP tunnel is called a bypass tunnel. The bypass tunnel must intersect the path of the original LSP(s) somewhere downstream of the point of local repair (PLR). This constrains the set of LSPs being backed up via that bypass tunnel to those LSPs that pass through a common downstream node. All LSPs that pass through the PLR and through this common node which do not also use the facilities involved in the bypass tunnel are candidates for this set of LSPs.

The **no** form of the **fast-reroute** command removes the protection 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.

Default	no fast-reroute
----------------	------------------------

Parameters *frr-method* — specifies the fast reroute method to use

Values one-to-one, facility

Default facility

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 protection 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**

Parameters *rate-in-mbps* — specifies the amount of bandwidth in Mb/s 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, this command defines how many more routers a protection tunnel is allowed to traverse compared with the LSP itself. For example, if an LSP traverses four routers, any protection tunnel for the LSP can be no more than 10 router hops, including the ingress and egress routers.

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

Default **16**

Parameters *limit* — specifies the maximum number of hops

Values 0 to 255

node-protect

Syntax	[no] node-protect
Context	config>router>mpls>lsp>fast-reroute
Description	<p>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.</p> <p>When node-protect is enabled, the 7705 SAR provides node protection on the specified LSP. If node protection cannot be provided, link protection is attempted. If link protection cannot be provided, there will be no protection.</p> <p>The no form of this command provides link protection. If link protection cannot be provided, there will be no protection.</p>
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 next-hop 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	system IP address				
Parameters	<p><i>ip-address</i> — specifies 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> <table> <tr> <td>Default</td><td>system IP address</td></tr> <tr> <td>Values</td><td>system IP or network interface IP addresses</td></tr> </table>	Default	system IP address	Values	system IP or network interface IP addresses
Default	system IP address				
Values	system IP or network interface IP addresses				

hop-limit

Syntax	hop-limit <i>number</i> no hop-limit
Context	config>router>mpls>lsp
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> <ul style="list-style-type: none"> • If the new value is less than the current number of hops of the established LSP, then the LSP is brought down. 7705 SAR OS software then tries to re-establish the LSP within the new hop-limit number. If the new value is equal to or greater than the current number of hops of the established LSP, then the LSP is not affected. <p>The no form of this command returns the parameter to the default value.</p>
Default	255
Parameters	<i>number</i> — specifies the number of hops the LSP can traverse, expressed as an integer
Values	2 to 255

include

Syntax	[no] include <i>group-name</i> [<i>group-name</i> ...(up to 5max)]
Context	config>router>mpls>lsp config>router>mpls>lsp>primary config>router>mpls>lsp>secondary
Description	<p>This command specifies the admin groups to be included when an LSP is set up. Up to 5 groups per operation can be specified, and up to 32 maximum.</p> <p>The no form of the command deletes the specified groups in the specified context.</p>
Default	no include
Parameters	<i>group-name</i> — specifies admin groups to be included when an LSP is set up

metric

Syntax	metric <i>metric</i>
Context	config>router>mpls>lsp
Description	This command specifies the metric for this LSP, which is used to select an LSP from among a set of LSPs that are destined for the same egress router. The LSP with the lowest metric will be selected.
Default	1
Parameters	<i>metric</i> — specifies the metric for this LSP
	Values 1 to 65535

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. 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 reverts the parameter to the default value.</p>
Default	0
Parameters	<p><i>number</i> — specifies the number of times that the 7705 SAR OS 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 to 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, between LSP re-establishment attempts after the LSP has failed.</p> <p>The no form of this command reverts to the default value.</p>
Default	30
Parameters	<i>seconds</i> — specifies the amount of time, in seconds, between attempts to re-establish the LSP after it has failed
Values	1 to 600

rsvp-resv-style

Syntax	rsvp-resv-style [se ff]
Context	config>router>mpls>lsp
Description	<p>This command specifies the RSVP-TE reservation style, shared explicit (se) or fixed filter (ff). 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	se
Parameters	<p>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.</p> <p>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.</p>

shutdown

Syntax	[no] shutdown
Context	config>router>mpls>lsp config>router>mpls>lsp>primary config>router>mpls>lsp>secondary
Description	<p>This lsp form of this command disables the existing LSP, including the primary and any standby secondary paths.</p> <p>The primary and secondary forms of this command administratively disables an LSP path and disables an existing LSP. Shutting down an LSP path does not change other configuration parameters for the LSP path.</p> <p>To shut down only the primary path enter the config>router>mpls>lsp <i>lsp-name</i>>primary <i>path-name</i>> shutdown command.</p> <p>To shut down a specific standby secondary path 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	lsp — shutdown primary — no shutdown secondary — no shutdown

to

Syntax	to <i>ip-address</i>
Context	config>router>mpls>lsp
Description	<p>This command specifies the system IP address of the egress router for the LSP. This command is mandatory to create an LSP.</p> <p>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.</p> <p>The to <i>ip-address</i> must be the system IP address of the egress router. If the to address does not match the SDP address, the LSP is not included in the SDP definition.</p>
Default	n/a
Parameters	<i>ip-address</i> — specifies the system IP address of the egress router

Primary and Secondary Path Commands

primary

Syntax	[no] primary <i>path-name</i>
Context	config>router>mpls>lsp
Description	<p>This command specifies a preferred path for the LSP. This command is optional only if the secondary path-name 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-name command override the comparable LSP attributes that are defined in the config>router>mpls>lsp context.</p> <p>The no form of this command deletes the association of this <i>path-name</i> from the lsp lsp-name. All configurations specific to this primary path, such as record, bandwidth, and hop limit, are deleted. The primary path must be shut down 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	n/a
Parameters	<i>path-name</i> — specifies 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 lsp-name>primary path-name command is specified. After the switchover from the primary path to the secondary path, the 7705 SAR OS 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 two secondary paths can be specified. Both secondary paths are considered equal, and the first available path is used. The 7705 SAR OS software will not switch back between secondary paths.</p> <p>The 7705 SAR OS software starts signaling 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>

The **no** form of this command removes the association between this *path-name* and *lsp-name*. All specific configurations for this association are deleted. The secondary path must be shut down first in 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	n/a
Parameters	<i>path-name</i> — specifies 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 (MBB) functionality for an LSP or a primary or secondary LSP path. When enabled for the LSP, a make-before-break operation will be performed for the primary path and all the secondary paths of the LSP.
Default	adaptive

bandwidth

Syntax	bandwidth <i>rate-in-mbps</i> 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> — specifies the amount of bandwidth reserved for the LSP path in Mb/s
Values	0 to 100000

exclude

Syntax	[no] exclude <i>group-name</i> [<i>group-name...</i> (up to 5 max)]
Context	config>router>mpls>lsp>primary config>router>mpls>lsp>secondary
Description	This command specifies the admin groups to be excluded when an LSP is set up. Up to 5 groups per operation can be specified, up to 32 maximum. The admin groups are defined in the config>router>mpls>admin-group context. Use the no form of the command to remove the exclude command.
Default	no exclude
Parameters	<i>group-name</i> — specifies the existing <i>group name</i> to be excluded when an LSP is set up

hop-limit

Syntax	hop-limit <i>number</i> no hop-limit
Context	config>router>mpls>lsp>primary config>router>mpls>lsp>secondary
Description	This optional command overrides the config>router>mpls>lsp <i>lsp-name</i> > 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: <ul style="list-style-type: none"> • If the new value is less than the current number of hops of the established LSP, then 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 to or greater 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 <i>lsp-name</i> > hop-limit command.
Default	no hop-limit
Parameters	<i>number</i> — specifies the number of hops the LSP can traverse, expressed as an integer Values 2 to 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

srlg

Syntax	[no] srlg
Context	config>router>mpls>lsp>secondary
Description	This command enables the use of the SRLG constraint in the CSPF computation of a secondary path for an LSP at the head-end LER. When this feature is enabled, CSPF includes the SRLG constraint in the computation of the secondary LSP path.

CSPF and SRLGs for Secondary Paths

CSPF requires that the primary LSP be established already and in the up state, since the head-end LER needs the most current ERO computed by CSPF for the primary path and CSPF includes the list of SRLGs in the ERO during the CSPF computation of the primary path. At a subsequent establishment of a secondary path with the SRLG constraint, the MPLS/RSVP-TE task queries CSPF again, which provides the list of SRLG numbers to be avoided. CSPF prunes all links with interfaces that belong to the same SRLGs as the interfaces included in the ERO of the primary path. If CSPF finds a path, the secondary path is set up. If CSPF does not find a path, MPLS/RSVP-TE keeps retrying the requests to CSPF.

If CSPF is not enabled on the LSP (using the **lsp lsp-name>cspf** command), then a secondary path of that LSP that includes the SRLG constraint is shut down and a specific failure code indicates the exact reason for the failure in the **show>router>mpls>lsp>path>detail** output.

Primary Path and Secondary Path Behavior

At initial primary LSP path establishment, if the primary path does not come up or is not configured, the SRLG secondary path is not signaled and is put in the down state. A specific failure code indicates the exact reason for the failure in the **show>router>mpls>lsp>path>detail** output. However, if a non-SRLG secondary path was configured, such as a secondary path with the SRLG option disabled, MPLS/RSVP-TE task signals it and the LSP uses it.

As soon as the primary path is configured and successfully established, MPLS/RSVP-TE moves the LSP to the primary path and signals all SRLG secondary paths.

Any time the primary path is reoptimized, has undergone a make-before-break (MBB) operation, or has come back up after being down, the MPLS/RSVP-TE task checks with CSPF to determine if the SRLG secondary path should be resignaled. If the MPLS/RSVP-TE task finds that the current secondary path is no longer SRLG disjoint — for example, the path became ineligible — it puts the path on a delayed make-before-break immediately after the expiry of the retry timer. If MBB fails on the first try, the secondary path is torn down and the path is put on retry.

At the next opportunity (that is, when the primary path goes down), the LSP uses of an eligible SRLG secondary path if the secondary path is in the up state. If all secondary eligible SRLG paths are in the down state, MPLS/RSVP-TE uses a non-SRLG secondary path if the path is configured and in the up state. If, while the LSP is using a non-SRLG secondary path, an eligible SRLG secondary path comes back up, MPLS/RSVP-TE will not switch the path of the LSP to it. As soon as the primary path is resignaled and comes up with a new SRLG list, MPLS/RSVP-TE resignals the secondary path using the new SRLG list.

A secondary path that becomes ineligible as a result of an update to the SRLG membership list of the primary path will have its ineligibility status removed when any of the following events occurs:

- A successful MBB operation of the standby SRLG path occurs, making it eligible again.
- The standby path goes down, in which case MPLS/RSVP-TE puts the standby on retry when the retry timer expires. If successful, it becomes eligible. If not successful after the retry timer expires or the number of retries reaches the configured retry-limit value, it is left down.
- The primary path goes down, in which case the ineligible secondary path is immediately torn down and will only be resigaled when the primary path comes back up with a new SRLG list.

Changes to SRLG Membership List

Once the primary path of the LSP is set up and is operationally up, any subsequent changes to the SRLG membership of an interface that the primary path is using is not considered until the next opportunity that the primary path is resigaled. The primary path may be resigaled due to a failure or to a make-before-break operation. A make-before-break operation occurs as a result of a global revertive operation, a timer-based or manual reoptimization of the LSP path, or a change by the user to any of the path constraints.

Once an SRLG secondary path is set up and is operationally up, any subsequent changes to the SRLG membership of an interface that the secondary path is using is not considered until the next opportunity that the secondary path is resigaled. The secondary path is resigaled due to a failure, to a resigaling of the primary path, or to a make-before-break operation. A make-before-break operation occurs as a result of a timer-based or manual reoptimization of the secondary path, or a change by the user to any of the path constraints of the secondary path, including enabling or disabling the SRLG constraint itself.

In addition, any user-configured **include** or **exclude** admin group statements for this secondary path are checked along with the SRLG constraints by CSPF.

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

Default **no srlg**

standby

Syntax **[no] standby**

Context config>router>mpls>lsp>secondary

Description The secondary path LSP is normally signaled if 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, the traffic is switched back to the primary path LSP.

The **no** form of this command specifies that the secondary LSP is signaled when the primary path LSP fails.

Default **n/a**

LSP Path Commands

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 the IGP (best effort) calculated shortest path to the egress router. Paths are created in a shutdown state. A path must be shut down before making any changes (adding or deleting hops) to the path. When a path is shut down, 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 shut down 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>
Default	n/a
Parameters	<i>path-name</i> — specifies the unique case-sensitive alphanumeric name label for the LSP path, up to 32 characters in length

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, 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>

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 shut down 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.

Default	n/a
Parameters	<p><i>hop-index</i> — specifies the hop index, which is used to order the specified hops. 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 to 1024</p> <p><i>ip-address</i> — specifies 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, 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>strict — 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 the 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> <p>loose — specifies that the route taken by the LSP from the previous hop to this hop can traverse other routers. Multiple hop entries with the same IP address are flagged as errors. Either the loose or strict keyword must be specified.</p>

Static LSP Commands

static-lsp

Syntax	[no] static-lsp <i>lsp-name</i>
Context	config>router>mpls
Description	<p>This command configures static LSPs on the ingress router. The static LSP is a manually configured LSP where the next-hop 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 shut down before it can be deleted. If the LSP is not shut down, the no static-lsp <i>lsp-name</i> command generates a warning message on the console indicating that the LSP is administratively up.</p>
Parameters	<i>lsp-name</i> — identifies the LSP. 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.

push

Syntax	push <i>label nexthop ip-address</i> no push <i>label</i>
Context	config>router>mpls>static-lsp
Description	<p>This command specifies the label to be pushed onto 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><i>label</i> — specifies the label to push on the label stack</p> <p>Label values 16 through 31 are 7705 SAR reserved Label values 32 through 1023 are available for static assignment Label values 1024 through 2047 are reserved for future use Label values 2048 through 18431 are statically assigned for services Label values 28672 through 131071 are dynamically assigned for both MPLS and services Label values 131072 through 1048575 are reserved for future use.</p> <p>Values 16 to 1048575</p> <p><i>ip-address</i> — 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 an ARP entry does not exist, the software sets the operational status of the static LSP to down and continues to send an ARP request for the configured next hop at fixed intervals.</p>

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. For LSPs that are used as transport tunnels for services, the to <i>ip-address</i> must be the system IP address. If the to <i>ip-address</i> does not match the SDP address, the LSP is not included in the SDP definition. This command is required when creating an LSP.
Parameters	<i>ip-address</i> — identifies the egress router system address
	Values a.b.c.d
Default	n/a

static-lsp-fast-retry

Syntax	static-lsp-fast-retry <i>seconds</i> no static-lsp-fast-retry
Context	config>router>mpls
Description	This command specifies the fast-retry timer that can be configured for static LSPs. When a static LSP is trying to come up, MPLS tries to resolve the ARP entry for the next hop of the LSP. If the next hop is still down or unavailable, the request may fail. In that case, MPLS starts a non-configurable timer of 30 seconds before making the next request. The fast-retry timer allows the user to configure a shorter retry timer so that the LSP comes up shortly after the next hop is available.
Parameters	<i>seconds</i> — fast-retry timer value, in seconds
	Values 1 to 30
Default	30

Configuration Commands (RSVP-TE)

- [Generic Commands on page 137](#)
- [Interface Commands on page 142](#)
- [Message Pacing Commands on page 148](#)

Generic Commands

rsvp

Syntax	[no] rsvp
Context	config>router
Description	<p>This command creates the RSVP-TE protocol instance and enables RSVP-TE configuration. RSVP-TE is enabled by default.</p> <p>RSVP-TE is used to set up LSPs. RSVP-TE should be enabled on all router interfaces that participate in signaled LSPs.</p> <p>The no form of this command deletes this RSVP-TE protocol instance and removes all configuration parameters for this RSVP-TE instance. To suspend the execution and maintain the existing configuration, use the shutdown command. RSVP-TE must be shut down before the RSVP-TE instance can be deleted. If RSVP-TE is not shut down, the no rsvp command does nothing except issue a warning message on the console indicating that RSVP-TE is still administratively enabled.</p>
Default	no shutdown

shutdown

Syntax	[no] shutdown
Context	config>router>rsvp config>router>rsvp>interface
Description	<p>This command disables the RSVP-TE protocol instance or the RSVP-related functions for the interface. The RSVP-TE configuration information associated with this interface is retained. When RSVP-TE is administratively disabled, all the RSVP-TE sessions are torn down.</p> <p>The no form of this command administratively enables RSVP-TE on the interface.</p>
Default	shutdown

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 (referred to as a maintenance interface) or all RSVP interfaces on the node (referred to as a maintenance node). When this command is executed, the node performs the following operations in no specific order.</p> <p>A PathErr message with an error sub-code of "Local Maintenance on TE Link required" is generated for each LSP that is in transit at this node and is using a maintenance interface as its outgoing interface. A PathErr message with the error code "Local node maintenance required" is generated if all interfaces are affected.</p> <p>A single make-before-break attempt is performed for all adaptive CSPF LSPs that originate on the node and whose paths make use of the maintenance interfaces listed in the PathErr message. If an alternative path for an affected LSP is not found, the LSP is maintained on its current path. The maintenance node also tears down and resignals any bypass or detour LSP that uses the maintenance interfaces as soon as they are not active. The maintenance node floods an IGP TE LSA/LSP containing a Link TLV for the links under graceful shutdown with the Traffic Engineering metric set to 0xffffffff and the Unreserved Bandwidth parameter set to zero (0).</p> <p>Upon receipt of the PathErr message, an intermediate LSR tears down and resignals any bypass LSP whose path makes use of the listed maintenance interfaces as soon as no associations with a protected LSP are active. The node does not take any action on a detour LSP whose path makes use of the listed maintenance interfaces.</p> <p>Upon receipt of the PathErr message, a head-end LER performs a single make-before-break attempt on the affected adaptive CSPF LSP. If an alternative path is not found, the LSP is maintained on its current path.</p> <p>A node does not take any action on the paths of the following originating LSPs after receiving the PathErr message:</p> <ul style="list-style-type: none">• 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 that can be found.• an adaptive CSPF LSP whose path has explicit hops defined using the listed maintenance interfaces or node• a CSPF LSP that has the adaptive option disabled and whose current path is over the listed maintenance interfaces in the PathErr message. These are not subject to make-before-break.• a non-CSPF LSP whose current path is over the listed maintenance interfaces in the PathErr message <p>Upon receipt of the updated IGP TE LSA/LSP for the maintenance interfaces, the head-end LER updates the TE database. This information will be used at the next scheduled CSPF computation for any LSP whose path might traverse any of the maintenance interfaces.</p>

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 n/a

keep-multiplier

Syntax	[no] keep-multiplier <i>number</i> no keep-multiplier
Context	config>router>rsvp
Description	The keep-multiplier <i>number</i> is an integer used by RSVP-TE to declare that a reservation is down or the neighbor is down. The keep-multiplier <i>number</i> is used with the refresh-time command to determine when RSVP-TE will declare the session down. The no form of this command reverts to the default value.
Default	3
Parameters	<i>number</i> — specifies the keep-multiplier value Values 1 to 255

rapid-retransmit-time

Syntax	rapid-retransmit-time <i>hundred-milliseconds</i> no rapid-retransmit-time
Context	config>router>rsvp
Description	This command is used to define the value of the rapid retransmission interval. This is used in the retransmission mechanism based on an exponential backoff timer in order to handle unacknowledged message- <i>_id</i> objects. The RSVP-TE message with the same message-id is retransmitted every $2 \times$ rapid-retransmit-time interval. The node will stop retransmission of unacknowledged RSVP-TE 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 rapid retransmission interval must be smaller than the regular refresh interval configured in config>router>rsvp>refresh-time . The no form of this command reverts to the default value.
Default	5 (which represents 500 msec)
Parameters	<i>hundred-milliseconds</i> — 1 to 100, in units of 100 msec

rapid-retry-limit

Syntax	rapid-retry-limit <i>number</i> no rapid-retry-limit
Context	config>router>rsvp
Description	<p>This command is used to define the value of the rapid retry limit. This is used in the retransmission mechanism based on an exponential backoff timer in order to handle unacknowledged message_id objects. The RSVP-TE message with the same message_id is retransmitted every $2 \times$ rapid-retransmit-time interval. The node will stop retransmission of unacknowledged RSVP-TE messages whenever the updated backoff interval exceeds the value of the regular refresh interval or the number of retransmissions reaches the value of the rapid-retry-limit parameter, whichever comes first.</p> <p>The no form of this command reverts to the default value.</p>
Default	3
Parameters	<i>number</i> — 1 to 6, integer values

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 7705 SAR PLR node or terminating on this 7705 SAR 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 that 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-TE messages pertaining to LSP paths tunneled over the bypass. It will also not send message-id in RSVP-TE messages. This effectively disables summary refresh.</p>
Default	disable

refresh-time

Syntax	refresh-time <i>seconds</i> no refresh-time
Context	config>router>rsvp
Description	<p>This command controls the interval, in seconds, between the successive PATH and RESV refresh messages. RSVP-TE declares the session down after it misses keep-multiplier <i>number</i> consecutive refresh messages.</p> <p>The no form of this command reverts to the default value.</p>
Default	30
Parameters	<i>seconds</i> — specifies the refresh time in seconds
	Values 1 to 65535

Interface Commands

interface

Syntax	[no] interface <i>ip-int-name</i>
Context	config>router>rsvp
Description	<p>This command enables RSVP-TE protocol support on an IP interface. No RSVP-TE commands are executed on an IP interface where RSVP-TE is not enabled.</p> <p>The no form of this command deletes all RSVP-TE commands such as hello-interval and subscription, which are defined for the interface. The RSVP-TE interface must be shut down before 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>
Parameters	<p><i>ip-int-name</i> — specifies the network IP interface. The interface name cannot be in the form of an IP address. 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>Values 1 to 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-TE neighbors to authenticate RSVP-TE messages. Authentication uses the MD5 message-based digest.</p> <p>When enabled on an RSVP-TE interface, authentication of RSVP-TE messages operates in both directions of the interface.</p> <p>A 7705 SAR 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• the key used with the authentication algorithm• the lifetime of the key; the user-entered key is valid until the user deletes it from the interface• the source address of the sending system• the latest sending sequence number used with this key identifier

A 7705 SAR RSVP-TE sender transmits an authenticating digest of the RSVP-TE message, computed using the shared authentication key and a keyed hash algorithm. The message digest is included in an integrity object that also contains a flags field, a key identifier field, and a sequence number field. The 7705 SAR RSVP-TE sender complies with the procedures for RSVP-TE message generation in RFC 2747, *RSVP Cryptographic Authentication*.

A 7705 SAR RSVP-TE receiver uses the key together with the authentication algorithm to process received RSVP-TE 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-TE messages sent over the bypass tunnel. If the PLR receives an RSVP-TE message with an integrity object, it will perform the digest verification for the key of the interface over which the packet was received. If this fails, the packet is dropped. If the received RSVP-TE message is an RESV message and does not have an integrity object, then the PLR node will accept it only if it originated from the MP node.

A 7705 SAR MP node will accept RSVP-TE messages received over the bypass tunnel with and without the integrity object. If an integrity object is present, the proper digest verification for the key of the interface over which the packet was received is performed. If this fails, the packet is dropped.

The 7705 SAR MD5 implementation does not support the authentication challenge procedures in RFC 2747.

The **no** form of this command disables authentication.

Default	no authentication-key — the authentication key value is the null string
Parameters	<p><i>authentication-key</i> — specifies the authentication key. The key can be any combination of ASCII characters up to 16 characters in length (unencrypted). If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.</p> <p><i>hash-key</i> — specifies 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 keyword is not used, the less-encrypted hash form is assumed.</p>

bfd-enable

Syntax	[no] bfd-enable
Context	config>router>rsvp>interface
Description	<p>This command enables the use of bidirectional forwarding (BFD) to control the state of the associated RSVP-TE interface. This causes RSVP-TE to register the interface with the BFD session on that interface.</p> <p>The user configures the BFD session parameters, such as transmit-interval, receive-interval, and multiplier, under the IP interface in the config>router> interface>bfd context.</p> <p>The BFD session on the interface might already have been started because of a prior registration with another protocol; for example, OSPF or IS-IS.</p> <p>The registration of an RSVP-TE interface with BFD is performed when a neighbor gets its first session, which means registration occurs when this node sends or receives a new PATH message over the interface. However, if the session did not come up due to not receiving an RESV for a new PATH message sent after the maximum number of retries, the LSP is shut down and the node deregisters with BFD. In general, the registration of RSVP-TE with BFD is removed as soon as the last RSVP-TE session is cleared.</p> <p>The registration of an RSVP-TE interface with BFD is performed independently of whether RSVP-TE hello is enabled on the interface or not. However, hello timeout clears all sessions toward the neighbor and RSVP-TE deregisters with BFD at the clearing of the last session.</p> <p>An RSVP-TE session is associated with a neighbor based on the interface address that the PATH message is sent to. If multiple interfaces exist to the same node, each interface is treated as a separate RSVP-TE neighbor. The user must enable BFD on each interface, and RSVP-TE will register with the BFD session running with each of those neighbors independently.</p> <p>Similarly, disabling BFD on the interface results in removing registration of the interface with BFD.</p> <p>When a BFD session transitions to the down state, the following actions are triggered. For RSVP-TE signaled LSPs, this triggers activation of FRR bypass or detour backup LSPs (PLR role), global revertive (head-end role), and switchover to secondary (if any) (head-end role) for affected LSPs with FRR enabled. It triggers a switchover to secondary (if any) and scheduling of retries for signaling the primary path of the non-FRR-affected LSPs (head-end role).</p> <p>The no form of this command removes BFD from the associated RSVP-TE protocol adjacency.</p>
Default	no bfd-enable

hello-interval

Syntax	hello-interval <i>milli-seconds</i> no hello-interval
Context	config>router>rsvp>interface
Description	<p>This command configures the time interval between RSVP-TE hello messages.</p> <p>RSVP-TE hello packets are used to detect loss of RSVP-TE connectivity with the neighboring node. Hello packets detect the loss of a neighbor more quickly than it would take for the RSVP-TE session to time out based on the refresh interval. After the loss of the of keep-multiplier <i>number</i> consecutive hello packets, the neighbor is declared to be in a down state.</p> <p>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.</p>
Default	3000
Parameters	<p><i>milli-seconds</i> — specifies the RSVP-TE hello interval in milliseconds, in multiples of 1000. A 0 (zero) value disables the sending of RSVP-TE hello messages.</p> <p>Values 0 to 60000 milliseconds (in multiples of 1000)</p>

refresh-reduction

Syntax	[no] refresh-reduction
Context	config>router>rsvp>interface
Description	This command enables the use of the RSVP-TE overhead refresh reduction capabilities on this RSVP-TE interface.

When this option is enabled, a 7705 SAR node will enable support for three capabilities:

- it will accept bundle RSVP-TE messages from its peer over this interface
- it will attempt to perform reliable RSVP-TE message delivery to its peer
- it will use summary refresh messages to refresh PATH and RESV states

The reliable message delivery must be explicitly enabled by the user after refresh reduction is enabled. The other two capabilities are enabled immediately.

A bundle RSVP-TE message is intended to reduce the overall message handling load. A bundle message consists of a bundle header followed by one or more bundle sub-messages. A sub-message can be any regular RSVP-TE message except another bundle message. A 7705 SAR node will only process received bundle RSVP-TE messages but will not generate them.

When reliable RSVP-TE message delivery is supported by both the node and its peer over the RSVP-TE interface, an RSVP-TE message is sent with a `message_id` object. A `message_id` object can be added to any RSVP-TE message when sent individually or as a sub-message of a bundle message.

If the sender sets the `ack_desired` flag in the `message_id` object, the receiver acknowledges the receipt of the RSVP-TE message by piggy-backing a `message_ack` object to the next RSVP-TE message it sends to its peer. Alternatively, an ACK message can also be used to send the `message_ack` object. In both cases, one or many `message_ack` objects could be included in the same message.

The 7705 SAR supports the sending of separate ACK messages only, but is capable of processing received `message_ack` objects piggy-backed to hop-by-hop RSVP-TE messages, such as PATH and RESV.

The 7705 SAR sets the `ack_desired` flag only in non-refresh RSVP-TE messages and in refresh messages that contain new state information.

A retransmission mechanism based on an exponential backoff timer is supported in order to handle unacknowledged `message_id` objects. The RSVP-TE message with the same `message_id` is retransmitted every $2 \times$ rapid-retransmit-time interval. The rapid-retransmit-time is referred to as the rapid retransmission interval because it must be smaller than the regular refresh interval configured in the **config>router>rsvp>refresh-time** context. There is also a maximum number of retransmissions of an unacknowledged RSVP-TE message **rapid-retry-limit**. The node will stop retransmission of unacknowledged RSVP-TE messages whenever the updated backoff interval exceeds the value of the regular **refresh-time** interval or the number of retransmissions reaches the value of the **rapid-retry-limit** parameter, whichever comes first. These two parameters are configurable globally on a system in the **config>router>rsvp** context.

Summary refresh consists of sending a summary refresh message containing a `message_id` list object. The fields of this object are populated each with the value of the `message_identifier` field in the `message_id` object of a previously sent individual PATH or RESV message. The summary refresh message is sent every refresh regular interval as configured by the user using the `refresh-time` command in the **config>router>rsvp** context. The receiver checks each `message_id` object against the saved PATH and RESV states. If a match is found, the state is updated as if a regular PATH or RESV refresh message was received from the peer. If a specific `message_identifier` field does not match, then the node sends a `message_id_nack` object to the originator of the message.

The above capabilities are referred to collectively as “refresh overhead reduction extensions”. When the refresh-reduction is enabled on a 7705 SAR RSVP-TE interface, the node indicates this to its peer by setting a “refresh-reduction-capable” bit in the flags field of the common RSVP-TE header. If both peers of an RSVP-TE interface set this bit, all the above three capabilities can be used. Furthermore, the node monitors the settings of this bit in received RSVP-TE messages from the peer on the interface. As soon as this bit is cleared, the 7705 SAR stops sending summary refresh messages. If a peer did not set the “refresh-reduction-capable” bit, a node does not attempt to send summary refresh messages.

However, if the peer did not set the “refresh-reduction-capable” bit, then a node with refresh reduction enabled and reliable message delivery enabled will still attempt to perform reliable message delivery with this peer. If the peer does not support the `message_id` object, it returns the error message “unknown object class”. In this case, the 7705 SAR node retransmits the RSVP-TE message without the `message_id` object and reverts to using this method for future messages destined for this 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	<p>This command enables reliable delivery of RSVP-TE messages over the RSVP-TE interface. When refresh-reduction is enabled on an interface and reliable-delivery is disabled, then the 7705 SAR will send a message_id and not set ACK desired in the RSVP-TE messages over the interface. Thus, the 7705 SAR does not expect an ACK 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-TE header, the node will enter summary refresh for a specific message_id it sent regardless of whether it received an ACK or not to this message from the neighbor.</p> <p>Finally, when the reliable-delivery option is enabled on any interface, RSVP-TE message pacing is disabled on all RSVP-TE interfaces of the system; for example, the user cannot enable the msg-pacing option in the config>router>rsvp context, and an 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 be generated in CLI after such an attempt.</p> <p>The no form of the command reverts to the default value.</p>
Default	no reliable-delivery

subscription

Syntax	subscription <i>percentage</i> no subscription
Context	config>router>rsvp>interface
Description	<p>This command configures the percentage of the link bandwidth that RSVP-TE can use for reservation and sets a limit for the amount of over-subscription or under-subscription allowed on the interface.</p> <p>When the subscription is set to zero, no new sessions are permitted on this interface. If the percentage is exceeded, the reservation is rejected and a log message is generated.</p> <p>The no form of this command reverts the percentage to the default value.</p>
Default	100
Parameters	<p><i>percentage</i> — specifies the percentage of the interface’s bandwidth that RSVP-TE allows to be used for reservations</p> <p>Values 0 to 1000</p>

Message Pacing Commands

msg-pacing

Syntax	[no] msg-pacing
Context	config>router>rsvp
Description	This command enables RSVP-TE message pacing, which is defined by the max-burst and period commands. 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-TE messages that can be sent under normal operating conditions, as specified by the period command. The no form of this command reverts to the default value.
Default	650
Parameters	<i>number</i> — maximum number of RSVP-TE messages Values 10 to 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, during which the router can send RSVP-TE messages, as specified by the max-burst command. The no form of this command reverts to the default value.
Default	100
Parameters	<i>milli-seconds</i> — the time interval during which the router can send RSVP-TE messages Values 10 to 1000 milliseconds, in increments of 10 milliseconds

Show Commands (MPLS)

admin-group

Syntax	admin-group <i>group-name</i>
Context	show>router>mpls
Description	This command displays MPLS administrative group information.
Parameters	<i>group-name</i> — specifies the administrative group name
Output	The following output is an example of MPLS administrative group information, and Table 8 describes the fields.

Sample Output

```
A:ALU-1# show router mpls admin-group
=====
MPLS Administrative Groups
=====
Group Name                               Group Value
-----
green                                    15
red                                      25
yellow                                  20
-----
No. of Groups: 3
=====
A:ALU-1#
```

Table 8: Show Router MPLS Admin-Group Output Fields

Label	Description
Group Name	The name of the administrative group. The name identifies the administrative group within a 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 administrative groups within the router instance

bypass-tunnel

Syntax	bypass-tunnel [to <i>ip-address</i>] [protected-lsp [<i>lsp-name</i>]] [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 that serves as a backup for a set of LSPs. This type of LSP tunnel is called a bypass tunnel.
Parameters	<p><i>ip-address</i> — specifies the IP address of the egress router</p> <p><i>lsp-name</i> — specifies the name of the LSP protected by the bypass tunnel</p> <p>dynamic — displays dynamically assigned labels for bypass protection</p> <p>manual — displays manually assigned labels for bypass protection</p> <p>detail — displays detailed information</p>
Output	The following output is an example of MPLS bypass tunnel information, and Table 9 describes the fields.

Sample Output

```

A:ALU-12>show>router>mpls# bypass-tunnel to 10.20.1.4
=====
Legend :  m - Manual                d - Dynamic
=====
To          State    Out I/F  Out Label  Reserved   Protected   Type
              BW (Kbps)   LSP Count
-----
10.20.1.4    Up        lag      *-*        131071      0
-----
Bypass Tunnels : 1
=====
A:ALU-12>show>router>mpls#

```

Table 9: Show Router MPLS Bypass-Tunnel Output Fields

Label	Description
To	The system IP address of the egress router
State	The LSP's administrative state
Out I/F	The name of the network IP interface
Out Label	The incoming MPLS label on which to match
Reserved BW (Kbps)	The amount of bandwidth in kilobytes per second (Kbps) reserved for the LSP

Table 9: Show Router MPLS Bypass-Tunnel Output Fields (Continued)

Label	Description
Protected LSP Count	The number of times this LSP has used a protected LSP
Type	The type of protected LSP

interface

Syntax	interface [<i>ip-int-name</i> <i>ip-address</i>] [label-map [<i>label</i>]] interface [<i>ip-int-name</i> <i>ip-address</i>] statistics
Context	show>router>mpls
Description	This command displays MPLS interface information.
Parameters	<p><i>ip-int-name</i> — identifies the network IP interface. The interface name cannot be in the form of an IP address. 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><i>ip-address</i> — specifies the system or network interface IP address</p> <p>label-map <i>label</i> — specifies the MPLS label on which to match</p> <p>Values 32 to 1023</p> <p>statistics — displays IP address and the number of packets and octets sent and received on an interface basis</p>
Output	The following output is an example of MPLS interface information, and Table 10 describes the fields.

Sample Output

```

ALU-12# show router mpls interface
=====
MPLS Interfaces
=====
Interface                Port-id      Adm      Opr      TE-metric
-----
system                   vport-1      Up       Up       None
  Admin Groups           None
  Srlg Groups            None
ip-10.10.1.2             1/1/1        Up       Up       None
  Admin Groups           None
  Srlg Groups            None
ip-10.10.4.2             1/1/2        Up       Up       None
  Admin Groups           None
  Srlg Groups            None
ip-10.10.3.2             1/1/3        Up       Up       None
  Admin Groups           None
  Srlg Groups            None

```

```

-----
Interfaces : 4
=====
ALU-12#

*A:ALU-48>config>router>mpls# show router mpls interface "to-104" label-map 35
=====
MPLS Interface : to-104 (Label-Map 35)
=====
In Label   In I/F      Out Label  Out I/F      Next Hop      Type      Adm   Opr
-----
35         1/1/1      n/a       n/a         n/a          Static    Up    Down

Interfaces : 1
=====
*A:ALU-48>config>router>mpls#

ALU-12# show router mpls interface statistics
=====
MPLS Interface (statistics)
=====
Interface   : ip-10.10.1.1
  Transmitted : Pkts - 6                      Octets - 540
  Received    : Pkts - 0                      Octets - 0
Interface   : ip-10.10.2.1
  Transmitted : Pkts - 0                      Octets - 0
  Received    : Pkts - 0                      Octets - 0
Interface   : ip-10.10.3.1
  Transmitted : Pkts - 0                      Octets - 0
  Received    : Pkts - 0                      Octets - 0
=====
ALU-12#

```

Table 10: Show Router MPLS Interface Output Fields

Label	Description
Interface	The interface name
Port-id	The port ID in the <i>slot/mda/port</i> format
Adm	The administrative state of the interface
Opr	The operational state of the interface
Te-metric	The traffic engineering metric used on the interface
Srlg Groups	The shared risk link group (SRLG)
Interfaces	The total number of interfaces
Transmitted	The number of packets and octets transmitted from the interface
Received	The number of packets and octets received

Table 10: Show Router MPLS Interface Output Fields (Continued)

Label	Description
In Label	The ingress label
In I/F	The ingress interface
Out Label	The egress label
Out I/F	The egress interface
Next Hop	The next-hop IP address for the static LSP
Type	Indicates whether the label value is statically or dynamically assigned

label

Syntax **label** *start-label* [*end-label* | **in-use** | *label-owner*]

Context show>router>mpls

Description This command displays MPLS labels exchanged.

Parameters *start-label* — specifies the label value assigned at the ingress router
end-label — specifies the label value assigned for the egress router
in-use — specifies the number of in-use labels displayed
label-owner — specifies the owner of the label

Values static, tldp

Output The following output is an example of MPLS label information, and [Table 11](#) describes the fields.

Sample Output

```
ALU-12# show router mpls label 32
=====
MPLS Label 32
=====
Label                Label Type          Label Owner
-----
32                   static-lsp         Not-in-use
-----
In-use labels in entire range : 7
=====
ALU-12#
```

Table 11: Show Router MPLS Label Output Fields

Label	Description
Label	The value of the label
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

label-range

Syntax **label-range**

Context show>router>mpls

Description This command displays the MPLS label range.

Output The following output is an example of MPLS label range information, and [Table 12](#) describes the fields.

Sample Output

```

ALU-12# show router mpls label-range
=====
Label Ranges
=====
Label Type      Start Label      End Label      Aging      Total Available
-----
static-lsp      32                1023          -           991
static-svc      2048              18431         -           16383
dynamic         32768             131071        0           98301
=====
ALU-12#

```

Table 12: Show Router MPLS Label Range Output Fields

Label	Description
Label Type	Displays 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

Table 12: Show Router MPLS Label Range Output Fields (Continued)

Label	Description
Aging	The number of labels released from a service that are transitioning back to the label pool. Labels are aged 15 seconds.
Total Available	The number of label values available

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-name* — specifies 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 LSPs that transit the router
terminate — displays the LSPs that terminate at the router
name — displays the IP address of the named LSP
count — displays the total number of LSPs
activepath — displays the present path being used to forward traffic
path-name — specifies the name of the path carrying the LSP
mbb — displays make-before-break (MBB) information
detail — displays detailed information

Output The following outputs are examples of MPLS LSP information:

- MPLS LSP ([Sample Output, Table 13](#))
- MPLS LSP Detail ([Sample Output, Table 14](#))
- MPLS LSP Path Detail ([Sample Output, Table 15](#))
- MPLS LSP Path MBB ([Sample Output, Table 16](#))

Sample Output

```
A:ALU-48# show router mpls lsp
=====
MPLS LSPs (Originating)
=====
```

LSP Name	To	Fastfail Config	Adm	Opr
to-104	10.10.10.104	Yes	Up	Up
to-103	0.0.0.0	Yes	Up	Up
to-99	10.10.10.99	No	Up	Up
to-100	10.10.10.100	No	Up	Up
to-49	10.20.30.49	No	Dwn	Up

```
-----
LSPs : 5
=====
A:ALU-48#

*A:ALU-48# show router mpls lsp to-104
=====
MPLS LSPs (Originating)
=====
```

LSP Name	To	Fastfail Config	Adm	Opr
to-104	10.10.10.104	Yes	Up	Dwn

```
-----
LSPs : 1
=====
*A:ALU-48#
```

Table 13: Show Router 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
FastFail Config	enabled — fast reroute is enabled. In the event of a failure, traffic is immediately rerouted on the pre-computed protection LSP, thus minimizing packet loss disabled — there is no protection LSP from each node on the primary path
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
LSPs	The total number of LSPs configured

Sample Output

```

*A:ALU-48# show router mpls lsp to-104 detail
=====
MPLS LSPs (Originating) (Detail)
-----
Type : Originating
-----
LSP Name      : to-104                      LSP Tunnel ID : 1
From          : 10.10.10.103                To           : 10.10.10.104
Adm State     : Up                          Oper State    : Down
LSP Up Time   : 0d 00:00:00                 LSP Down Time : 0d 00:46:50
Transitions   : 0                           Path Changes   : 0
Retry Limit   : 0                           Retry Timer    : 30 sec
Signaling     : RSVP                        Resv. Style    : FF
Hop Limit     : 10                          Negotiated MTU : 0
Adaptive      : Enabled
FastReroute   : Enabled                     Oper FR        : Disabled
FR Method     : Facility                     FR Hop Limit   : 16
FR Bandwidth  : 0 Mbps                       FR Node Protect: Enabled
FR Object     : Enabled
CSPF          : Enabled                      ADSPEC         : Enabled
Metric        : 1                           Use TE metric   : Disabled
Include Grps  :                               Exclude Grps    :
None                                                  None
Type          : RegularLsp

Secondary     : secondary-path                Down Time      : 0d 00:46:50
Bandwidth     : 50000 Mbps
Primary       : to-NYC                        Down Time      : 0d 00:46:50
Bandwidth     : 0 Mbps
=====

```

Table 14: Show Router MPLS LSP Detail Output Fields

Label	Description
LSP Name	The name of the LSP used in the path
From	The IP address of the ingress router for the LSP
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
LSP Up Time	The length of time the LSP has been operational
LSP Down Time	The total time in increments that the LSP path has not been operational
Transitions	The number of transitions that have occurred for the LSP
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 Limit	The number of attempts that the software should make to re-establish the LSP after it has failed
Retry Timer	The time, in seconds, for LSP re-establishment attempts after an LSP failure
Signaling	Specifies the signaling style
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.
	ff — specifies a shared reservation environment with an explicit reservation scope. Specifies an explicit list of senders and a distinct reservation for each of them.
Hop Limit	The maximum number of hops that an LSP can traverse, including the ingress and egress routers
Negotiated MTU	The size of the maximum transmission unit (MTU) that is negotiated during establishment of the LSP
Adaptive	Indicates whether make-before-break is enabled or disabled for resigaled paths

Table 14: Show Router MPLS LSP Detail Output Fields (Continued)

Label	Description
Fast Reroute	Enabled — fast reroute is enabled. In the event of a failure, traffic is immediately rerouted on the pre-computed protection LSP, thus minimizing packet loss.
	Disabled — there is no protection LSP from each node on the primary path
Oper FR	Indicates whether FRR has been enabled or disabled
FR Method	The type of Fast Reroute (FRR) that is used by the path
FR Hop Limit	The total number of hops a protection LSP can take before merging back onto the main LSP path
FR Bandwidth	The amount of bandwidth reserved for fast reroute
FR Node Protect	Indicates whether FRR has node protection enabled or disabled
FR Object	Indicates whether signaling the fr-object is on or off
CSPF	Indicates whether CSPF has been enabled or disabled
ADSPEC	enabled — the LSP will include advertising data (ADSPEC) objects in RSVP-TE messages
	disabled — the LSP will not include advertising data (ADSPEC) objects in RSVP-TE messages
Metric	The TE metric value
Use TE metric	Indicates whether the use of the TE metric is enabled or disabled
Include Grps	The admin groups that are to be included by an LSP when signaling a path
Exclude Grps	The admin groups that are to be avoided by an LSP when signaling a path
Type	The type of LSP
Secondary	The alternate path that the LSP will use if the primary path is not available
Down Time	The length of time that the path has been down
Bandwidth	The amount of bandwidth in megabits per second (Mbps) reserved for the LSP path
Primary	The preferred path for the LSP

Sample Output

```
*A:ALU-48# show router mpls lsp path detail
=====
MPLS LSP Path (Detail)
=====
Legend :
    @ - Detour Available          # - Detour In Use
    b - Bandwidth Protected       n - Node Protected
=====
LSP 1 Path 1
-----
LSP Name      : 1                      Path LSP ID : 30226
From          : 10.20.1.1              To          : 10.20.1.2
Adm State     : Up                     Oper State  : Up
Path Name     : 1                      Path Type   : Primary
Path Admin    : Up                     Path Oper   : Up
OutInterface  : 1/1/1                  Out Label   : 131071
Path Up Time  : 0d 00:59:39            Path Dn Time: 0d 00:00:00
Retry Limit   : 20                     Retry Timer : 30 sec
RetryAttempt  : 0                      Next Retry  *: 0 sec
Bandwidth     : 200 Mbps                Oper Bandwi*: 50 Mbps
Hop Limit     : 255
Record Route  : Record                  Record Label: Record
Oper MTU      : 1500                    Neg MTU     : 1500
Adaptive      : Enabled
Include Grps  :                         Exclude Grps:
None                                                  None
Path Trans    : 9                       CSPF Queries: 205
Failure Code  : noError                  Failure Node: n/a
ExplicitHops  :
    No Hops Specified
Actual Hops   :
    10.10.1.1(10.20.1.1)                Record Label : N/A
    -> 10.10.1.2(10.20.1.2)            Record Label : 131071
ComputedHops  :
    10.10.1.1 -> 10.10.1.2
LastResignalAttempt: 2008/04/08 11:42:33.22 PST Metric : 1000

Last MBB:
MBB Type      : Timer-based Resignal    MBB State   : Success/Failed
Ended at     : 2008/04/08 11:12:23.76 PST Old Metric  : 3000

In Progress MBB:
MBB Type      : Config Change           NextRetryIn : 16 sec
Started at   : 2008/04/08 12:01:02.20 PST RetryAttempt: 3
Failure Code : noCspfRouteToDestination Failure Node: 10.20.1.1
=====
*A:ALU-48#
```

Table 15: Show Router MPLS LSP Path Detail Output Fields

Label	Description
LSP Name	The name of the LSP used in the path
Path LSP ID	The LSP ID for the path
From	The IP address of the ingress router for the LSP
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
Path Name	The alphanumeric name of the path
Path Type	The type of path: primary or secondary
Path Admin	The administrative status of the path
Path Oper	The operational status of the path
OutInterface	The output interface of the LSP
Out Label	The output label of the LSP
Path Up Time	The length of time that the path has been operationally up
Path Down Time	The length of time that the path has been operationally down
Retry Limit	The number of times an LSP will retry before giving up completely
Retry Timer	The length of time between LSP signaling attempts
Retry Attempt	The number of attempts that have been made to re-establish the LSP
Next Retry	The time when the next attempt to re-establish the LSP will occur
Bandwidth	The amount of bandwidth in megabits per second (Mbps) reserved for the LSP path
Oper Bandwidth	The bandwidth reserved by the LSP
Hop Limit	The limit on the number of hops taken by the LSP
Record Route	Indicates whether a list of routers for the LSP has been recorded
Record Label	Indicates whether a list of router labels has been recorded
Oper MTU	The operational MTU of the connection to the next hop

Table 15: Show Router MPLS LSP Path Detail Output Fields (Continued)

Label	Description
Neg MTU	The MTU negotiated between the router and its next hop
Adaptive	Indicates whether make-before-break is enabled or disabled for resigaled paths
Include Grps	The admin groups that are to be included by an LSP when signaling a path
Exclude Grps	The admin groups that are to be avoided by an LSP when signaling a path
Path Trans	The number of times a path has made a transition between up and down states
CSPF Queries	The number of requests made by the LSP to the TE database
Failure Code	The reason code for in-progress MBB failure. A value of none indicates that no failure has occurred.
Failure Node	The IP address of the node in the LSP path at which the in-progress MBB failed. If no failure has occurred, this value is none .
Explicit Hops	The hops that have been specified by the user
Actual Hops	The hops that the route has taken
Record Label	The label recorded at the given hop
Computed Hops	The hops computed and returned from the routing database
LastResignalAttempt	The system up time when the last attempt to resignal this LSP was made
Last Resignal	The last time the route was resigaled
Metric	The value of the metric
Last MBB	Header for the last make-before-break (MBB) information
MBB Type	An enumerated integer that specifies the type of make-before-break (MBB) operation. If none displays, then there is no MBB in progress or no last MBB.
MBB State	The state of the most recent invocation of the make-before-break functionality
Ended at	The system up time when the last MBB ended
Old Metric	The cost of the traffic engineered path for the LSP path prior to MBB
In Progress MBB	Header for the currently in-progress MBB information

Table 15: Show Router MPLS LSP Path Detail Output Fields (Continued)

Label	Description
MBB Type	An enumerated integer that specifies the type of make-before-break (MBB) operation. If none displays, then there is no MBB in progress or no last MBB.
NextRetryIn	The amount of time remaining, in seconds, before the next attempt is made to retry the in-progress MBB
Started At	The time the current MBB began
RetryAttempt	The number of attempts for the MBB in progress
Failure Code	The reason code for in-progress MBB failure. A value of none indicates that no failure has occurred.
Failure Node	The IP address of the node in the LSP path at which the in-progress MBB failed. If no failure has occurred, this value is none .

Sample Output

```
*A:sim1# show router mpls lsp path mbb
=====
MPLS LSP Path MBB
=====
LSP 1 Path 1
-----
LastResignalAttempt: 2008/04/08 11:42:33.22 PST   CSPF Metric   : 0

Last MBB:
MBB Type      : Timer-based Resignal                MBB State      : Success/Failed
Ended at      : 2008/04/08 11:12:23.76 PST          Old Metric     : 3000

In Progress MBB:
MBB Type      : Config Change                      NextRetryIn    : 16 sec
Started at    : 2008/04/08 12:01:02.20 PST          RetryAttempt   : 3
Failure Code  : noCspfRouteToDestination           Failure Node    : 10.20.1.1

-----
LSP 2 Path 1
-----
LastResignalAttempt: 2008/04/08 11:42:33.54 PST   CSPF Metric     : 0

Last MBB:
MBB Type      : Timer-based Resignal                MBB State      : Success/Failed
Ended at      : 2008/04/08 11:12:24.76 PST          Old Metric     : 2000
```

```

-----
LSP 4 Path 1
-----
LastResignalAttempt: 2008/04/08 11:42:34.12 PST   CSPF Metric   : 0

In Progress MBB:
MBB Type      : Global Revertive                    NextRetryIn : 10 sec
Started at    : 2008/04/08 11:45:02.20 PST          RetryAttempt: 2
Failure Code: noCspfRouteToDestination              Failure Node: 10.20.1.1
=====
*A:sim1#

```

Table 16: Show Router MPLS LSP Path MBB Output Fields

Label	Description
LastResignalAttempt	The system up time when the last attempt to resignal this LSP was made
CSPF Metric	The value of the CSPF metric
Last MBB	Header for the last make-before-break (MBB) information
MBB Type	An enumerated integer that specifies the type of make-before-break (MBB) operation. If none displays, then there is no MBB in progress or no last MBB.
MBB State	The state of the most recent invocation of the make-before-break functionality
Ended at	The system up time when the last MBB ended
Old Metric	The cost of the traffic-engineered path for the LSP path prior to MBB
In Progress MBB	Header for the currently in-progress MBB information
MBB Type	An enumerated integer that specifies the type of make-before-break (MBB) operation. If none displays, then there is no MBB in progress or no last MBB.
NextRetryIn	The amount of time remaining, in seconds, before the next attempt is made to retry the in-progress MBB
Started At	The time that the current MBB began
RetryAttempt	The number of attempts for the MBB in progress
Failure Code	The reason code for in-progress MBB failure. A value of none indicates that no failure has occurred.
Failure Node	The IP address of the node in the LSP path at which the in-progress MBB failed. When no failure has occurred, this value is none .

path

Syntax	path [<i>path-name</i>] [lsp-binding]
Context	show>router>mpls
Description	This command displays MPLS paths.
Parameters	<i>path-name</i> — the unique name label for the LSP path lsp-binding — displays binding information
Output	The following output is an example of MPLS path information, and Table 17 describes the fields.

Sample Output

```

A:ALU-12# show router mpls path
=====
MPLS Path:
=====
Path Name                               Adm  Hop Index   IP Address      Strict/Loose
-----
nyc_to_sjc_via_dfw                      Up   20          100.20.1.4      Strict
                                           30          100.20.1.6      Strict
                                           40          100.20.1.8      Strict
                                           50          100.20.1.10     Strict

nyc_to_sjc_via_den                      Up   10          100.20.1.5      Strict
                                           20          100.20.1.7      Loose
                                           30          100.20.1.9      Loose
                                           40          100.20.1.11     Loose
                                           50          100.20.1.13     Strict

secondary_path2                         Down no hops    n/a             n/a
-----
Paths : 3
=====
A:ALU-12#

A:ALU-12# show router mpls path lsp-binding
=====
MPLS Path:
=====
Path Name                               Opr   LSP Name      Binding
-----
nyc_to_sjc_via_dfw                      Up    NYC_SJC_customer1  Primary
nyc_to_sjc_via_den                      Up    NYC_SJC_customer1  Standby
secondary_path2                         Down  NYC_SJC_customer1  Seconda*
-----
Paths : 3
=====
A:ALU-12#

```

Table 17: Show Router 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 other routers
Opr	The operational status of the path (up or down)
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

srlg-group

Syntax	srlg-group [<i>group-name</i>]
Context	show>router>mpls
Description	This command displays MPLS shared risk link groups (SRLGs)
Parameters	<i>group-name</i> — specifies the name of the SRLG within a router instance.
Output	The following output is an example of MPLS SRLG group information, and Table 18 describes the fields.

Sample Output

```
*A:ALU-48>show>router>mpls# srlg-group test2
=====
MPLS Srlg Groups
=====
Group Name                Group Value  Interfaces
-----
test2                     2           to-104
-----
No. of Groups: 1
=====
*A:ALU-48>show>router>mpls#
```

Table 18: Show Router MPLS SRLG Group Output Fields

Label	Description
Group Name	The name of the SRLG group within a router instance
Group Value	The group value associated with this SRLG group
Interfaces	The interface where the SRLG group is associated
No. of Groups	The total number of SRLG groups associated with the output

static-lsp

Syntax

static-lsp [*lsp-name*]
static-lsp [*lsp-type*]
static-lsp count

Context

show>router>mpls

Description

This command displays MPLS static LSP information.

Parameters

lsp-name — name that identifies the LSP. The LSP name can be up to 32 characters long and must be unique.

lsp-type — type that identifies the LSP. The LSP type is one of the keywords **transit** or **terminate**, where **terminate** displays the number of static LSPs that terminate at the router, and **transit** displays the number of static LSPs that transit the router.

count — the number of static LSPs that originate and terminate at the router

Output

The following output is an example of MPLS static LSP information, and [Table 19](#) describes the fields.

Sample Output - static-lsp

```
ALU-12# show router mpls static-lsp
=====
MPLS Static LSPs (Originating)
=====
LSP Name      To          Next Hop      Out Label Up/Down Time  Adm  Opr
ID            Out Port
-----
to131         10.9.9.9     10.1.2.2      131      30d 02:42:53  Up   Down
1             n/a
to121         10.8.8.8     10.1.3.2      121      30d 02:42:53  Up   Down
2             n/a
static-lsp_- 10.9.9.9     10.1.2.2      35        0d 01:39:34  Up   Down
cc
3             n/a
-----
LSPs : 3
=====
*A:ANOTHERNAME>show>router>mpls#
```

Sample Output - static-lsp transit

```
A:ALU-12# show router mpls static-lsp transit
=====
MPLS Static LSPs (Transit)
=====
In Label      In I/F      Out Label     Out I/F      Next Hop      Adm  Opr
-----
1020          1/1/1       1021          1/1/5        10.10.10.6    Up   Up
-----
LSPs : 1
=====
```

Sample Output - static-lsp terminate

```
*A:ALU-12>show>router>mpls# static-lsp terminate
=====
MPLS Static LSPs (Terminate)
=====
In Label      In Port      Out Label     Out Port      Next Hop      Adm  Opr
-----
131           1/3/1        n/a           n/a           n/a           Up   Down
121           1/2/1        n/a           n/a           n/a           Up   Down
35            1/3/1        n/a           n/a           n/a           Up   Down
-----
LSPs : 3
=====
```

Sample Output - static-lsp count

```
*A:ALU-12>show>router>mpls# static-lsp count
=====
MPLS Static-LSP Count
=====
Originate          Transit           Terminate
-----
0                  0                0
=====
*A:ALU-12>show>router>mpls# static-lsp
```

Table 19: Show Router 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
Out Label	The egress label
Adm	Down — indicates that the path is administratively disabled
	Up — indicates that the path is administratively enabled
Opr	Down — indicates that the path is operationally down
	Up — indicates that the path is operationally up
LSPs	The total number of static LSPs
In Label	The ingress label
In Port	The ingress port
Out Port	The egress port
Up/Down Time	The duration that the LSP is either operationally up or down
Static-LSP Count	The number of originating, transit, and terminating static LSPs

status

Syntax	status
Context	show>router>mpls
Description	This command displays MPLS operation information.
Output	The following output is an example of MPLS status information, and Table 20 describes the fields.

Sample Output

```

A:ALU-48>show router mpls status
=====
MPLS Status
=====
Admin Status      : Up                Oper Status      : Up
Oper Down Reason  : n/a
FR Object         : Enabled           Resignal Timer   : Disabled
Hold Timer        : 1 seconds         Next Resignal    : N/A
Srlg Frr          : Disabled          Srlg Frr Strict  : Disabled
Dynamic Bypass    : Enabled

LSP Counts        Originate        Transit          Terminate
-----
Static LSPs       0                  0                0
Dynamic LSPs      0                  0                0
Detour LSPs       0                  0                0
=====
A:ALU-48>config>router>mpls#

```

Table 20: Show Router MPLS Status Output Fields

Label	Description
Admin Status	Down — indicates that MPLS is administratively disabled
	Up — indicates that MPLS is administratively enabled
Oper Status	Down — indicates that MPLS is operationally down
	Up — indicates that 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

Table 20: Show Router MPLS Status Output Fields (Continued)

Label	Description
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	The amount of time that the ingress node holds before programming its data plane and declaring the LSP up to the service module
Oper Down Reason	The reason that MPLS is operationally down
Next Resignal	The amount of time until the next resignal for the LSP
Dynamic Bypass	Indicates whether dynamic bypass is enabled or disabled
LSP Counts	The number of originate, transit, and terminate LSPs that are static, dynamic, or detour

Show Commands (RSVP)

interface

Syntax	interface [<i>ip-int-name</i> <i>ip-address</i>] statistics [detail]
Context	show>router>rsvp
Description	<p>This command shows RSVP-TE interface information.</p> <p><i>ip-int-name</i> — identifies the network IP interface. The interface name cannot be in the form of an IP address. 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><i>ip-address</i> — the system or network interface IP address</p> <p>statistics — the IP address and the number of packets sent and received on an per-interface basis</p> <p>detail — displays detailed information</p>
Output	<p>The following outputs are examples of RSVP-TE interface information:</p> <ul style="list-style-type: none"> • RSVP-TE Interface (Sample Output, Table 21) • RSVP-TE Interface Detail (Sample Output, Table 22) • RSVP-TE Interface Statistics (Sample Output, Table 23)

Sample Output

```
A:ALU-12# show router rsvp interface
=====
RSVP Interfaces
=====
Interface                               Total    Active    Total BW  Resv BW  Adm Opr
Sessions Sessions (Mbps)    (Mbps)
-----
system                                 -         -         -         -         Up   Up
ip-10.10.1.1                          1         1        100         0         Up   Up
ip-10.10.2.1                          1         1        100         0         Up   Up
ip-10.10.3.1                          0         0        100         0         Up   Up
-----
Interfaces : 4
=====
A:ALU-12#
```

Table 21: Show Router RSVP-TE Interface Output Fields

Label	Description
Interface	The name of the IP interface
Total Sessions	The total number of RSVP-TE 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-TE sessions on this interface
Total BW (Mbps)	The amount of bandwidth in megabits per second (Mbps) available to be reserved for the RSVP-TE protocol on the interface
Resv BW (Mbps)	The amount of bandwidth in megabits per second (Mbps) reserved on this interface. A value of zero (0) indicates that no bandwidth is reserved.
Adm	Down — the RSVP-TE interface is administratively disabled
	Up — the RSVP-TE interface is administratively enabled
Opr	Down — the RSVP-TE interface is operationally down
	Up — the RSVP-TE interface is operationally up
Interfaces	The number of interfaces listed in the display

Sample Output

```

A: ALU-12# show router rsvp interface detail
=====
RSVP Interfaces (Detailed)
-----
Interface : system
-----
Interface      : system                      Port ID      : system
Admin State    : Up                        Oper State    : Up
Active Sessions: 0                        Active Resvs  : 0
Total Sessions : 0
Subscription   : 100 %
Unreserved BW  : 0 Mbps
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
No Neighbors.

Interface      : system                      Port ID      : system
Port Speed     : 0 Mbps
Reserved BW    : 0 Mbps
Aggregate      : Dsabl
Hello Timeouts : 0
Bfd Enabled    : Yes
Auth Key Id    : n/a
Auth Win Size  : n/a
Reliable Deli. : Disabled

A: ALU-12#

```

Table 22: Show Router RSVP-TE Interface Detail Output Fields

Label	Description
Interface	The name of the network IP interface
Port ID	The physical port bound to the interface
Admin State	Down — the RSVP-TE interface is administratively disabled
	Up — the RSVP-TE interface is administratively enabled
Oper State	Down — the RSVP-TE interface is operationally down
	Up — the RSVP-TE interface is operationally up
Active Sessions	The total number of active RSVP-TE sessions on this interface
Active Resvs	The total number of active RSVP-TE sessions that have reserved bandwidth
Total Sessions	The total number of RSVP-TE 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.
Subscription	The percentage of the link bandwidth that RSVP-TE can use for reservation. When the value is zero (0), no new sessions are permitted on this interface.
Port Speed	The speed for the interface
Unreserved BW	The amount of unreserved bandwidth
Reserved BW	The amount of bandwidth in megabits per second (Mbps) reserved by the RSVP-TE session on this interface. A value of zero (0) indicates that no bandwidth is reserved.
Total BW	The amount of bandwidth in megabits per second (Mbps) available to be reserved for the RSVP-TE protocol on this interface
Hello Interval	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.
Hello Timeouts	The total number of hello messages that timed out on this RSVP-TE interface
Authentication	Enabled — MD5 authentication is enabled
	Disabled — MD5 authentication is disabled

Table 22: Show Router RSVP-TE Interface Detail Output Fields (Continued)

Label	Description
Bfd Enabled	Yes — BFD is enabled on the RSVP-TE interface
	No — BFD is disabled on the RSVP-TE interface
Auth Rx Seq Num	The received MD5 sequence number
Auth Key Id	The MD5 key identifier
Auth Tx Seq Num	The transmitted MD5 sequence number
Auth Win Size	The MD5 window size
Refresh Reduc.	Enabled — refresh reduction capabilities are enabled
	Disabled — refresh reduction capabilities are disabled
Reliable Deli.	Enabled — reliable delivery is enabled
	Disabled — reliable delivery is disabled
Bfd Enabled	Yes — BFD is enabled on the RSVP-TE interface
	No — BFD is disabled on the RSVP-TE interface
No. of Neighbors	The IP addresses of the RSVP-TE neighbors

Sample Output

```

A:ALU-12# show router rsvp interface statistics
=====
RSVP Interface (statistics)
=====
Interface system
-----
Interface                : Up
Total Packets            (Sent) : 0                (Recd.): 0
Bad Packets              (Sent) : 0                (Recd.): 0
Paths                   (Sent) : 0                (Recd.): 0
Path Errors              (Sent) : 0                (Recd.): 0
Path Tears               (Sent) : 0                (Recd.): 0
Resvs                   (Sent) : 0                (Recd.): 0
Resv Confirms            (Sent) : 0                (Recd.): 0
Resv Errors              (Sent) : 0                (Recd.): 0
Resv Tears               (Sent) : 0                (Recd.): 0
Refresh Summaries        (Sent) : 0                (Recd.): 0
Refresh Acks             (Sent) : 0                (Recd.): 0
Bundle Packets           (Sent) : 0                (Recd.): 0
Hellos                   (Sent) : 0                (Recd.): 0
Auth Errors              (Sent) : 0                (Recd.): 0
-----

```

Table 23: Show Router RSVP-TE Interface Statistics Output Fields

Label	Description
Interface	The name of the IP interface displayed in the header
Interface (status)	The status of the interface (up or down)
Sent	The total number of error-free RSVP-TE packets that have been transmitted on the RSVP-TE interface
Recd	The total number of error-free RSVP-TE packets received on the RSVP-TE interface
Total Packets	The total number of RSVP-TE packets, including errors, received on the RSVP-TE interface
Bad Packets	The total number of RSVP-TE packets with errors transmitted on the RSVP-TE interface
Paths	The total number of RSVP-TE PATH messages received on the RSVP-TE interface
Path Errors	The total number of RSVP-TE PATH ERROR messages transmitted on the RSVP-TE interface
Path Tears	The total number of RSVP-TE PATH TEAR messages received on the RSVP-TE interface
Resvs	The total number of RSVP-TE RESV messages received on the RSVP-TE interface
Resv Confirms	The total number of RSVP-TE RESV CONFIRM messages received on the RSVP-TE interface
Resv Errors	The total number of RSVP-TE RESV ERROR messages received on the RSVP-TE interface
Resv Tears	The total number of RSVP-TE RESV TEAR messages received on the RSVP-TE interface
Refresh Summaries	The total number of RSVP-TE RESV summary refresh messages received on the RSVP-TE interface
Refresh Acks	The total number of RSVP-TE RESV acknowledgement messages received when refresh reduction is enabled on the RSVP-TE interface
Bundle Packets	The total number of RSVP-TE RESV bundle packets received on the RSVP-TE interface

Table 23: Show Router RSVP-TE Interface Statistics Output Fields (Continued)

Label	Description
Hellos	The total number of RSVP-TE RESV HELLO REQ messages received on the RSVP-TE interface
Auth Errors	The number of authentication errors

neighbor

Syntax	neighbor [<i>ip-address</i>] [detail]
Context	show>router>rsvp
Description	This command displays RSVP-TE neighbors.
Parameters	<i>ip-address</i> — the IP address of the originating router detail — displays detailed information
Output	The following output is an example of RSVP-TE neighbor information, and Table 24 describes the fields.

Sample Output

```
*A:ALU-12>show>router>rsvp# neighbor
=====
RSVP Neighbors
=====
Legend :
      LR - Local Refresh Reduction          RR - Remote Refresh Reduction
      LD - Local Reliable Delivery          RM - Remote Node supports Message ID
=====
Neighbor      Interface      Hello  Last Oper  Flags
                               Change
=====
No Matching Entries
=====
```

Table 24: Show Router RSVP-TE Neighbor Output Fields

Label	Description
Neighbor	The IP address of the RSVP-TE neighbor
Interface	The interface ID of the RSVP-TE neighbor
Hello	The status of the Hello message
Last Oper Change	The time of the last operational change to the connection
Flags	Any flags associated with the connection to the neighbor

session

Syntax	session [<i>session-type</i>] [from <i>ip-address</i> to <i>ip-address</i> lsp-name <i>name</i>] [status { up down }] [detail]
Context	show>router>rsvp
Description	This command shows RSVP-TE session information.
Parameters	<p><i>session-type</i> — specifies the session type</p> <p>Values originate, transit, terminate, detour, detour-transit, detour-terminate, bypass-tunnel, manual-bypass</p> <p>from <i>ip-address</i> — specifies the IP address of the originating router</p> <p>to <i>ip-address</i> — specifies the IP address of the egress router</p> <p><i>name</i> — specifies the name of the LSP used in the path</p> <p>status up — specifies to display a session that is operationally up</p> <p>status down — specifies to display a session that is operationally down</p> <p>detail — displays detailed information</p>
Output	The following output is an example of RSVP-TE session information, and Table 25 describes the fields.

Sample Output

```

A:ALU-12# show router rsvp session
=====
RSVP Sessions
=====
From          To          Tunnel LSP   Name                               State
            ID      ID
-----
10.20.1.3     10.20.1.1    1      37   C_A_1::C_A_1                      Up
10.20.1.3     10.20.1.1    2      38   C_A_2::C_A_2                      Up
10.20.1.3     10.20.1.1    3      39   C_A_3::C_A_3                      Up
10.20.1.3     10.20.1.1    4      40   C_A_4::C_A_4                      Up
10.20.1.1     10.20.1.3    2      40   A_C_2::A_C_2                      Up
10.20.1.1     10.20.1.3    3      41   A_C_3::A_C_3                      Up
10.20.1.1     10.20.1.3    4      42   A_C_4::A_C_4                      Up
10.20.1.1     10.20.1.3    5      43   A_C_5::A_C_5                      Up
10.20.1.1     10.20.1.3    6      44   A_C_6::A_C_6                      Up
10.20.1.1     10.20.1.3    7      45   A_C_7::A_C_7                      Up
10.20.1.1     10.20.1.3    8      46   A_C_8::A_C_8                      Up
10.20.1.3     10.20.1.1    5      41   C_A_5::C_A_5                      Up
10.20.1.3     10.20.1.1    6      42   C_A_6::C_A_6                      Up
10.20.1.3     10.20.1.1    7      43   C_A_7::C_A_7                      Up
10.20.1.3     10.20.1.1    8      44   C_A_8::C_A_8                      Up
...
-----
Sessions : 65
=====
A:ALU-12#

```

```

A:ALU-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:ALU-12#

```

Table 25: Show Router RSVP-TE Session Output Fields

Label	Description
From	The IP address of the originating router
To	The IP address of the egress router
Tunnel ID	The ID of the ingress node of the tunnel supporting this RSVP-TE session
LSP ID	The ID assigned by the agent to this RSVP-TE session
Name	The administrative name assigned to the RSVP-TE session by the agent
State	Down — the operational state of this RSVP-TE session is down
	Up — the operational state of this RSVP-TE session is up

statistics

Syntax	statistics
Context	show>router>rsvp
Description	This command displays global statistics in the RSVP-TE instance.
Output	The following output is an example of RSVP-TE statistics information, and Table 26 describes the fields.

Sample Output

```
A:ALU-12# show router rsvp statistics
=====
RSVP Global Statistics
=====
PATH Timeouts : 0                      RESV Timeouts : 0
=====
```

Table 26: Show Router RSVP-TE Statistics Output Fields

Label	Description
PATH Timeouts	The total number of PATH timeouts
RESV Timeouts	The total number of RESV timeouts

status

Syntax	status
Context	show>router>rsvp
Description	This command displays RSVP-TE operational status.
Output	The following output is an example of RSVP-TE status information, and Table 27 describes the fields.

Sample Output

```
A:ALU-12# 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   : Enabled
Rapid Retransmit  : 5 hmsec            Rapid Retry Limit : 3
=====
```

Table 27: Show Router RSVP-TE Status Output Fields

Label	Description
Admin Status	Down — RSVP-TE is administratively disabled
	Up — RSVP-TE is administratively enabled
Oper Status	Down — RSVP-TE is operationally down
	Up — RSVP-TE is operationally up
Keep Multiplier	The keep-multiplier <i>number</i> used by RSVP-TE to declare that a reservation is down or the neighbor is down
Refresh Time	The refresh-time <i>interval</i> , in seconds, between the successive PATH and RESV refresh messages
Message Pacing	Enabled — RSVP-TE messages, specified in the max-burst command, are sent in a configured interval, specified in the period command
	Disabled — message pacing is disabled. RSVP-TE message transmission is not regulated.
Pacing Period	The time interval, in milliseconds, during which the router can send the number of RSVP-TE messages specified in the max-burst command
Max Packet Burst	The maximum number of RSVP-TE messages that are sent under normal operating conditions in the period specified
Refresh Bypass	Enabled — the refresh-reduction-over-bypass command is enabled
	Disabled — the refresh-reduction-over-bypass command is disabled
Rapid Retransmit	The time interval for the rapid retransmission time, which is used in the retransmission mechanism that handles unacknowledged message_id objects (the units “hmsec” represent hundreds of msec; for example, 5 hmsec represents 500 msec)
Rapid Retry Limit	The value of the rapid retry limit, which is used in the retransmission mechanism that handles unacknowledged message_id objects

Clear Commands

interface

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

lsp

Syntax	lsp [<i>lsp-name</i>]
Context	clear>router>mpls
Description	This command resets and restarts an LSP.
Parameters	<i>lsp-name</i> — specifies the name of the LSP to clear

interface

Syntax	interface [<i>ip-int-name</i>] [statistics]
Context	clear>router>rsvp
Description	This command resets or clears statistics for an RSVP-TE interface.
Parameters	<p><i>ip-int-name</i> — identifies the IP interface to clear. The interface name cannot be in the form of an IP address. 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>statistics — clears only statistics</p>

statistics

Syntax	statistics
Context	clear>router>rsvp
Description	This command clears global statistics for the RSVP-TE instance; for example, clears path and resv timeout counters.

Debug Commands

mpls

Syntax	[no] mpls [lsp <i>lsp-name</i>] [sender <i>source-address</i>] [endpoint <i>endpoint-address</i>] [tunnel-id <i>tunnel-id</i>] [lsp-id <i>lsp-id</i>] [interface <i>ip-int-name</i>]
Context	debug>router
Description	This command enables and configures debugging for MPLS.
Parameters	<p><i>lsp-name</i> — the name that identifies the LSP. The LSP name can be up to 32 characters long and must be unique.</p> <p><i>source-address</i> — specifies the system IP address of the sender</p> <p><i>endpoint-address</i> — specifies the far-end system IP address</p> <p><i>tunnel-id</i> — specifies the MPLS SDP ID</p> <p>Values 0 to 4294967295</p> <p><i>lsp-id</i> — specifies the LSP ID</p> <p>Values 1 to 65535</p> <p><i>ip-int-name</i> — identifies the interface. The interface name cannot be in the form of an IP address. 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>

event

Syntax	[no] event
Context	debug>router>mpls debug>router>rsvp
Description	<p>This command enables debugging for specific events.</p> <p>The no form of the command disables the debugging.</p>

all

Syntax	all [detail] no all
Context	debug>router>mpls>event debug>router>rsvp>event
Description	This command debugs all events. The no form of the command disables the debugging.
Parameters	detail — displays detailed information about all events

frr

Syntax	frr [detail] no frr
Context	debug>router>mpls>event
Description	This command debugs fast reroute events. The no form of the command disables the debugging.
Parameters	detail — displays detailed information about reroute 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

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	[no] rsvp [lsp <i>lsp-name</i>] [sender <i>sender-address</i>] [endpoint <i>endpoint-address</i>] [tunnel-id <i>tunnel-id</i>] [lsp-id <i>lsp-id</i>] [interface <i>ip-int-name</i>] no rsvp
Context	debug>router
Description	This command enables and configures debugging for RSVP.
Parameters	<p><i>lsp-name</i> — name that identifies the LSP. The LSP name can be up to 80 characters long and must be unique.</p> <p><i>sender-address</i> — specifies the system IP address of the sender (a.b.c.d)</p> <p><i>endpoint-address</i> — specifies the far-end system IP address (a.b.c.d)</p> <p><i>tunnel-id</i> — specifies the RSVP-TE tunnel ID</p> <p>Values 0 to 4294967295</p> <p><i>lsp-id</i> — specifies the LSP ID</p> <p>Values 1 to 65535</p> <p><i>ip-int-name</i> — identifies the interface. The interface name cannot be in the form of an IP address. 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>

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

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-TE reservation events. The no form of the command disables the debugging.
Parameters	detail — displays detailed information about RSVP-TE 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 context to debug packets.

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

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-TE 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

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-TE 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-TE RESV packets. The no form of the command disables the debugging.
Parameters	detail — displays detailed information about RSVP-TE RESV packets

resvterr

Syntax	resvterr [detail] no resvterr
Context	debug>router>rsvp>packet
Description	This command debugs ResvErr packets. The no form of the command disables the debugging.
Parameters	detail — displays detailed information about ResvErr packets

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 the Label Distribution Protocol (LDP).

Topics in this chapter include:

- [Label Distribution Protocol on page 196](#)
 - [LDP and MPLS on page 196](#)
 - [LDP Architecture on page 197](#)
 - [LDP Subsystem Interrelationships on page 198](#)
 - [Execution Flow on page 200](#)
 - [Label Exchange on page 202](#)
 - [LDP Filters on page 202](#)
 - [ECMP Support for LDP on page 203](#)
 - [Graceful Restart Helper on page 205](#)
- [LDP Process Overview on page 206](#)
- [Configuration Notes on page 207](#)
 - [Reference Sources on page 207](#)
- [Configuring LDP with CLI on page 209](#)
- [LDP Command Reference on page 221](#)

Label Distribution Protocol

Label Distribution Protocol (LDP) is used to distribute labels in non-traffic-engineered applications. LDP allows routers to establish LSPs through a network by mapping network-layer routing information directly to data link LSPs.

An LSP is defined by the set of labels from the ingress LER to the egress LER. LDP associates a Forwarding Equivalence Class (FEC) with each LSP it creates. An FEC is a collection of common actions associated with a class of packets. When an ingress LER assigns a label to an 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 by each LSR, where each LSR splices incoming labels for the 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 an FEC label binding in response to an explicit request from another LSR. This is known as Downstream On Demand (DOD) label distribution. LSRs can also distribute label bindings to LSRs that have not explicitly requested them. This is called Downstream Unsolicited (DU). For LDP on the 7705 SAR, Downstream Unsolicited (DU) mode is implemented.

LDP and MPLS

LDP performs dynamic label distribution in MPLS environments. The LDP operation begins with a Hello discovery process network to form an adjacency with an LDP peer in the network. LDP peers are two MPLS routers 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 and to distribute its own label information (LDP is bidirectional), and 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 pseudowires (PWs) 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 is 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 that enable fast access and packet identification are maintained in the forwarding plane.

When an unlabeled packet ingresses the 7705 SAR, classification policies associate it with an FEC, the appropriate label is imposed on the packet, and then the packet is forwarded. Other actions can also take place on a packet before it is forwarded, including imposing additional labels, other encapsulations, or learning actions. Once 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 performed on the packet and then the packet is forwarded.

The LDP implementation provides support for DU, ordered control, and liberal label retention mode.

For LDP label advertisement, DU mode is supported. To prevent filling the uplink bandwidth with unassigned label information, Ordered Label Distribution Control mode is supported.

A PW/VLL label can be dynamically assigned by targeted LDP operations. Targeted LDP allows the inner labels (that is, the VLL labels) in the MPLS headers to be managed automatically. This makes it easier for operators to manage the VLL connections. There is, however, additional signaling and processing overhead associated with this targeted LDP dynamic label assignment.

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 11](#). 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 interprocess communication through the event queue, as well as through updates to the various data structures. The following list describes the primary data structures that LDP maintains:

- FEC/label database — this database contains all the FEC-to-label mappings, including both sent and received. It also contains both address FECs (prefixes and host addresses) as well as service FECs (L2 VLLs).
- 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

LDP Subsystem Interrelationships

[Figure 11](#) shows the relationships between LDP subsystems and other 7705 SAR OS subsystems. The following sections describe how the subsystems work to provide services.

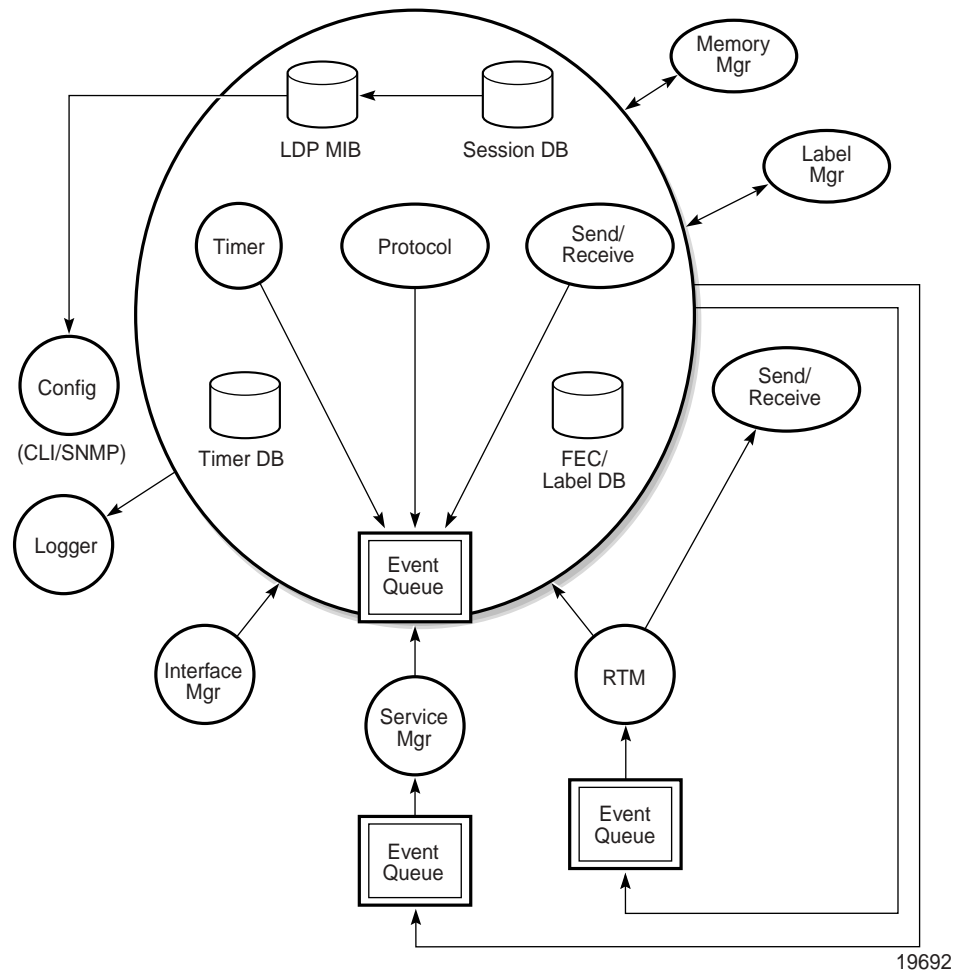
Memory Manager and LDP

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 large 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 bootup; hence, anything that causes it to fail will likely indicate that the system is not functional. The 7705 SAR uses a label range from 28 672 (28K) to 131 071 (128K-1) to allocate all dynamic labels, including VC labels.

Figure 11: LDP Subsystem Interrelationships

LDP Configuration

The 7705 SAR 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, targeted LDP 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. Refer to the 7705 SAR OS System Management Guide for logger configuration information.

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 the service manager of relevant LDP events, such as connection setups and failures, timeouts, and labels signaled or withdrawn.

Execution Flow

LDP activity in the 7705 SAR 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 are met, such as ensuring that the system IP interface and the label manager are operational, and ensuring that there is memory available. It then allocates a pool of memory to itself and initializes its databases.

Session Lifetime

In order for a targeted LDP session to be established, an adjacency has to be created. The LDP extended discovery mechanism requires hello messages to be exchanged between two peers for session establishment. Once the adjacency is established, session setup is attempted.

Adjacency Establishment

In the 7705 SAR, adjacency management is done through the establishment of a Service Destination Point (SDP) object, which is a service entity in the Alcatel-Lucent service model.

The Alcatel-Lucent service model uses logical entities that interact to provide a service. The service model requires the service provider to create and configure four main entities:

- customers
- services
- Service Access Points (SAPs) on local 7705 SAR routers
- SDPs that connect to one or more remote 7705 SAR routers or 77x0 SR routers

An SDP is the network-side termination point for a tunnel to a remote 7705 SAR or 77x0 SR router. An SDP defines a local entity that includes the system IP address of the remote 7705 SAR routers and 77x0 SR routers, and a path type.

Each SDP comprises:

- the SDP ID
- the transport encapsulation type, MPLS
- the far-end system IP address

If the SDP is identified as using LDP signaling, then an LDP extended hello adjacency is attempted.

If another SDP is created to the same remote destination and if LDP signaling is enabled, no further action is taken, since only one adjacency and one LDP session exists between the pair of nodes.

An SDP is a unidirectional object, so a pair of SDPs pointing at each other must be configured in order for an LDP adjacency to be established. Once an adjacency is established, it is maintained through periodic hello messages.

Session Establishment

When the LDP adjacency is established, the session setup follows as per the LDP specification. Initialization and keepalive messages complete the session setup, followed by address messages to exchange all interface IP addresses. Periodic keepalives or other session messages maintain the session liveliness.

Since TCP is back-pressured by the receiver, it is necessary to be able to push that back-pressure all the way into the protocol. Packets that cannot be sent are buffered on the session object and reattempted as the back-pressure eases.

Label Exchange

Label exchange is initiated by the service manager. When an SDP is attached to a service (that is, once 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

Label actions can also occur for the following reasons:

- MTU changes — LDP withdraws the previously assigned label and resignals the FEC with the new Maximum Transmission Unit (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, the received label 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 that it uses no memory (0 bytes) when it is not running.

LDP Filters

Only inbound LDP label binding filtering are supported.

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

- neighbor — match on bindings received from the specified peer
- prefix-list — match on bindings with the specified prefix/prefixes

The default import behavior is to accept all FECs received from peers.

The default export behavior is to originate label bindings for the system address and propagate all FECs received.



Note: In order for the 7705 SAR to consider a received label to be active, there must be an exact match to the FEC advertised together with the label found in the routing table. This can be achieved by configuring a static route pointing to the prefix encoded in the FEC.

ECMP Support for LDP

Equal Cost Multipath Protocol (ECMP) support for LDP performs load balancing for VLL-type services that use LDP-based LSPs as transport tunnels, by having multiple equal-cost outgoing next hops for an IP prefix.

There is only one next-hop peer for a network link. To offer protection from a network link or next-hop peer failure, multiple network links can be configured to connect to different next-hop peers, or multiple links to the same peer. For example, an MLPPP link and an Ethernet link can be connected to two peers, or two Ethernet links can be connected to the same peer. ECMP occurs when the cost of each link reaching a target IP prefix is equal.

The 7705 SAR uses a liberal label retention mode, which retains all labels for an IP prefix from all next-hop peers. If ECMP is not enabled, the label from only one of the next-hop peers is selected and installed in the forwarding plane. In this case, the algorithm used to distribute the traffic flow looks up the route information, and selects the network link with the lowest IP address. If the selected network link or next-hop peer fails, another next-hop peer is selected, and LDP reprograms the forwarding plane to use the label sent by the newly selected peer.



Note: The 7705 SAR does not support load balancing of pure IP traffic or pure MPLS traffic over ECMP routes; that is, it does not support FIB ECMP or LFIB ECMP. The 7705 SAR does, however, support VLL ECMP and VPRN transport tunnel ECMP on the LER node.



Notes:

- Because timeout is built into dynamic ARP, the MAC address of the remote peer needs to be renewed periodically. The flow of IP traffic resets the timers back to their maximum values. In the case of LDP ECMP, one link could be used for transporting user MPLS (pseudowire) traffic but the LDP session could possibly be using a different equal cost link. For LDPs using ECMP and for static LSPs, it is important to ensure that the remote MAC address is learned and does not expire. Configuring static ARP entries or running continuous IP traffic ensures that the remote MAC address is always known. Running BFD for fast detection of Layer 2 faults or running any OAM tools with SAA ensures that the learned MAC addresses do not expire.
- ARP entries are refreshed by static ARP and BFD, SAA, OSPF, IS-IS, or BGP.
- For information on configuring static ARP and running BFD, refer to the 7705 SAR OS Router Configuration Guide.

Label Operations

If an LSR is the ingress router for a given IP prefix, LDP programs a PUSH operation for the prefix in the IOM. This creates an LSP ID to the Next Hop Label Forwarding Entry (NHLFE) mapping (LTN mapping) and an LDP tunnel entry in the forwarding plane. LDP will also inform the Tunnel Table Manager (TTM) about this tunnel. Both the LSP ID to NHLFE (LTN) entry and the tunnel entry will have an NHLFE for the label mapping that the LSR received from each of its next-hop peers.

If the LSR is to behave as a transit router for a given IP prefix, LDP will program a SWAP operation for the prefix in the IOM. This involves creating an Incoming Label Map (ILM) entry in the forwarding plane. The ILM entry might need to map an incoming label to multiple NHLFEs.

If an LSR is an egress router for a given IP prefix, LDP will program a POP entry in the IOM. This too will result in an ILM entry being created in the forwarding plane, but with no NHLFEs.

When unlabeled packets arrive at the ingress LER, the forwarding plane consults the LTN entry and uses a hashing algorithm to map the packet to one of the NHLFEs (PUSH label) and forward the packet to the corresponding next-hop peer. For a labeled packet arriving at a transit or egress LSR, the forwarding plane consults the ILM entry and either uses a hashing algorithm to map it to one of the NHLFEs if they exist (SWAP label) or simply routes the packet if there are no NHLFEs (POP label).

Graceful Restart Helper

Graceful Restart (GR) is part of the LDP handshake process (that is, the LDP peering session initialization) and needs to be supported by both peers. GR provides a mechanism that allows the peers to cope with a service interruption due to a CSM switchover, which is a period of time when the standby CSM is not capable of synchronizing the states of the LDP sessions and labels being advertised and received.

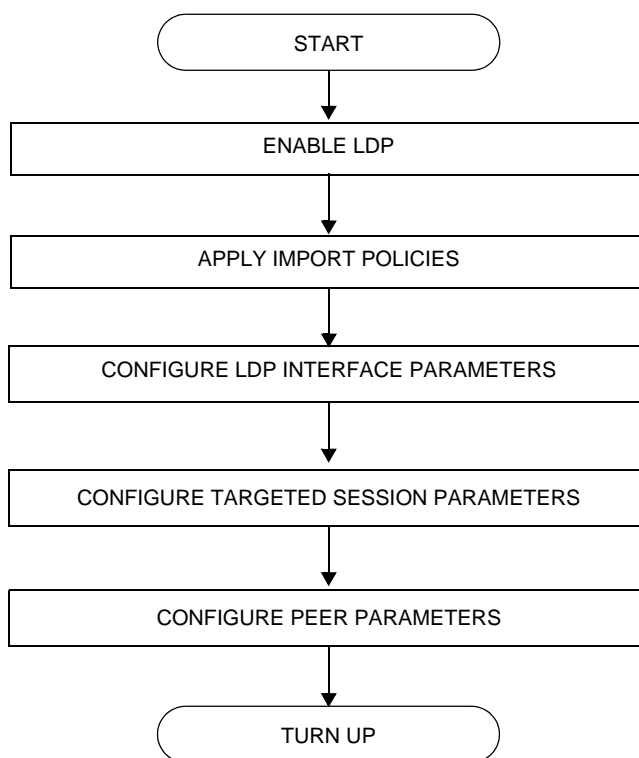
Graceful Restart Helper (GR-Helper) decouples the data plane from the control plane so that if the control plane is not responding (that is, there is no LDP message exchange between peers), then the data plane can still forward frames based on the last known (advertised) labels.

Because the 7705 SAR supports non-stop services / high-availability for LDP (and MPLS), the full implementation of GR is not needed. However, GR-Helper is implemented on the 7705 SAR to support non-high-availability devices. With GR-Helper, if an LDP peer of the 7705 SAR requests GR during the LDP handshake, the 7705 SAR agrees to it but does not request GR. For the duration of the LDP session, if the 7705 SAR LDP peer fails, the 7705 SAR continues to forward MPLS packets based on the last advertised labels and will not declare the peer dead until the GR timer expires.

LDP Process Overview

Figure 12 displays the process to provision basic LDP parameters.

Figure 12: LDP Configuration and Implementation



Configuration Notes

Refer to the 7705 SAR OS Services Guide for information about signaling.

Reference Sources

For information on supported IETF drafts and standards, as well as standard and proprietary MIBs, refer to [Standards and Protocol Support](#).

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 210](#)
- [Basic LDP Configuration on page 211](#)
- [Common Configuration Tasks on page 212](#)
- [LDP Configuration Management Tasks on page 219](#)

LDP Configuration Overview

When the 7705 SAR implementation of LDP is instantiated, the protocol is in the no shutdown state. In addition, targeted sessions are then enabled. The default parameters for LDP are set to the documented values for targeted sessions in *draft-ietf-mpls-ldp-mib-09.txt*.

LDP must be enabled in order for signaling to be used to obtain the ingress and egress labels in frames transmitted and received on the service destination point (SDP). When signaling is off, labels must be manually configured when the SDP is bound to a service.

Basic LDP Configuration

This section provides information to configure LDP and gives configuration examples of common configuration tasks.

The LDP protocol instance is created in the `no shutdown` (enabled) state.

The following example displays the default LDP configuration output.

```
ALU-1>config>router>ldp# info
-----
          interface-parameters
          exit
          targeted-session
          exit
-----
ALU-1>config>router>ldp#
```

Common Configuration Tasks

This section provides a brief overview of the following common configuration tasks to configure LDP:

- [Enabling LDP](#)
- [Configuring Graceful Restart Helper Parameters](#)
- [Applying Import Policies](#)
- [Configuring Interface Parameters](#)
- [Specifying Targeted Session Parameters](#)
- [Specifying Peer Parameters](#)
- [Enabling LDP Signaling and Services](#)

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 CLI syntax to enable LDP on a 7705 SAR router:

CLI Syntax: `ldp`

Example: `config>router# ldp`

The following example displays the enabled LDP configuration output.

```
ALU-1>config>router# info
-----
...
#-----
echo "LDP Configuration"
#-----
      ldp
        interface-parameters
        exit
        targeted-session
        exit
      exit
-----
...
ALU-1>config>router#
```

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 SAR recovery is self-contained and relies on information stored internally to self-heal.

Maximum recovery time is the time (in seconds) that 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) that the LSR is willing to retain its MPLS forwarding state. The time should be long enough to allow the neighboring LSRs to resynchronize 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
```

Example:

```
config>router>ldp
config>router>ldp# graceful-restart
config>router>ldp>graceful-restart# maximum-recovery-
time 120
config>router>ldp>graceful-restart# neighbor-liveness-
time 60
config>router>ldp# exit
config>router#
```

The following example displays the import policy configuration output.

```
ALU-1>config>router>ldp>graceful-restart# info
-----
maximum-recovery-time 120
neighbor-liveness-time 60
-----
ALU-1>config>router>ldp>graceful-restart#
```

Applying Import Policies

Inbound label binding filtering is supported. Inbound filtering allows a route policy to control the label bindings that 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 the following:

- neighbor — match on bindings received from the specified peer
- prefix-list — match on bindings with the specified prefix or prefixes

The import policy must already exist. Policies are configured in the `config>router>policy-options` context. Refer to the 7705 SAR OS Router Configuration Guide for details.



Note: The 7705 SAR supports 1024 inbound labels (maximum number of LSPs, including protection and redundant LSPs). If the supported number of labels in the 7705 SAR is exceeded, all LDP sessions will be shut down and all labels will be removed. To recover the LDP sessions, perform a `shutdown/no shutdown` combination of commands in the `configure>router>ldp` context.

Use the following CLI syntax to apply import policies:

CLI Syntax: `config>router>ldp`
`import policy-name [policy-name...(up to 5 max)]`

Example: `config>router>ldp`
`config>router>ldp# import LDP-import`
`config>router>ldp# exit`
`config>router#`

The following example displays the import policy configuration output.

```
ALU-1>config>router>ldp# info
-----
import "LDP-import"
interface-parameters
exit
targeted-session
exit
-----
ALU-1>config>router>ldp#
```

Configuring Interface Parameters

Use the following CLI syntax to configure LDP interface parameters:

CLI Syntax:

```
config>router# ldp
      interface-parameters
      hello timeout factor
      interface ip-int-name
      hello timeout factor
      keepalive timeout factor
      local-lsr-id {system|interface}
      transport-address {system|interface}
no shutdown
keepalive timeout factor
transport-address {system|interface}
```

Example:

```
config>router# ldp
config>router>ldp# interface-parameters
config>router>ldp>if-params# interface to-104
config>router>ldp>if-params>if# hello 15 3
config>router>ldp>if-params>if# local-lsr-id system
config>router>ldp>if-params>if# no shutdown
config>router>ldp>if-params>if# exit
config>router>ldp>if-params# exit
config>router>ldp#
```

The following example displays the LDP interface parameter configuration output.

```
ALU-1>config>router>ldp# info
-----
import "LDP-import"
  interface-parameters
    hello 15 3
    keepalive 30 3
    interface "to-104"
      hello 15 3
      keepalive 30 3
      local-lsr-id system
      no shutdown
    exit
  exit
targeted-session
exit
no shutdown
-----
ALU-1>config>router>ldp#
```

Specifying Targeted Session Parameters

Use the following CLI 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
```

Example:

```
config>router# ldp
config>router>ldp# targeted-session
config>router>ldp>targ-session# hello 5000 255
config>router>ldp>targ-session# keepalive 5000 255
config>router>ldp>targ-session# peer 10.10.10.104
config>router>ldp>targ-session>peer# hello 2500 100
config>router>ldp>targ-session>peer# keepalive 15 3
config>router>ldp>targ-session>peer# no shutdown
config>router>ldp>targ-session>peer# exit
config>router>ldp>targ-session# exit
config>router>ldp#
```

The following example displays the LDP targeted session configuration output.

```
ALU-1>config>router>ldp# info
-----

import "LDP-import"
interface-parameters
    hello 15 3
    keepalive 30 3
    interface "to-104"
        hello 15 3
        keepalive 30 3
        no shutdown
    exit
exit
    targeted-session
        hello 5000 255
        keepalive 5000 255
        peer 10.10.10.104
            hello 2500 100
            keepalive 15 3
        exit
    exit
-----

ALU-1>config>router>ldp#
```

Specifying Peer Parameters

Use the following CLI syntax to specify LDP peer parameters:

CLI Syntax:

```
config>router# ldp
      peer-parameters
      peer ip-address
          authentication-key {authentication-key|hash-key} [hash|hash2]
```

Example:

```
config>router# ldp
config>router>ldp# peer-parameters
config>router>ldp>peer-params# peer 10.10.10.104
config>router>ldp>peer-params>peer$ authentication-key
testuser
config>router>ldp>peer-params>peer$ exit
```

The following example displays the LDP peer parameters configuration output.

```
ALU-1>config>router>ldp# info
-----
import "LDP-import"
graceful-restart
exit
import "LDP-import"
peer-parameters
  peer 10.10.10.104
    authentication-key "nGjXyHQtcGhxbBm.kDeYdzSmPZy9KK03" hash2
  exit
exit
interface-parameters
  interface "test"
  exit
  interface "to-104"
    hello 15 3
  exit
exit
targeted-session
  hello 5000 255
  keepalive 5000 255
  peer 10.10.10.104
    hello 2500 100
    keepalive 15 3
  exit
exit
-----
ALU-1>config>router>ldp#
```

Enabling 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 destination points (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 triggers session establishment. The SDP's signaling default enables `tldp`. The SDP uses the targeted-session parameters configured in the `config>router>ldp>targeted-session` context.

The `service>sdp>ldp` and `router>lsp` commands are mutually exclusive; you can either specify an LSP or enable an LDP. There cannot be two methods of transport in a single SDP.

To enable LDP on the SDP when an LSP is already specified, the LSP must be removed from the configuration using the `no lsp lsp-name` command. For further information about configuring SDPs, refer to the 7705 SAR OS Services Guide.

Use the following CLI syntax to enable LDP on an MPLS SDP:

CLI Syntax:

```
config>service>sdp#
    ldp
    signaling {off|tldp}
```

The following example displays an SDP configuration output with the signaling default `tldp` enabled.

```
ALU-1>config>service>sdp# info detail
-----
description "MPLS: to-99"
far-end 10.10.10.99
ldp
signaling tldp
path-mtu 4462
keep-alive
    hello-time 10
    hold-down-time 10
    max-drop-count 3
    timeout 5
    no message-length
    no shutdown
exit
no shutdown
-----
ALU-1>config>service>sdp#
```

LDP Configuration Management Tasks

This section discusses the following LDP configuration management tasks:

- [Disabling LDP](#)
- [Modifying Targeted Session Parameters](#)
- [Modifying Interface Parameters](#)

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 CLI syntax to disable LDP:

CLI Syntax: `no ldp
shutdown`

Example: `config>router# ldp
config>router>ldp# shutdown
config>router>ldp# exit
config>router# no ldp`

Modifying Targeted Session Parameters

You can modify LDP parameters without shutting down entities. The changes take effect immediately. 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 CLI syntax to revert targeted session parameters back to the default values.

Example: `config>router# ldp
config>router>ldp# targeted-session
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 example displays the default value output.

```
ALU-1>config>router>ldp>targeted# info detail
-----
no disable-targeted-session
hello 45 3
keepalive 40 4
-----
ALU-1>config>router>ldp>targeted#
```

Modifying Interface Parameters

You can modify LDP interface parameters without shutting down entities. The changes take effect immediately. Individual parameters cannot be deleted. The `no` form of an `interface-parameter` command reverts modified values back to the defaults.

The following example displays the CLI syntax to revert targeted session parameters back to the default values.

Example:

```
config>router# ldp
config>router>ldp>interface-parameters
config>router>ldp>if-params# no hello
config>router>ldp>if-params# no interface to-104
config>router>ldp>if-params# no keepalive
config>router>ldp>if-params# no transport-address
config>router>ldp>if-params# info
```

The following example displays the default value output.

```
ALU-1>config>router>ldp>targeted# info detail
-----
hello 15 3
keepalive 30 3
no transport-address
-----
ALU-1>config>router>ldp>targeted#
```

LDP Command Reference

Command Hierarchies

- [LDP Commands](#)
- [Show Commands](#)
- [Clear Commands](#)
- [Debug Commands](#)

LDP Commands

```

config
  — router [router-name]
    — [no] ldp
      — [no] graceful-restart
        — maximum-recovery-time interval
        — no maximum-recovery-time
        — neighbor-liveness-time interval
        — no neighbor-liveness-time
      — import policy-name [policy-name...(up to 5 max)]
      — no import
      — interface-parameters
        — hello timeout factor
        — no hello
        — [no] interface ip-int-name
          — hello timeout factor
          — no hello
          — keepalive timeout factor
          — no keepalive
          — local-lsr-id {system | interface}
          — no local-lsr-id
          — [no] shutdown
          — transport-address {system | interface}
          — no transport-address
        — keepalive timeout factor
        — no keepalive
        — transport-address {system | interface}
        — no transport-address
      — peer-parameters
        — [no] peer ip-address
          — authentication-key {authentication-key | hash-key} [hash | hash2]
          — no authentication-key
      — [no] shutdown
      — targeted-session
        — [no] disable-targeted-session
        — hello timeout factor
        — no hello
        — keepalive timeout factor
        — no keepalive
        — [no] peer ip-address
          — hello timeout factor
          — no hello
          — keepalive timeout factor
          — no keepalive
          — [no] shutdown

```

Show Commands

```

show
  — router [router-instance]
    — 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 vc-type vc-type [vc-id vc-id [session ip-addr[:label-space]]]
      — bindings service-id service-id [detail]
      — 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 [router-instance]
    — ldp
      — instance
      — interface ip-int-name [statistics]
      — peer ip-address [statistics]
      — session ip-addr[:label-space] [statistics]
      — statistics

```

Debug Commands

```
[no] debug
  — router [router-instance]
    — [no] ldp
      — [no] interface interface-name
        — [no] event
          — [no] messages
        — [no] packet
          — hello [detail]
          — no hello
      — [no] 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
```

Command Descriptions

- [Configuration Commands on page 226](#)
- [Show Commands on page 239](#)
- [Clear Commands on page 262](#)
- [Debug Commands on page 264](#)

Configuration Commands

- [Generic Commands on page 227](#)
- [LDP Global Commands on page 228](#)
- [Interface Parameters Commands on page 234](#)
- [Targeted Session Commands on page 236](#)
- [Peer Parameters Commands on page 237](#)

Generic Commands

shutdown

Syntax	[no] shutdown
Context	config>router>ldp config>router>ldp>if-params>if config>router>ldp>targ-session>peer
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>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 can be deleted.</p> <p>The no form of this command administratively enables an entity.</p> <p>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.</p>
Default	no shutdown

LDP Global Commands

ldp

Syntax	[no] ldp
Context	config>router
Description	<p>This command creates the context to configure an LDP protocol instance.</p> <p>When an LDP instance is created, the protocol is enabled (in the no shutdown state). To suspend the LDP protocol, use the shutdown command. Configuration parameters are not affected.</p> <p>The no form of the command deletes the LDP protocol instance, removing all associated configuration parameters. The LDP instance must first be disabled with the shutdown command before being deleted.</p>
Default	none — LDP must be explicitly enabled

graceful-restart

Syntax	[no] graceful-restart
Context	config>router>ldp
Description	<p>This command enables graceful restart helper.</p> <p>The no form of the command disables graceful restart.</p>
Default	graceful-restart

maximum-recovery-time

Syntax	maximum-recovery-time <i>interval</i> no maximum-recovery-time
Context	config>router>ldp>graceful-restart
Description	<p>This command configures the local maximum recovery time, which is the time (in seconds) that the sender of the TLV would like the receiver to wait, after detecting the failure of LDP communication with the sender.</p> <p>The no form of the command returns the default value.</p>
Default	120

Parameters *interval* — specifies the maximum length of recovery time, in seconds
Values 15 to 1800

neighbor-liveness-time

Syntax **neighbor-liveness-time** *interval*
no neighbor-liveness-time

Context config>router>ldp>graceful-restart

Description This command configures the neighbor liveness time, which is the time (in seconds) that the LSR is willing to retain its MPLS forwarding state. The time should be long enough to allow the neighboring LSRs to resynchronize all the LSPs in a graceful manner, without creating congestion in the LDP control plane.

The **no** form of the command returns the default value.

Default 120

Parameters *interval* — specifies the length of time, in seconds
Values 5 to 300

import

Syntax **import** *policy-name* [*policy-name* ...(up to 5 max)]
no import

Context config>router>ldp

Description This command configures import route policies to determine which routes are accepted from LDP neighbors. Policies are configured in the **config>router>policy-options** context. Refer to the 7705 SAR OS Router Configuration Guide.

If no import policy is specified, LDP accepts all routes 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 specified name(s) must already be defined.

The **no** form of the command removes all policies from the configuration.

Default **no import**

Parameters *policy-name* — specifies 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.

hello

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

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

Description This command configures the hold time. This is 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 is sent in the hello messages to the neighbor. Hold time cannot be less than three times the hello interval. The hold time can be configured globally (applies to all LDP interfaces) or per interface. The most specific value is used.

When an LDP session is being set up, the hold time is negotiated to the lower of the two peers. Once an operational value is agreed upon, the **hello factor** is used to derive the value of the hello interval.

The **no** form of the command:

- at the interface-parameters and targeted-session levels, sets the **hello timeout** and the **hello factor** to the default values
- at the interface level, sets the **hello timeout** and the **hello factor** to the value defined under the interface-parameters level
- at the peer level, sets the **hello timeout** and the **hello factor** to the value defined under the targeted-session level

Default The default value is dependent upon the CLI context. [Table 28](#) lists the **hello timeout factor** default values.

Table 28: Hello Timeout Factor Default Values

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	<i>timeout</i> — configures the time interval, in seconds, that LDP waits before declaring a neighbor down
Values	1 to 65535
	<i>factor</i> — specifies the number of keepalive messages that should be sent on an idle LDP session in the hello timeout interval
Values	1 to 255

keepalive

Syntax	keepalive <i>timeout factor</i> no keepalive
Context	config>router>ldp>if-params config>router>ldp>if-params>if config>router>ldp>targ-session config>router>ldp>targ-session>peer
Description	<p>This command configures the time interval, in seconds, that LDP waits before tearing down the session. The <i>factor</i> parameter derives the keepalive interval.</p> <p>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.</p> <p>When an 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 <i>factor</i> is used to derive the value of the keepalive interval.</p> <p>The no form of the command:</p> <ul style="list-style-type: none"> • at the interface-parameters and targeted-session levels, sets the keepalive <i>timeout</i> and the keepalive <i>factor</i> to the default value • at the interface level, sets the keepalive <i>timeout</i> and the keepalive <i>factor</i> to the value defined under the interface-parameters level • at the peer level, sets the keepalive <i>timeout</i> and the keepalive <i>factor</i> to the value defined under the targeted-session level

Default The default value is dependent upon the CLI context. [Table 29](#) lists the **keepalive timeout factor** default values.

Table 29: Keepalive Timeout Factor Default Values

Context	Timeout	Factor
config>router>ldp>if-params	30	3
config>router>ldp>targ-session	40	4
config>router>ldp>if-params>if	Inherits values from interface-parameters context	
config>router>ldp>targ-session>peer	Inherits values from targeted-session context	

Parameters *timeout* — configures the time interval, expressed in seconds, that LDP waits before tearing down the session

Values 1 to 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 to 255

local-lsr-id

Syntax **local-lsr-id {system | interface}**
no local-lsr-id

Context config>router>ldp>if-params>if

Description This command enables the use of the address of the link LDP interface as the LSR-ID in order to establish an LDP adjacency and session with a directly connected LDP peer.

By default, the LDP session uses the system interface address as the LSR-ID. This means that targeted LDP (T-LDP) and interface LDP share a common LDP TCP session and therefore a common LDP label space. The system interface must be configured on the router or the LDP protocol will not come up on the node. At initial configuration, the LDP session to the peer remains down while the interface is down. If the user changes the LSR-ID while the LDP session is up, LDP immediately tears down the session and attempts to re-establish it using the new LSR-ID. If the interface used for the local LSR-ID goes down, then the LDP session will also go down.

The **interface** option is the recommended setting when static route-LDP synchronization is enabled.

When the **interface** option is selected, the transport connection (TCP) for the link LDP session configured by the [transport-address](#) command is automatically set to interface. Having both the **local-lsr-id** and transport address set to the local interface creates two TCP sessions to the peer and therefore two different LDP label spaces: one to the interface IP address for link LDP (L-LDP) and one to the system IP address for T-LDP.

The **no** form of the command resets the **local-lsr-id** to the default value.

Default **system**

Parameters **system** — specifies that the system IP address is used to set up the LDP session between neighbors
 interface — specifies that the IP interface address is used to set up the LDP session between neighbors

Interface Parameters Commands

interface-parameters

Syntax	interface-parameters
Context	config>router>ldp
Description	This command enables the context to configure LDP interfaces and parameters applied to LDP interfaces.

interface

Syntax	[no] interface <i>ip-int-name</i>
Context	config>router>ldp>if-params
Description	<p>This command enables LDP on the specified IP interface.</p> <p>The no form of the command deletes the LDP interface and all configuration information associated with the LDP interface.</p> <p>The LDP interface must be disabled using the shutdown command before it can be deleted.</p>
Parameters	<i>ip-int-name</i> — specifies 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 {system interface} 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 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 that 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 address selection situation can also occur when there is a link and a targeted adjacency, since targeted adjacencies request the session to be set up only to the system IP address.</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 as **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 the adjacency can be matched to the session.

The **no** form of the command:

- at the global level, sets the transport address to the default value
- at the interface level, sets the transport address to the value defined under the global level

Default **system**

Parameters **interface** — specifies that the IP interface address is used to set up the LDP session between neighbors. The transport address interface cannot be used if multiple interfaces exist between two neighbors, since only one LDP session is set up between two neighbors.

system — specifies that the system IP address is used to set up the LDP session between neighbors

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>targeted-session
Description	<p>This command disables support for 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 setup 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> — specifies the LDP peer in dotted-decimal notation

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-address</i> — specifies the LDP peer in dotted-decimal notation

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	<p>This command specifies the authentication key to be used between LDP peers before establishing sessions. Authentication uses the MD5 message-based digest.</p> <p>The no form of this command disables authentication.</p>
Default	none
Parameters	<p><i>authentication-key</i> — specifies the authentication key. Allowed values are any string up to 16 characters long (unencrypted) 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><i>hash-key</i> — specifies the hash key. Allowed values are any string up to 33 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>This is useful when a user must configure the parameter; however, for security purposes, the actual unencrypted key value is not provided.</p>

hash — specifies that the key is entered and stored on the node in encrypted form

hash2 — specifies that the key is entered and stored on the node in a more complex encrypted form



Note: If neither the **hash** or **hash2** keyword is specified, the key is entered in clear text. However, for security purposes, the key is stored on the node using hash encryption.

Show 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:ALU-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           Enabled      eta_keychain1
-----
No. of Peers: 1
=====
*A:ALU-48>config>router>ldp#
```

bindings

Syntax	bindings [fec-type <i>fec-type</i> [detail]] [session <i>ip-addr[:label-space]</i>] bindings <i>label-type start-label [end-label]</i> bindings { prefix <i>ip-prefix/mask</i> [detail]} [session <i>ip-addr[:label-space]</i>] bindings active [prefix <i>ip-prefix/mask</i>] bindings vc-type <i>vc-type</i> [vc-id <i>vc-id</i> [session <i>ip-addr[:label-space]</i>]] bindings service-id <i>service-id</i> [detail]
Context	show>router>ldp
Description	This command displays the contents of the label information base.
Parameters	<i>ip-addr</i> — specifies the IP address of the next hop <p style="margin-left: 40px;">Values a.b.c.d</p> <i>fec-type</i> — specifies the forwarding class type <p style="margin-left: 40px;">Values prefixes, services</p>

ip-prefix — specifies the IP prefix in dotted-decimal notation

Values a.b.c.d (host bits must be 0)

mask — specifies the 32-bit address mask used to indicate the bits of an IP address that are being used for the subnet address

Values 0 to 32

label-space — specifies the label space identifier that the router is advertising on the interface

Values 0 to 65535

label-type — specifies the label type to display

Values ingress-label, egress-label

start-label — specifies a label value to begin the display

Values 16 to 1048575

end-label — specifies a label value to end the display

Values 17 to 1048575

vc-type — specifies the VC type to display

Values atmvc, atmvp, cesopsn, cesopsn-cas, satop-e1, satop-t1, ethernet, ipipe

vc-id — specifies the VC ID to display

Values 1 to 4294967295

service-id — specifies the service ID number to display

Values 1 to 2147483647

Output The following output is an example of LDP bindings information, and [Table 30](#) describes the fields. Following the table are output examples for:

- LDP bindings detail
- LDP bindings session
- LDP bindings active

Sample Output - show router ldp bindings

```

A:cpm-a# show router ldp bindings
=====
LDP LSR ID: 1.1.1.30
=====
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, C - Cpipe Service
        TLV - (Type, Length: Value)
=====
LDP Prefix Bindings
=====
Prefix          Peer          IngLbl  EgrLbl  EgrIntf/LspId  EgrNextHop
-----
1.1.1.30/32     1.1.1.33       131071U  --      --              --
1.1.1.30/32     1.1.1.57       131071U  --      --              --
1.1.1.33/32     1.1.1.33       --      131071  1/2/3:1        10.4.1.33
1.1.1.33/32     1.1.1.57       131061U  131059  --              --
1.1.1.57/32     1.1.1.33       131060U  131067  --              --
1.1.1.57/32     1.1.1.57       --      131071  LspId 1         --
1.1.1.58/32     1.1.1.33       131059U  131066  --              --
1.1.1.58/32     1.1.1.57       131059N  131070  LspId 1         --
-----
No. of Prefix Bindings: 8

=====
LDP Service FEC 128 Bindings
=====
Type  VCId      SvcId      SDPIId  Peer          IngLbl  EgrLbl  LMTU  RMTU
-----
E-Eth 100      1          1        1.1.1.57      131069U  131068D  1500  1500
E-Eth 101      2          1        1.1.1.57      --      131067D  1500  1500
E-Eth 102      3          1        1.1.1.57      131067U  131066  1500  1500
E-Eth 103      4          1        1.1.1.57      131066W  131065  1500  1500
E-Eth 104      5          1        1.1.1.57      131065U  --      1500  0
E-Eth 105      5          1        1.1.1.57      131064U  --      1500  0
E-Eth 106      6          1        1.1.1.57      131063U  131064D  1500  1500
E-Eth 107      7          1        1.1.1.57      131062U  --      1500  0
-----
No. of VC Labels: 8

=====
LDP Service FEC 129 Bindings
=====
AGI          SvcId      SDPIId  SAPII  TAPII
Type          Peer          IngLbl  EgrLbl  LMTU  RMTU
-----
No Matching Entries Found
=====
A:cpm-a#

```

Table 30: LDP Bindings Output 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 A: Apipe service C: Cpipe service P: Ipipe service TLV: (Type, Length: Value)
Type	The service type exchanging labels in the SDP. The possible types displayed are Epipe, Spoke, and Unknown.	
VCId	The value used by each end of an SDP tunnel to identify the VC	
SvcID	Identifies the service in the service domain	
SDPId	Identifies the SDP in the service domain	
Peer	The IP address of the peer	
IngLbl	The ingress LDP label	
	U — indicates that the label is in use	
	R — indicates that the label has been released	
EgrLbl	The egress LDP label	
LMTU	The local MTU value	
RMTU	The remote MTU value	
No. of Prefix Bindings	The total number of LDP bindings on the router	
EgrIntf/LspId	The egress interface LSP ID	
EgrNextHop	The egress next-hop address	
No. of VC Labels	The total number of VC labels	
No. of Service Bindings	The total number of service bindings	
AGI Type	The address group identifier (AGI)	
SAII Peer	The source attachment individual identifier (SAII)	
TAII EgrLbl	The target attachment individual identifier (TAII)	
Vc-switching	Not applicable — always indicates no	
Egr. Flags	Specifies egress flag, if any	
Egr. Ctl Word	Indicates whether egress control words are used	

Table 30: LDP Bindings Output Fields (Continued)

Label	Description
Egr. Status Bits	Indicates whether egress status bits are supported
Igr. Flags	Specifies ingress flag, if any
Igr. Ctl Word	Indicates whether ingress control words are used
Igr. Status Bits	Indicates whether ingress status bits are supported
Op	The operation performed on the ingress or egress label in the LDP stack (push or pop)

Sample Output - show router ldp bindings detail

```

A:cpm-a# show router ldp bindings detail
=====
LDP LSR ID: 1.1.1.30
=====
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, C - Cpipe Service
        TLV - (Type, Length: Value)
=====
LDP Prefix Bindings
=====
-----
Prefix                : 1.1.1.30/32
-----
Ing Lbl                : 131071U                Peer                : 1.1.1.33
Egr. Flags             : None                    Ing. Flags             : None
-----
Prefix                : 1.1.1.30/32
-----
Ing Lbl                : 131071U                Peer                : 1.1.1.57
Egr. Flags             : None                    Ing. Flags             : None
-----
Prefix                : 1.1.1.33/32
-----
Ing Lbl                : --                      Peer                : 1.1.1.33
Egr Lbl                : 131071                  Egr Int/LspId         : 1/2/3:1
EgrNextHop             : 10.4.1.33
Egr. Flags             : None                    Ing. Flags             : None
-----
Prefix                : 1.1.1.33/32
-----
Ing Lbl                : 131061U                Peer                : 1.1.1.57
Egr Lbl                : 131059                  Egr Int/LspId         : --
EgrNextHop             : --
Egr. Flags             : None                    Ing. Flags             : None
-----
Prefix                : 1.1.1.57/32

```

LDP Command Reference

```

-----
Ing Lbl       : 131060U                               Peer       : 1.1.1.33
Egr Lbl       : 131067                               Egr Int/LspId : --
EgrNextHop    : --
Egr. Flags    : None                                Ing. Flags    : None
-----
Prefix        : 1.1.1.57/32
-----
Ing Lbl       : --                                    Peer       : 1.1.1.57
Egr Lbl       : 131071                               Egr Int/LspId : LspId 1
EgrNextHop    : --
Egr. Flags    : None                                Ing. Flags    : None
Lsp Name      : lsp_both2
-----
Prefix        : 1.1.1.58/32
-----
Ing Lbl       : 131059U                               Peer       : 1.1.1.33
Egr Lbl       : 131066                               Egr Int/LspId : --
EgrNextHop    : --
Egr. Flags    : None                                Ing. Flags    : None
-----
Prefix        : 1.1.1.58/32
-----
Ing Lbl       : 131059N                               Peer       : 1.1.1.57
Egr Lbl       : 131070                               Egr Int/LspId : LspId 1
EgrNextHop    : --
Egr. Flags    : None                                Ing. Flags    : None
Lsp Name      : lsp_both2
=====
No. of Prefix Bindings: 8

=====
LDP Service Bindings
=====
-----
Type          : E-Eth                                VcId        : 100
SvcId         : 1                                    SdpId        : 1
Peer Address   : 1.1.1.57                            Vc-switching : No
LMTU          : 1500                                RMTU         : 1500
Egr. Lbl      : 131068D                            Egr. Ctl Word : No
Egr. Flags    : None                                Egr. Status Bits : Supported (0x16)
Ing. Lbl      : 131069U                            Ing. Ctl Word  : No
Ing. Flags    : None                                Ing. Status Bits : Supported (0x0)
-----
Type          : E-Eth                                VcId        : 101
SvcId         : 2                                    SdpId        : 1
Peer Address   : 1.1.1.57                            Vc-switching : No
LMTU          : 1500                                RMTU         : 1500
Egr. Lbl      : 131067D                            Egr. Ctl Word : Yes
Egr. Flags    : None                                Egr. Status Bits : Supported (0x16)
Ing. Lbl      : --                                    Ing. Ctl Word  : No
Ing. Flags    : Released                            Ing. Status Bits : N/A
-----

```

Type	: E-Eth	VcId	: 102
SvcId	: 3	SdpId	: 1
Peer Address	: 1.1.1.57	Vc-switching	: No
LMTU	: 1500	RMTU	: 1500
Egr. Lbl	: 131066	Egr. Ctl Word	: No
Egr. Flags	: None	Egr. Status Bits	: N/A
Ing. Lbl	: 131067U	Ing. Ctl Word	: No
Ing. Flags	: None	Ing. Status Bits	: Supported (0x0)

Type	: E-Eth	VcId	: 103
SvcId	: 4	SdpId	: 1
Peer Address	: 1.1.1.57	Vc-switching	: No
LMTU	: 1500	RMTU	: 1500
Egr. Lbl	: 131065	Egr. Ctl Word	: No
Egr. Flags	: None	Egr. Status Bits	: N/A
Ing. Lbl	: 131066W	Ing. Ctl Word	: No
Ing. Flags	: None	Ing. Status Bits	: Supported (0x16)

Type	: E-Eth	VcId	: 104
SvcId	: 5	SdpId	: 1
Peer Address	: 1.1.1.57	Vc-switching	: Yes 1:105
LMTU	: 1500	RMTU	: 0
Egr. Lbl	: --	Egr. Ctl Word	: No
Egr. Flags	: None	Egr. Status Bits	: N/A
Ing. Lbl	: 131065U	Ing. Ctl Word	: No
Ing. Flags	: None	Ing. Status Bits	: Supported (0x18)

Type	: E-Eth	VcId	: 105
SvcId	: 5	SdpId	: 1
Peer Address	: 1.1.1.57	Vc-switching	: Yes 1:104
LMTU	: 1500	RMTU	: 0
Egr. Lbl	: --	Egr. Ctl Word	: No
Egr. Flags	: None	Egr. Status Bits	: N/A
Ing. Lbl	: 131064U	Ing. Ctl Word	: No
Ing. Flags	: None	Ing. Status Bits	: Supported (0x18)

Type	: E-Eth	VcId	: 106
SvcId	: 6	SdpId	: 1
Peer Address	: 1.1.1.57	Vc-switching	: No
LMTU	: 1500	RMTU	: 1500
Egr. Lbl	: 131064D	Egr. Ctl Word	: Yes
Egr. Flags	: None	Egr. Status Bits	: Supported (0x16)
Ing. Lbl	: 131063U	Ing. Ctl Word	: No
Ing. Flags	: None	Ing. Status Bits	: Supported (0x0)

Type	: E-Eth	VcId	: 107
SvcId	: 7	SdpId	: 1
Peer Address	: 1.1.1.57	Vc-switching	: No
LMTU	: 1500	RMTU	: 0
Egr. Lbl	: --	Egr. Ctl Word	: No
Egr. Flags	: None	Egr. Status Bits	: N/A
Ing. Lbl	: 131062U	Ing. Ctl Word	: No
Ing. Flags	: None	Ing. Status Bits	: Supported (0x0)

```
=====
No. of VC Labels: 8
=====
A:cpm-a#
```

Sample Output - show router ldp bindings session

```
ALU-12# show router ldp bindings session 10.10.10.104
=====
LDP LSR ID: 10.10.10.103
=====
Legend:  U - Label In Use,  R - Label Released
=====
LDP Prefix Bindings
=====
Prefix                Peer                IngLbl  EgrLbl  EgrIntf        EgrNextHop
-----
No Matching Entries Found

=====
LDP Service FEC 128 Bindings
=====
Type  VCId      SvcId      SDPIId  Peer                IngLbl  EgrLbl  LMTU  RMTU
-----
Ukwn  222        Ukwn       Ukwn    10.10.10.104      --      131071  0      0
VPLS  700        700        2       10.10.10.104      131071U 131070  1514  0
-----
No. of Service Bindings: 2

=====
LDP Service FEC 129 Bindings
=====
AGI                                SAII                                TAII
Type                SvcId      SDPIId  Peer                IngLbl  EgrLbl  LMTU  RMTU
-----
No Matching Entries Found
=====
ALU-12#
```

Sample Output - show router ldp bindings active

```
ALU-12# show router ldp bindings active
=====
LDP Prefix Bindings (Active)
=====
Prefix                Op                IngLbl  EgrLbl  EgrIntf        EgrNextHop
-----
10.20.1.3/32          Push             --      131069  1/1/6          20.1.1.1
10.20.1.10/32         Pop              131069  --      --              --
-----
No. of Prefix Bindings: 2
=====
ALU-12#
```

discovery

Syntax	discovery [{ peer [<i>ip-address</i>]} { interface [<i>ip-int-name</i>]}] [state <i>state</i>] [detail]
Context	show>router>ldp
Description	This command displays the status of the interfaces participating in LDP discovery.
Parameters	<p><i>ip-address</i> — specifies the IP address of the peer</p> <p><i>ip-int-name</i> — specifies an existing interface. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.</p> <p><i>state</i> — specifies the current operational state of the adjacency</p> <p>detail — displays detailed information</p>
Output	The following outputs are examples of LDP discovery information, and Table 31 describes the fields.

Sample Output - show router ldp discovery

```

ALU-12# show router ldp discovery
=====
LDP Hello Adjacencies
=====
Interface Name                Local Addr      Peer Addr      AdjType State
-----
N/A                          10.10.10.103    10.10.10.93    Targ    Trying
N/A                          10.10.10.103    10.10.10.104    Targ    Estab
to-104                       10.0.0.103      224.0.0.2      Link    Trying
-----
No. of Hello Adjacencies: 3
=====
ALU-12#

```

Sample Output - show router ldp discovery detail

```

ALU-12# show router ldp discovery detail
=====
LDP Hello Adjacencies (Detail)
=====
Peer 10.10.10.93
-----
Local Address      : 10.10.10.103      Peer Address      : 10.10.10.93
Adjacency Type     : Targeted          State             : Trying
-----
Peer 10.10.10.104
-----
Local Address      : 10.10.10.103      Peer Address      : 10.10.10.104
Adjacency Type     : Targeted          State             : Established
Up Time           : 0d 18:26:36      Hold Time Remaining: 38
Hello Mesg Recv    : 76616920      Hello Mesg Sent   : 466580812
Remote Cfg Seq No  : 159            Remote IP Address  : 10.10.10.104
Local Cfg Seq No   : 1674451      Local IP Address   : 0.224.173.172

```

```

-----
Interface "to-104"
-----
Local Address      : 10.0.0.103      Peer Address      : 224.0.0.2
Adjacency Type    : Link           State             : Trying
=====
ALU-12#

```

Table 31: 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
State	Established — indicates that the adjacency is established
	Trying — indicates that 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 message 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 message 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 [<i>ip-int-name</i> <i>ip-address</i>] [detail]
Context	show>router>ldp
Description	This command displays configuration information about LDP interfaces.
Parameters	<p><i>ip-int-name</i> — specifies an existing interface. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.</p> <p><i>ip-address</i> — identifies the LDP neighbor by IP address</p> <p>detail — displays detailed information</p>
Output	The following output is an example of LDP interface information, and Table 32 describes the fields.

Sample Output

```

A:ALU-12# show router ldp interface
=====
LDP Interfaces
=====
Interface                               Adm Opr  Hello  Hold  KA      KA      Transport
                               Factor Time  Factor Timeout Address
-----
i2_1/1                               UpUp    3      15    3       30     System
-----
No. of Interfaces: 1
=====
A:ALU-12#

A:ALU-12>show>router>ldp# interface detail
=====
LDP Interfaces (Detail)
=====
Interface "back"
-----
Admin State      : Up          Oper State      : Down
Oper Down Reason : interfaceDown
Hold Time        : 1000        Hello Factor    : 15
Keepalive Timeout : 1000        Keepalive Factor : 15
Transport Addr   : System      Last Modified   : 08/08/2007 09:50:15
Active Adjacencies : 0
Tunneling        : Disabled
Lsp Name         : None
=====
A:ALU-12>show>router>ldp#

```

Table 32: LDP Interface Output Fields

Label	Description
Interface	The interface associated with the LDP instance
Adm	Up — indicates that the LDP is administratively enabled
	Down — indicates that the LDP is administratively disabled
Opr	Up — indicates that the LDP is operationally enabled
	Down — indicates that the LDP is operationally disabled
Hello Factor	The value by which the hello timeout should be divided to give the hello time; that is, 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 time interval, in seconds, that LDP waits before declaring a neighbor to be down. Hold time (also known as Hello time) 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; that is, 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 three times the keepalive time (the time interval between successive LDP keepalive messages).
Transport Address	The transport address entity
No. of Interfaces	The total number of LDP interfaces
Oper Down Reason	The reason for the LSP being in the down state
Active Adjacencies	The number of active adjacencies
Last Modified	The time of the last modification to the LDP interface
Lsp Name	The LSP name

parameters

Syntax	parameters
Context	show>router>ldp
Description	This command displays configuration information about LDP parameters.
Output	The following output is an example of LDP parameters information, and Table 33 describes the fields.

Sample Output

```

A:ALU-12# show router ldp parameters
=====
LDP Parameters (LSR ID 10.10.10.103)
=====
-----
Graceful Restart Parameters
-----
Nbor Liveness Time : 120 sec           Max Recovery Time : 120
-----
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   : 5000 sec         Keepalive Factor    : 255
Hold Time           : 5000 sec         Hello Factor        : 255
Passive Mode        : False            Targeted Sessions   : Enabled
=====
A:ALU-12#

```

Table 33: LDP Parameters Output Fields

Label	Description
Graceful Restart Parameters	
Nbr Liveliness Time	The neighbor liveliness time
Max Recovery Time	The local maximum recovery time
Interface Parameters	
Keepalive Timeout	The time interval, in seconds, that LDP waits before tearing down a session. If no LDP messages are exchanged during this time interval, the LDP session is torn down. Generally the value is configured to be three times the keepalive time (the time interval between successive LDP keepalive messages).
Keepalive Factor	The value by which the keepalive timeout should be divided to give the keepalive time; that is, the time interval, in seconds, between LDP keepalive messages. LDP keepalive messages are sent to keep the LDP session from timing out when no other LDP traffic is being sent between the neighbors.
Hold Time	The time interval, in seconds, that LDP waits before declaring a neighbor to be down. Hold time (also known as Hello time) is local to the system and is sent in the hello messages to a neighbor.
Hello Factor	The value by which the hello timeout should be divided to give the hello time; that is, 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.
Propagate Policy	Specifies whether the LSR should generate FECs and which FECs it should generate
	system — indicates that the LDP will distribute label bindings only for the router's system IP address
	interface — indicates that the LDP will distribute label bindings for all LDP interfaces
	all — indicates that the LDP will distribute label bindings for all prefixes in the routing table
	none — indicates that the LDP will not distribute any label bindings

Table 33: LDP Parameters Output Fields (Continued)

Label	Description
Transport Address	<code>interface</code> — the interface IP address is used to set up the LDP session between neighbors. If multiple interfaces exist between two neighbors, the interface mode cannot be used since only one LDP session is actually set up between the two neighbors.
	<code>system</code> — the system IP address is used to set up the LDP session between neighbors
Label-Distribution	The label distribution method
Label-Retention	<code>liberal</code> — all advertised label mappings are retained whether they are from a valid next hop or not. When the label distribution value is downstream unsolicited, a router may receive label bindings for the same destination for all its neighbors. Labels for the non-next-hops for the FECs are retained in the software but not used. When a network topology change occurs where a non-next-hop becomes a true next hop, the label received earlier is then used.
	<code>conservative</code> — advertised label mappings are retained only if they will be used to forward packets; for example if the label came from a valid next hop. Label bindings received from non-next-hops for each FEC are discarded.
Control Mode	<code>ordered</code> — label bindings are not distributed in response to a label request until a label binding has been received from the next hop for the destination
	<code>independent</code> — label bindings are distributed immediately in response to a label request even if a label binding has not yet been received from the next hop for the destination
Route Preference	The route preference assigned to LDP routes. When multiple routes are available to a destination, the route with the lowest preference will be used. This value is only applicable to LDP interfaces and not for targeted sessions.
Targeted Session Parameters	
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

Table 33: LDP Parameters Output Fields (Continued)

Label	Description
Hello Factor	The value by which the hello timeout should be divided to give the hello time; that is, 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.
	Disabled — indicates that no authentication is being used
Passive Mode	True — indicates that LDP responds only when it gets a connect request from a peer and will not attempt to actively connect to its neighbors
	False — indicates that LDP actively tries to connect to its peers
Targeted Sessions	Enabled — indicates that targeted sessions are enabled
	Disabled — indicates that targeted sessions are disabled

peer

Syntax **peer** [*ip-address*] [**detail**]

Context show>router>ldp

Description This command displays configuration information about LDP peers.

Parameters *ip-address* — specifies the IP address of the LDP peer
detail — displays detailed information

Output The following output is an example of LDP peer information, and [Table 34](#) describes the fields.

Sample Output

```
A:ALU-12# show router ldp peer
=====
LDP Peers
=====
Peer           Adm  Opr  Hello  Hold   KA    KA    Passive  Auto
                Factor Time  Factor Timeout Mode   Created
-----
10.10.10.93    Up   Up   3      45     4     40     Disabled Yes
10.10.10.104   Up   Up   3      45     4     40     Disabled Yes
-----
No. of Peers: 2
=====
A:ALU-12#
```

```

A:ALU-12# show router ldp peer detail
=====
LDP Peers (Detail)
=====
-----
Peer 1.2.3.4
-----
Admin State      : Up           Oper State       : Down
Hold Time       : 45           Hello Factor     : 3
Keepalive Timeout : 40        Keepalive Factor : 4
Passive Mode    : Disabled    Last Modified    : 05/01/2008 21:44:17
Active Adjacencies : 0       Auto Created     : No
Tunneling       : None
Lsp Name        : None
=====
A:ALU-12#

```

Table 34: LDP Peer Output Fields

Label	Description
Peer	The IP address of the peer
Adm	Up — indicates that LDP is administratively enabled
	Down — indicates that LDP is administratively disabled
Opr	Up — indicates that LDP is operationally enabled
	Down — indicates that LDP is operationally disabled
Hello Factor	The value by which the hello timeout should be divided to give the hello time; that is, 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 time interval, in seconds, that LDP waits before declaring a neighbor to be down. Hold time (also known as Hello time) 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; that is, 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 three times the keepalive time (the time interval between successive LDP keepalive messages).

Table 34: LDP Peer Output Fields (Continued)

Label	Description
Passive Mode	The mode used to set up LDP sessions. This value is only applicable to targeted sessions and not to LDP interfaces. This mode is always set to False.
	True — indicates that LDP responds only when it gets a connect request from a peer and will not attempt to actively connect to its neighbors
	False — indicates that LDP actively tries to connect to its peers
Auto Create	Specifies whether or not a targeted peer was automatically created through a Service Manager. For an LDP interface, this value is always false.
No. of Peers	The total number of LDP peers
LSP	The LSP name

peer-parameters

Syntax	peer-parameters <i>peer-ip-address</i>
Context	show>router>ldp
Description	This command displays LDP peer information.
Parameters	<i>peer-ip-address</i> — specifies the peer IP address
Output	The following output is an example of LDP peer-parameters information, and Table 35 describes the fields.

Sample Output

```

A:ALU-214># show router ldp peer-parameters
=====
LDP Peers
=====
Peer          TTL Security Min-TTL-Value Authentication Auth key chain
-----
10.10.10.104   Disabled      n/a           Enabled      n/a
-----
No. of Peers: 1
=====
A:ALU-214>#

```

Table 35: LDP Peer-Parameter Output Fields

Label	Description
Peer	The IP address of the peer
Authentication	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*]]

Context show>router>ldp

Description This command displays configuration information about LDP sessions.

Parameters *ip-addr* — specifies the IP address of the LDP peer

label-space — specifies the label space identifier that the router is advertising on the interface

Values 0 to 65535

detail — displays detailed information

packet-type — specifies the packet type

Values hello, keepalive, init, label, notification, address

Output The following output is an example of LDP session information, and [Table 36](#) describes the fields.

Sample Output

```
ALU-12# show router ldp session
=====
LDP Sessions
=====
Peer LDP Id           Adj Type State           Msg Sent  Msg Recv  Up Time
-----
10.10.10.104:0        Targeted Established  13943     13947     0d 21:12:41
-----
No. of Sessions: 1
=====
ALU-12#
```

```
A:cpm-a# show router ldp session detail
=====
LDP Sessions (Detail)
=====
Session with Peer 1.1.1.33:0
-----
Adjacency Type      : Link          State          : Established
Up Time             : 0d 00:03:51
Max PDU Length      : 4096          KA/Hold Time Remaining: 26
Link Adjacencies    : 1            Targeted Adjacencies : 0
Local Address       : 1.1.1.30      Peer Address      : 1.1.1.33
Local TCP Port      : 646          Peer TCP Port     : 50232
Local KA Timeout    : 30           Peer KA Timeout   : 30
Msg Sent            : 89           Msg Recv         : 126
FECs Sent           : 3           FECs Recv        : 3
GR State            : Not Capable
Nbr Liveness Time   : 0            Max Recovery Time  : 0
Number of Restart   : 0            Last Restart Time  : Never
Advertise           : Address
-----
Session with Peer 1.1.1.57:0
-----
Adjacency Type      : Targeted      State          : Established
Up Time             : 0d 00:03:49
Max PDU Length      : 4096          KA/Hold Time Remaining: 36
Link Adjacencies    : 0            Targeted Adjacencies : 1
Local Address       : 1.1.1.30      Peer Address      : 1.1.1.57
Local TCP Port      : 646          Peer TCP Port     : 49574
Local KA Timeout    : 40           Peer KA Timeout   : 40
Msg Sent            : 55           Msg Recv         : 61
FECs Sent           : 11           FECs Recv        : 8
GR State            : Not Capable
Nbr Liveness Time   : 0            Max Recovery Time  : 0
Number of Restart   : 0            Last Restart Time  : Never
Advertise           : Address/Servi*
=====
A:cpm-a#
```

Table 36: 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 that 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
Msg Sent	The number of messages sent

Table 36: LDP Session Output Fields (Continued)

Label	Description
Msg Rcvd	The number of messages received
Up Time	The amount of time the adjacency has been enabled

status

Syntax	status
Context	show>router>ldp
Description	This command displays LDP status information.
Output	The following output is an example of LDP status information, and Table 37 describes the fields.

Sample Output

```
*A:csasim2>show>router>ldp# status

=====
LDP Status for LSR ID 10.10.10.32
=====
Admin State      : Up          Oper State      : Up
Created at       : 05/01/2008 16:12:07 Up Time         : 3d 23:31:22
Oper Down Reason : n/a         Oper Down Events : 0
Last Change      : 05/02/2008 16:49:01 Tunn Down Damp Time : 3 sec
Import Policies  :              Export Policies    :
    test-policy1              None
Active Adjacencies : 0          Active Sessions : 0
Active Interfaces  : 0          Inactive Interfaces : 1
Active Peers       : 0          Inactive Peers     : 0
Addr FECs Sent     : 0          Addr FECs Recv     : 0
Serv FECs Sent     : 0          Serv FECs Recv     : 0
Attempted Sessions : 0
No Hello Err       : 0          Param Adv Err      : 0
Max PDU Err        : 0          Label Range Err    : 0
Bad LDP Id Err     : 0          Bad PDU Len Err    : 0
Bad Mesg Len Err   : 0          Bad TLV Len Err    : 0
Malformed TLV Err  : 0          Keepalive Expired Err: 0
Shutdown Notif Sent: 0          Shutdown Notif Recv : 0
=====
*A:csasim2>show>router>ldp#
```

Table 37: LDP Status Output Fields

Label	Description
Admin State	Up — indicates that LDP is administratively enabled
	Down — indicates that LDP is administratively disabled
Oper State	Up — indicates that LDP is operationally enabled
	Down — indicates that LDP is operationally disabled
Created at	The date and time that the LDP instance was created
Up Time	The time, in hundredths of seconds, that the LDP instance has been operationally up
Oper Down Time	The time, in hundredths of seconds, that the LDP instance has been operationally down
Oper Down Events	The number of times the LDP instance has gone operationally down since the instance was created
Last Change	The date and time that the LDP instance was last modified
Import Policies	The import policy associated with the LDP instance
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

Table 37: LDP Status Output Fields (Continued)

Label	Description
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
Keepalive Expired Err	The number of session keepalive timer expired 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
Shutdown Notif Recv	The number of shutdown notifications received related to sessions associated with this LDP instance

Clear Commands

instance

Syntax	instance
Context	clear>router>ldp
Description	This command resets the LDP instance.

interface

Syntax	interface <i>ip-int-name</i> [statistics]
Context	clear>router>ldp
Description	This command restarts or clears statistics for LDP interfaces.
Parameters	<i>ip-int-name</i> — specifies an existing interface. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes. statistics — clears only the statistics for an interface

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> — specifies a targeted peer statistics — clears only the statistics for a targeted peer

session

Syntax	session <i>ip-addr</i> [: <i>label-space</i>] [statistics]
Context	clear>router>ldp
Description	This command restarts or clears statistics for LDP sessions.
Parameters	<i>ip-addr</i> — specifies the IP address of the LDP peer <i>label-space</i> — specifies the label space identifier that the router is advertising on the interface
Values	0 to 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.

```
ALU-12# debug router ldp peer 10.10.10.104
ALU-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
ALU-12>debug>router>ldp#
```

ldp

Syntax	[no] ldp
Context	debug>router
Description	This command configures LDP debugging.

interface

Syntax	[no] interface <i>interface-name</i>
Context	debug>router>ldp
Description	This command configures debugging for a specific LDP interface.
Parameters	<i>interface-name</i> — specifies an existing interface

peer

Syntax	[no] peer <i>ip-address</i>
Context	debug>router>ldp
Description	This command configures debugging for a specific LDP peer.
Parameters	<i>ip-address</i> — specifies the LDP peer to debug

event

Syntax	[no] event
Context	debug>router>ldp>interface <i>interface-name</i> debug>router>ldp>peer <i>ip-address</i>
Description	This command configures debugging for specific LDP events.

bindings

Syntax	[no] bindings
Context	debug>router>ldp>peer <i>ip-address</i> >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>interface <i>interface-name</i> >event debug>router>ldp>peer <i>ip-address</i> >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	[no] packet
Context	debug>router>ldp>interface <i>interface-name</i> debug>router>ldp>peer <i>ip-address</i>
Description	This command enables debugging for specific LDP packets. The no form of the command disables the debugging output.

hello

Syntax	hello [detail] no hello
Context	debug>router>ldp>interface <i>interface-name</i> >packet debug>router>ldp>peer <i>ip-address</i> >packet
Description	This command enables debugging for sent and received 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 <i>ip-address</i> >packet
Description	This command enables debugging for LDP Init packets. The detail option displays detailed information on the type length value (TLV) included in mac-flush 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 <i>ip-address</i> >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 neighbor <i>ip-address</i> >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 802.1ag	Service Layer OAM
IEEE 802.1p/q	VLAN Tagging
IEEE 802.3	10BaseT
IEEE 802.3ah	Ethernet OAM
IEEE 802.3u	100BaseTX
IEEE 802.3x	Flow Control
IEEE 802.3z	1000BaseSX/LX
IEEE 802.3-2008	Revised base standard
ITU-T Y.1731	OAM functions and mechanisms for Ethernet-based networks

Telecom Compliance

IC CS-03 Issue 9	Spectrum Management and Telecommunications
ACTA TIA-968-A	
AS/ACIF S016 (Australia/New Zealand)	Requirements for Customer Equipment for connection to hierarchical digital interfaces
ITU-T G.703	Physical/electrical characteristics of hierarchical digital interfaces
ITU-T G.707	Network node interface for the Synchronous Digital Hierarchy (SDH)
ITU-T G.712-2001	Transmission performance characteristics of pulse code modulation channels
ITU-T G.957	Optical interfaces for equipments and systems relating to the synchronous digital hierarchy
ITU-T V.24	List of definitions for interchange circuits between data terminal equipment (DTE) and data circuit- terminating equipment (DCE)
ITU-T V.36	Modems for synchronous data transmission using 60-108 kHz group band circuits
ITU-T X.21	Interface between Data Terminal Equipment and Data Circuit- Terminating Equipment for Synchronous Operation on Public Data Networks

Protocol Support

ATM

RFC 2514	Definitions of Textual Conventions and OBJECT_IDENTITIES for ATM Management, February 1999
RFC 2515	Definition of Managed Objects for ATM Management, February 1999
RFC 2684	Multiprotocol Encapsulation over ATM Adaptation Layer 5
af-tm-0121.000	Traffic Management Specification Version 4.1, March 1999
ITU-T Recommendation I.610	B-ISDN Operation and Maintenance Principles and Functions version 11/95
ITU-T Recommendation I.432.1	B-ISDN user- network interface - Physical layer specification: General characteristics
GR-1248-CORE	Generic Requirements for Operations of ATM Network Elements (NEs). Issue 3 June 1996
GR-1113-CORE	Bellcore, Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements, Issue 1, July 1994
AF-PHY-0086.001	Inverse Multiplexing for ATM (IMA)

BFD

draft-ietf-bfd-mib-00.txt	Bidirectional Forwarding Detection Management Information Base
draft-ietf-bfd-base-o5.txt	Bidirectional Forwarding Detection
draft-ietf-bfd-v4v6-1hop-06.txt	BFD IPv4 and IPv6 (Single Hop)
draft-ietf-bfd-multihop-06.txt	BFD for Multi-hop Paths

BGP

- RFC 1397 BGP Default Route Advertisement
- RFC 1997 BGP Communities Attribute
- RFC 2385 Protection of BGP Sessions via MDS
- RFC 2439 BGP Route Flap Dampening
- RFC 2547bis BGP/MPLS VPNs
- RFC 2918 Route Refresh Capability for BGP-4
- RFC 3107 Carrying Label Information in BGP-4
- RFC 3392 Capabilities Advertisement with BGP-4
- RFC 4271 BGP-4 (previously RFC 1771)
- RFC 4360 BGP Extended Communities Attribute
- RFC 4364 BGP/MPLS IP Virtual Private Networks (VPNs) (previously RFC 2574bis BGP/MPLS VPNs)
- RFC 4456 BGP Route Reflection: Alternative to Full-mesh IBGP (previously RFC 1966 and RFC 2796)
- RFC 4724 Graceful Restart Mechanism for BGP - GR Helper
- RFC 4760 Multi-protocol Extensions for BGP (previously RFC 2858)
- RFC 4893 BGP Support for Four-octet AS Number Space

DHCP/DHCPv6

- RFC 1534 Interoperation between DHCP and BOOTP
- RFC 2131 Dynamic Host Configuration Protocol (REV)
- RFC 3046 DHCP Relay Agent Information Option (Option 82)
- RFC 3315 Dynamic Host Configuration Protocol for IPv6

DIFFERENTIATED SERVICES

- RFC 2474 Definition of the DS Field in the IPv4 and IPv6 Headers
- RFC 2597 Assured Forwarding PHB Group
- RFC 2598 An Expedited Forwarding PHB
- RFC 3140 Per-Hop Behavior Identification Codes

DIGITAL DATA NETWORK MANAGEMENT V.35

- RS-232 (also known as EIA/TIA-232)

GRE

- RFC 2784 Generic Routing Encapsulation (GRE)

IPv6

- RFC 2460 Internet Protocol, Version 6 (IPv6) Specification
- RFC 2462 IPv6 Stateless Address Autoconfiguration
- RFC 2464 Transmission of IPv6 Packets over Ethernet Networks
- RFC 3587 IPv6 Global Unicast Address Format
- RFC 3595 Textual Conventions for IPv6 Flow Label
- RFC 4007 IPv6 Scoped Address Architecture
- RFC 4193 Unique Local IPv6 Unicast Addresses
- RFC 4291 IPv6 Addressing Architecture
- RFC 4443 Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 Specification
- RFC 4649 DHCPv6 Relay Agent Remote-ID Option
- RFC 4861 Neighbor Discovery for IP version 6 (IPv6)

LDP

- RFC 5036 LDP Specification

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 (IS-IS) 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 4205 for Shared Risk Link Group (SRLG) TLV draft-ietf-isis-igp-p2p-over-lan-05.txt
- RFC 5309 Point-to-Point Operation over LAN in Link State Routing Protocols

MPLS

- RFC 3031 MPLS Architecture
- RFC 3032 MPLS Label Stack Encoding
- RFC 3815 Definitions of Managed Objects for the Multiprotocol Label Switching (MPLS), Label Distribution Protocol (LDP)
- RFC 4379 Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures

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 1305 Network Time Protocol (Version 3) Specification, Implementation and Analysis
- RFC 1850 OSPF-MIB
- RFC 1907 SNMPv2-MIB
- RFC 2011 IP-MIB
- RFC 2012 TCP-MIB
- RFC 2013 UDP-MIB
- RFC 2030 Simple Network Time Protocol (SNTP) Version 4 for IPv4, IPv6 and OSI
- 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-BASED-SMMIB
- RFC 2575 SNMP-VIEW-BASED ACM-MIB
- RFC 2576 SNMP-COMMUNITY-MIB
- RFC 2588 SONET-MIB
- RFC 2665 EtherLike-MIB
- RFC 2819 RMON-MIB
- RFC 2863 IF-MIB
- RFC 2864 INVERTED-STACK-MIB
- RFC 3014 NOTIFICATION-LOG MIB
- RFC 3164 The BSD Syslog Protocol
- RFC 3273 HCRMON-MIB
- RFC 3411 An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks

- RFC 3412 Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)
- RFC 3413 Simple Network Management Protocol (SNMP) Applications
- RFC 3414 User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)
- RFC 3418 SNMP MIB
- draft-ietf-disman-alarm-mib-04.txt
- draft-ietf-mpls-ldp-mib-07.txt
- draft-ietf-ospf-mib-update-04.txt
- draft-ietf-mpls-lsr-mib-06.txt
- draft-ietf-mpls-te-mib-04.txt
- IANA-IFType-MIB

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 3630 Traffic Engineering (TE) Extensions to OSPF
- RFC 4203 Shared Risk Link Group (SRLG) sub-TLV

PPP

- RFC 1332 PPP Internet Protocol Control Protocol (IPCP)
- RFC 1570 PPP LCP Extensions
- RFC 1619 PPP over SONET/SDH
- RFC 1661 The Point-to-Point Protocol (PPP)
- RFC 1662 PPP in HDLC-like Framing
- RFC 1989 PPP Link Quality Monitoring
- RFC 1990 The PPP Multilink Protocol (MP)
- RFC 2686 The Multi-Class Extension to Multi-Link PPP

PSEUDOWIRES

- RFC 3550 RTP: A Transport Protocol for Real-Time Applications
- RFC 3985 Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture
- RFC 4385 Pseudowire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN
- RFC 4446 IANA Allocation for PWE3
- RFC 4447 Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)

RFC 4448 Encapsulation Methods for Transport of Ethernet over MPLS Networks
RFC 4553 Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP)
RFC 4717 Encapsulation Methods for Transport of Asynchronous Transfer Mode (ATM) over MPLS Networks
RFC 5085 Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires
RFC 5086 Structure-Aware Time Division Multiplexed (TDM) Circuit Emulation Service over Packet Switched Network (CESoPSN)
draft-ietf-pwe3-redundancy-02 Pseudowire (PW) Redundancy

RADIUS

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

RSVP-TE and FRR

RFC 2430 A Provider Architecture for DiffServ & TE
RFC 2961 RSVP Refresh Overhead Reduction Extensions
RFC 2702 Requirements for Traffic Engineering over MPLS
RFC 2747 RSVP Cryptographic Authentication
RFC 3097 RSVP Cryptographic Authentication - Updated Message Type Value
RFC 3209 Extensions to RSVP for LSP Tunnels
RFC 3210 Applicability Statement for Extensions to RSVP for LSP Tunnels
RFC 4090 Fast Reroute Extensions to RSVP-TE for LSP Tunnels

SONET/SDH

GR-253-CORE SONET Transport Systems: Common Generic Criteria. Issue 3, September 2000
ITU-T Recommendation G.841 Telecommunication Standardization Section of ITU, Types and Characteristics of SDH Networks Protection Architecture, issued in October 1998 and as augmented by Corrigendum1 issued in July 2002

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

SYNCHRONIZATION

G.813 Timing characteristics of SDH equipment slave clocks (SEC)
G.8261 Timing and synchronization aspects in packet networks
G.8262 Timing characteristics of synchronous Ethernet equipment slave clock
GR 1244 CORE Clocks for the Synchronized Network: Common Generic Criteria
IEEE 1588v2 1588 PTP 2008

TACACS+

IETF draft-grant-tacacs-02.txt The TACACS+ Protocol

TCP/IP

RFC 768 User Datagram Protocol
RFC 791 Internet Protocol
RFC 792 Internet Control Message Protocol
RFC 793 Transmission Control Protocol
RFC 826 Ethernet Address Resolution Protocol
RFC 854 Telnet Protocol Specification
RFC 1350 The TFTP Protocol (Rev. 2)
RFC 1812 Requirements for IPv4 Routers

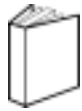
VPLS

RFC 4762 Virtual Private LAN Services Using LDP

Proprietary MIBs

TIMETRA-ATM-MIB.mib
TIMETRA-CAPABILITY-7705-V1.mib
TIMETRA-CFLOWD-MIB.mib
TIMETRA-CHASSIS-MIB.mib
TIMETRA-CLEAR-MIB.mib
TIMETRA-FILTER-MIB.mib
TIMETRA-GLOBAL-MIB.mib
TIMETRA-LDP-MIB.mib
TIMETRA-LOG-MIB.mib
TIMETRA-MPLS-MIB.mib
TIMETRA-OAM-TEST-MIB.mib
TIMETRA-PORT-MIB.mib
TIMETRA-PPP-MIB.mib
TIMETRA-QOS-MIB.mib
TIMETRA-ROUTE-POLICY-MIB.mib
TIMETRA-RSVP-MIB.mib
TIMETRA-SAP-MIB.mib
TIMETRA-SDP-MIB.mib
TIMETRA-SECURITY-MIB.mib
TIMETRA-SERV-MIB.mib
TIMETRA-SYSTEM-MIB.mib
TIMETRA-TC-MIB.mib

Customer documentation and product support



Customer documentation

<http://www.alcatel-lucent.com/myaccess>

Product manuals and documentation updates are available at [alcatel-lucent.com](http://www.alcatel-lucent.com). If you are a new user and require access to this service, please contact your Alcatel-Lucent sales representative.



Technical Support

<http://support.alcatel-lucent.com>



Documentation feedback

documentation.feedback@alcatel-lucent.com

