MPLS Basic Configuration Guide, Cisco IOS XE Everest 3.18SP (Cisco NCS 4200 Series)

First Published: 2017-05-17

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

CONTENTS

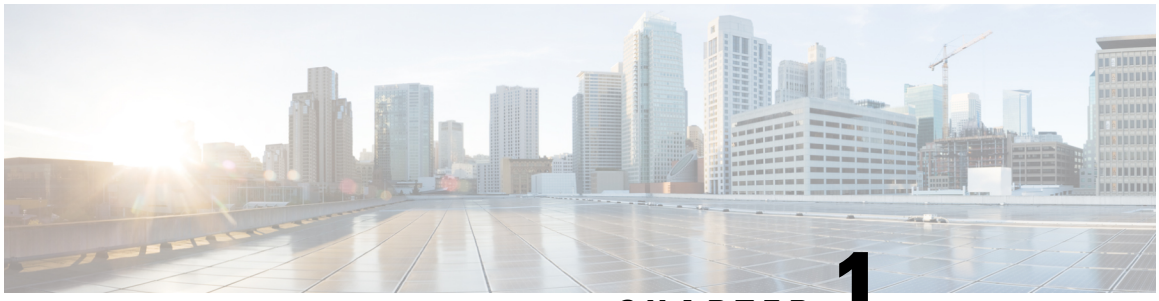

CHAPTER 1

Multiprotocol Label Switching (MPLS) on Cisco Routers

This document describes commands for configuring and monitoring Multiprotocol Label Switching (MPLS) functionality on Cisco routers and switches. This document is a companion to other feature modules describing other MPLS applications.

Information About MPLS

MPLS Overview

Multiprotocol label switching (MPLS) combines the performance and capabilities of Layer 2 (data link layer) switching with the proven scalability of Layer 3 (network layer) routing. MPLS enables service providers to meet the challenges of explosive growth in network utilization while providing the opportunity to differentiate services without sacrificing the existing network infrastructure. The MPLS architecture is flexible and can be employed in any combination of Layer 2 technologies. MPLS support is offered for all Layer 3 protocols, and scaling is possible well beyond that typically offered in today's networks.

MPLS efficiently enables the delivery of IP services over an ATM switched network. MPLS supports the creation of different routes between a source and a destination on a purely router-based Internet backbone. By incorporating MPLS into their network architecture, service providers can save money, increase revenue and productivity, provide differentiated services, and gain competitive advantages.

Note In the Cisco IOS XE Release 16.x, the ASR 1000 routers only support fragmentation of the MPLS packets from the IP to MPLS direction.

Functional Description of MPLS

Label switching is a high-performance packet forwarding technology that integrates the performance and traffic management capabilities of data link layer (Layer 2) switching with the scalability, flexibility, and performance of network layer (Layer 3) routing.

Label Switching Functions

In conventional Layer 3 forwarding mechanisms, as a packet traverses the network, each router extracts all the information relevant to forwarding the packet from the Layer 3 header. This information is then used as an index for a routing table lookup to determine the next hop for the packet.

In the most common case, the only relevant field in the header is the destination address field, but in some cases, other header fields might also be relevant. As a result, the header analysis must be done independently at each router through which the packet passes. In addition, a complicated table lookup must also be done at each router.

In label switching, the analysis of the Layer 3 header is done only once. The Layer 3 header is then mapped into a fixed length, unstructured value called a *label* .

Many different headers can map to the same label, as long as those headers always result in the same choice of next hop. In effect, a label represents a *forwarding equivalence class* --that is, a set of packets which, however different they may be, are indistinguishable by the forwarding function.

The initial choice of a label need not be based exclusively on the contents of the Layer 3 packet header; for example, forwarding decisions at subsequent hops can also be based on routing policy.

Once a label is assigned, a short label header is added at the front of the Layer 3 packet. This header is carried across the network as part of the packet. At subsequent hops through each MPLS router in the network, labels are swapped and forwarding decisions are made by means of MPLS forwarding table lookup for the label carried in the packet header. Hence, the packet header does not need to be reevaluated during packet transit through the network. Because the label is of fixed length and unstructured, the MPLS forwarding table lookup process is both straightforward and fast.

Distribution of Label Bindings

Each label switching router (LSR) in the network makes an independent, local decision to determine a label value to represent a forwarding equivalence class. This association is known as a label binding. Each LSR informs its neighbors of the label bindings it has made.

When a labeled packet is being sent from LSR A to the neighboring LSR B, the label value carried by the IP packet is the label value that LSR B assigned to represent the forwarding equivalence class of the packet. Thus, the label value changes as the IP packet traverses the network.

The awareness of label bindings by neighbouring routers is facilitated using the following protocols:

- Label Distribution Protocol (LDP) - Enables peer LSRs in an MPLS network to exchange label binding information for supporting hop-by-hop forwarding in an MPLS network.
- Tag Distribution Protocol (TDP) - Supports MPLS forwarding along normally routed paths.
- Resource Reservation Protocol (RSVP) - Supports MPLS traffic engineering.
- Border Gateway Protocol (BGP) - Supports MPLS virtual private networks (VPNs) .

Benefits of MPLS

MPLS provides the following major benefits to service provider networks:

Scalable support for Virtual Private Networks (VPNs)--MPLS enables VPN services to be supported in service provider networks, thereby greatly accelerating Internet growth.

The use of MPLS for VPNs provides an attractive alternative to the building of VPNs by means of either ATM or Frame Relay permanent virtual circuits (PVCs) or various forms of tunneling to interconnect routers at customer sites.

Unlike the PVC VPN model, the MPLS VPN model is highly scalable and can accommodate increasing numbers of sites and customers. The MPLS VPN model also supports "any-to-any" communication among VPN sites without requiring a full mesh of PVCs or the backhauling (suboptimal routing) of traffic across the service provider network. For each MPLS VPN user, the service provider's network appears to function as a private IP backbone over which the user can reach other sites within the VPN organization, but not the sites of any other VPN organization.

From a user perspective, the MPLS VPN model enables network routing to be dramatically simplified. For example, rather than having to manage routing over a topologically complex virtual backbone composed of many PVCs, an MPLS VPN user can generally employ the service provider's backbone as the default route in communicating with all of the other VPN sites.

Explicit routing capabilities (also called constraint-based routing or traffic engineering)--Explicit routing employs "constraint-based routing," in which the path for a traffic flow is the shortest path that meets the resource requirements (constraints) of the traffic flow.

In MPLS traffic engineering, factors such as bandwidth requirements, media requirements, and the priority of one traffic flow versus another can be taken into account. These traffic engineering capabilities enable the administrator of a service provider network to

- Control traffic flow in the network
- Reduce congestion in the network
- Make best use of network resources

Thus, the network administrator can specify the amount of traffic expected to flow between various points in the network (thereby establishing a traffic matrix), while relying on the routing system to

- Calculate the best paths for network traffic
- Set up the explicit paths to carry the traffic

Support for IP routing on ATM switches (also called IP and ATM integration)--MPLS enables an ATM switch to perform virtually all of the functions of an IP router. This capability of an ATM switch stems from the fact that the MPLS forwarding paradigm, namely, label swapping, is exactly the same as the forwarding paradigm provided by ATM switch hardware.

The key difference between a conventional ATM switch and an ATM label switch is the control software used by the latter to establish its virtual channel identifier (VCI) table entries. An ATM label switch uses IP routing protocols and the Tag Distribution Protocol (TDP) to establish VCI table entries.

An ATM label switch can function as a conventional ATM switch. In this dual mode, the ATM switch resources (such as VCI space and bandwidth) are partitioned between the MPLS control plane and the ATM control plane. The MPLS control plane provides IP-based services, while the ATM control plane supports ATM-oriented functions, such as circuit emulation or PVC services.

How to Configure MPLS

This section explains how to perform the basic configuration required to prepare a router for MPLS switching and forwarding.

Configuration tasks for other MPLS applications are described in the feature module documentation for the application.

Configuring a Router for MPLS Switching

MPLS switching on Cisco routers requires that Cisco Express Forwarding be enabled.

For more information about Cisco Express Forwarding commands, see the Cisco IOS Switching Command Reference.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **ip cef distributed**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**ip cef distributed** **Example:** `Device(config)# ip cef distributed`	Enables Cisco Express Forwarding on the route processor card.

Verifying Configuration of MPLS Switching

To verify that Cisco Express Forwarding has been configured properly, issue the **show ip cef summary** command, which generates output similar to that shown below:

SUMMARY STEPS

1. **show ip cef summary**

DETAILED STEPS

show ip cef summary

Example:

```
Router# show ip cef summary
IP CEF with switching (Table Version 49), flags=0x0
  43 routes, 0 resolve, 0 unresolved (0 old, 0 new)
  43 leaves, 49 nodes, 56756 bytes, 45 inserts, 2 invalidations
  2 load sharing elements, 672 bytes, 2 references
  1 CEF resets, 4 revisions of existing leaves
  4 in-place modifications
  refcounts:  7241 leaf, 7218 node
Adjacency Table has 18 adjacencies
Router#
```

Configuring a Router for MPLS Forwarding

MPLS forwarding on Cisco routers requires that forwarding of IPv4 packets be enabled.

For more information about MPLS forwarding commands, see the *Multiprotocol Label Switching Command Reference*.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **interface** *type slot/subslot /port* [**.** *subinterface*]
4. **mpls ip**
5. **end**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**interface** *type slot/subslot /port* [**.** *subinterface*] **Example:** `Device(config)# interface gigabitethernet 4/0/0`	Specifies the Gigabit Ethernet interface and enters interface configuration mode.

	Command or Action	Purpose
Step 4	**mpls ip** **Example:** `Device(config-if)# mpls ip`	Enables MPLS forwarding of IPv4 packets along normally routed paths for the Gigabit Ethernet interface.
Step 5	**end** **Example:** `Device(config-if)# end`	Exits interface configuration mode and returns to privileged EXEC mode.

What to do next

Configure either of the following:

- MPLS Label Distribution Protocol (LDP). For information about configuring MPLS LDP, see the *MPLS Label Distribution Protocol Configuration Guid*e.
- Static labels. For information about configuring static labels, see *MPLS Static Labels*.

Verifying Configuration of MPLS Forwarding

To verify that MPLS forwarding has been configured properly, issue the **show mpls interfaces detail** command, which generates output similar to that shown below:

SUMMARY STEPS

1. **show mpls interfaces detail**

DETAILED STEPS

show mpls interfaces detail

Example:

```
Device# show mpls interfaces detail

Interface GigabitEthernet1/0/0:
         IP labeling enabled (ldp)
         LSP Tunnel labeling not enabled
         MPLS operational
         MTU = 1500
Interface POS2/0/0:
         IP labeling enabled (ldp)
         LSP Tunnel labeling not enabled
         MPLS not operational
         MTU = 4470
```

Additional References

Related Documents

Related Topic	Document Title
Cisco IOS commands	*Cisco IOS Master Commands List, All Releases*
MPLS commands	*Cisco IOS Multiprotocol Label Switching Command Reference*

Standards

Standard	Title
The supported standards applicable to the MPLS applications appear in the respective feature module for the application.	--

MIBs

MIB	MIBs Link
The supported MIBs applicable to the MPLS applications appear in the respective feature module for the application.	To locate and download MIBs for selected platforms, Cisco software releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

RFCs

RFC	Title
The supported RFCs applicable to the MPLS applications appear in the respective feature module for the application.	--

Technical Assistance

Description	Link
The Cisco Support and Documentation website provides online resources to download documentation, software, and tools. Use these resources to install and configure the software and to troubleshoot and resolve technical issues with Cisco products and technologies. Access to most tools on the Cisco Support and Documentation website requires a Cisco.com user ID and password.	*Support & Downloads*

Glossary

BGP --Border Gateway Protocol. The predominant interdomain routing protocol used in IP networks.

Border Gateway Protocol --See BGP.

FIB --Forwarding Information Base. A table that contains a copy of the forwarding information in the IP routing table.

Forwarding Information Base --See FIB.

label --A short, fixed-length identifier that tells switching nodes how the data (packets or cells) should be forwarded.

label binding --An association between a label and a set of packets, which can be advertised to neighbors so that a label switched path can be established.

Label Distribution Protocol --See LDP.

Label Forwarding Information Base --See LFIB.

label imposition --The act of putting the first label on a packet.

label switching router --See LSR.

LDP --Label Distribution Protocol. The protocol that supports MPLS hop-by-hop forwarding by distributing bindings between labels and network prefixes.

LFIB --Label Forwarding Information Base. A data structure in which destinations and incoming labels are associated with outgoing interfaces and labels.

LSR --label switching router. A Layer 3 router that forwards a packet based on the value of an identifier encapsulated in the packet.

MPLS --Multiprotocol Label Switching. An industry standard on which label switching is based.

MPLS hop-by-hop forwarding --The forwarding of packets along normally routed paths using MPLS forwarding mechanisms.

Multiprotocol Label Switching --See MPLS.

Resource Reservation Protocol --See RSVP.

RIB --Routing Information Base. A common database containing all the routing protocols running on a router.

Routing Information Base --See RIB.

RSVP --Resource Reservation Protocol. A protocol for reserving network resources to provide quality of service guarantees to application flows.

traffic engineering --Techniques and processes used to cause routed traffic to travel through the network on a path other than the one that would have been chosen if standard routing methods were used.

Virtual Private Network --See VPN.

VPN --Virtual Private Network. A network that enables IP traffic to use tunneling to travel securely over a public TCP/IP network.

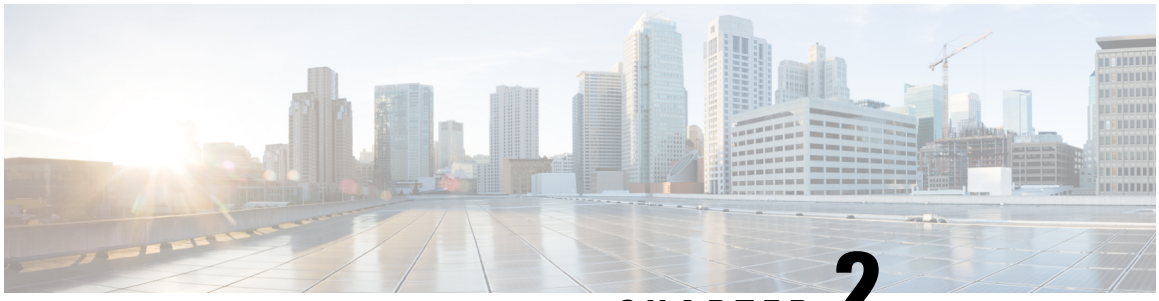

CHAPTER 2

MPLS LSP Ping, Traceroute, and AToM VCCV

As Multiprotocol Label Switching (MPLS) deployments increase and the traffic types they carry increase, the ability of service providers to monitor label switched paths (LSPs) and quickly isolate MPLS forwarding problems is critical to their ability to offer services. The MPLS LSP Ping, Traceroute, and AToM VCCV feature helps them mitigate these challenges.

The MPLS LSP Ping, Traceroute, and AToM VCCV feature can detect when an LSP fails to deliver user traffic.

- You can use MPLS LSP Ping to test LSP connectivity for IPv4 Label Distribution Protocol (LDP) prefixes, traffic engineering (TE) Forwarding Equivalence Classes (FECs), and AToM FECs.
- You can use MPLS LSP Traceroute to trace the LSPs for IPv4 LDP prefixes and TE tunnel FECs.
- Any Transport over MPLS Virtual Circuit Connection Verification (AToM VCCV) allows you to use MPLS LSP Ping to test the pseudowire (PW) section of an AToM virtual circuit (VC).

Internet Control Message Protocol (ICMP) ping and trace are often used to help diagnose the root cause when a forwarding failure occurs. The MPLS LSP Ping, Traceroute, and AToM VCCV feature extends this diagnostic and troubleshooting ability to the MPLS network and aids in the identification of inconsistencies between the IP and MPLS forwarding tables, inconsistencies in the MPLS control and data plane, and problems with the reply path.

The MPLS LSP Ping, Traceroute, and AToM VCCV feature uses MPLS echo request and reply packets to test LSPs. The Cisco implementation of MPLS echo request and echo reply are based on the Internet Engineering Task Force (IETF) Internet-Draft *Detecting MPLS Data Plane Failures.*

Prerequisites for MPLS LSP Ping, Traceroute, and AToM VCCV

Before you use the MPLS LSP Ping, Traceroute, and AToM VCCV feature, you should:

- Determine the baseline behavior of your Multiprotocol Label Switching (MPLS) network. For example:
  - What is the expected MPLS experimental (EXP) treatment?
  - What is the expected maximum size packet or maximum transmission unit (MTU) of the label switched path?

- What is the topology? What are the expected label switched paths? How many links in the label switching path (LSP)? Trace the paths of the label switched packets including the paths for load balancing.

- Understand how to use MPLS and MPLS applications, including traffic engineering, Any Transport over MPLS (AToM), and Label Distribution Protocol (LDP). You need to
  - Know how LDP is configured
  - Understand AToM concepts
- Understand label switching, forwarding, and load balancing.

Restrictions for MPLS LSP Ping, Traceroute, and AToM VCCV

- You cannot use MPLS LSP Traceroute to trace the path taken by Any Transport over Multiprotocol Label Switching (AToM) packets. MPLS LSP Traceroute is not supported for AToM. (MPLS LSP Ping is supported for AToM.) However, you can use MPLS LSP Traceroute to troubleshoot the Interior Gateway Protocol (IGP) LSP that is used by AToM.
- You cannot use MPLS LSP Ping or Traceroute to validate or trace MPLS Virtual Private Networks (VPNs).
- You cannot use MPLS LSP Traceroute to troubleshoot label switching paths (LSPs) that employ time-to-live (TTL) hiding.

Information About MPLS LSP Ping, Traceroute, and AToM VCCV

MPLS LSP Ping Operation

MPLS LSP Ping uses Multiprotocol Label Switching (MPLS) echo request and reply packets to validate a label switched path (LSP). Both an MPLS echo request and an MPLS echo reply are User Datagram Protocol (UDP) packets with source and destination ports set to 3503.

The MPLS echo request packet is sent to a target device through the use of the appropriate label stack associated with the LSP to be validated. Use of the label stack causes the packet to be switched inband of the LSP (that is, forwarded over the LSP itself). The destination IP address of the MPLS echo request packet is different from the address used to select the label stack. The destination address of the UDP packet is defined as a 127.*x* .*y* .*z* /8 address. This prevents the IP packet from being IP switched to its destination if the LSP is broken.

An MPLS echo reply is sent in response to an MPLS echo request. It is sent as an IP packet and forwarded using IP, MPLS, or a combination of both types of switching. The source address of the MPLS echo reply packet is an address from the device generating the echo reply. The destination address is the source address of the device in the MPLS echo request packet.

The figure below shows the echo request and echo reply paths for MPLS LSP Ping.

Figure 1: MPLS LSP Ping Echo Request and Echo Reply Paths

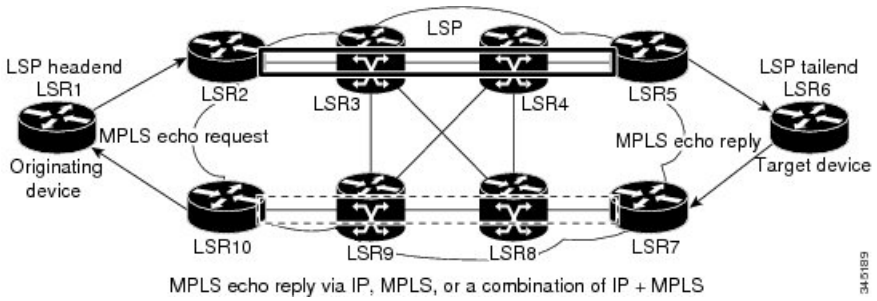

If you initiate an MPLS LSP Ping request at LSR1 to a Forwarding Equivalence Class (FEC), at LSR6, you get the results shown in the table below .

Table 1: MPLS LSP Ping Example

Step	Device	Action
1.	LSR1	Initiates an MPLS LSP Ping request for an FEC at the target device LSR6 and sends an MPLS echo request to LSR2.
1.	LSR2	Receives and forwards the MPLS echo request packet through transit devices LSR3 and LSR4 to the penultimate device LSR5.
1.	LSR5	Receives the MPLS echo request, pops the MPLS label, and forwards the packet to LSR6 as an IP packet.
1.	LSR6	Receives the IP packet, processes the MPLS echo request, and sends an MPLS echo reply to LSR1 through an alternate route.
1.	LSR7 to LSR10	Receive and forward the MPLS echo reply back toward LSR1, the originating device.
1.	LSR1	Receives the MPLS echo reply in response to the MPLS echo request.

You can use MPLS LSP Ping to validate IPv4 Label Distribution Protocol (LDP), Any Transport over MPLS (AToM), and IPv4 Resource Reservation Protocol (RSVP) FECs by using appropriate keywords and arguments with the command:

MPLS LSP Traceroute Operation

MPLS LSP Traceroute also uses Multiprotocol Label Switching (MPLS) echo request and reply packets to validate a label switched path (LSP). The echo request and echo reply are User Datagram Protocol (UDP) packets with source and destination ports set to 3503.

The MPLS LSP Traceroute feature uses time-to-live (TTL) settings to force expiration of the TTL along an LSP. MPLS LSP Traceroute incrementally increases the TTL value in its MPLS echo requests (TTL = 1, 2, 3, 4, ...) to discover the downstream mapping of each successive hop. The success of the LSP traceroute depends on the transit device processing the MPLS echo request when it receives a labeled packet with a TTL of 1. On Cisco devices, when the TTL expires, the packet is sent to the Route Processor (RP) for processing. The transit device returns an MPLS echo reply containing information about the transit hop in response to the TTL-expired MPLS packet.

The figure below shows an MPLS LSP Traceroute example with an LSP from LSR1 to LSR4.

Figure 2: MPLS LSP Traceroute Example

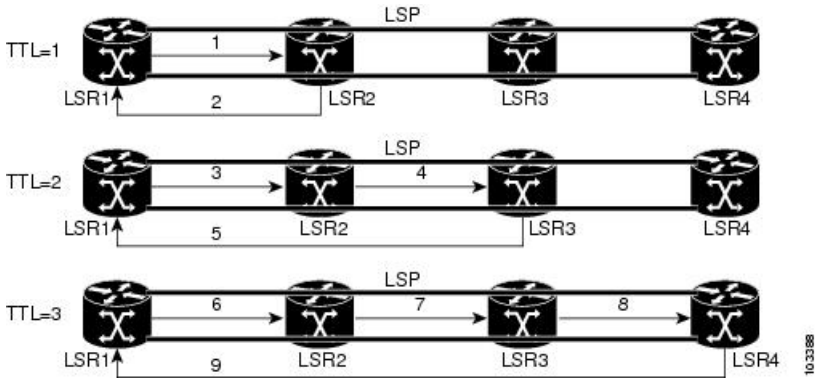

If you enter an LSP traceroute to a Forwarding Equivalence Class (FEC) at LSR4 from LSR1, you get the results shown in the table below.

Table 2: MPLS LSP Traceroute Example

Step	Device	MPLS Packet Type and Description	Device Action
1.	LSR1	MPLS echo request—With a target FEC pointing to LSR4 and to a downstream mapping.	• Sets the TTL of the label stack to 1. • Sends the request to LSR2.
1.	LSR2	MPLS echo reply.	Receives packet with TTL = 1. • Processes the UDP packet as an MPLS echo request. • Finds a downstream mapping, replies to LSR1 with its own downstream mapping based on the incoming label, and sends a reply.
1.	LSR1	MPLS echo request—With the same target FEC and the downstream mapping received in the echo reply from LSR2.	• Sets the TTL of the label stack to 2. • Sends the request to LSR2.
1.	LSR2	MPLS echo request.	Receives packet with TTL = 2. • Decrements the TTL. • Forwards the echo request to LSR3.
1.	LSR3	MPLS reply packet.	Receives packet with TTL = 1. • Processes the UDP packet as an MPLS echo request. • Finds a downstream mapping and replies to LSR1 with its own downstream mapping based on the incoming label.
1.	LSR1	MPLS echo request—With the same target FEC and the downstream mapping received in the echo reply from LSR3.	• Sets the TTL of the packet to 3. • Sends the request to LSR2.

Step	Device	MPLS Packet Type and Description	Device Action
1.	LSR2	MPLS echo request.	Receives packet with TTL = 3. • Decrements the TTL. • Forwards the echo request to LSR3.
1.	LSR3	MPLS echo request.	Receives packet with TTL = 2 • Decrements the TTL. • Forwards the echo request to LSR4.
1.	LSR4	MPLS echo reply.	Receives packet with TTL = 1. • Processes the UDP packet as an MPLS echo request. • Finds a downstream mapping and also finds that the device is the egress device for the target FEC. • Replies to LSR1.

You can use MPLS LSP Traceroute to validate IPv4 Label Distribution Protocol (LDP) and IPv4 RSVP FECs by using appropriate keywords and arguments with the **trace mpls** command:

By default, the TTL is set to 30. Therefore, the traceroute output always contains 30 lines, even if an LSP problem exists. This might mean duplicate entries in the output, should an LSP problem occur. The device address of the last point that the trace reaches is repeated until the output is 30 lines. You can ignore the duplicate entries. The following example shows that the trace encountered an LSP problem at the device that has an IP address of 10.6.1.6:

```
Device# traceroute mpls ipv4 10.6.7.4/32
Tracing MPLS Label Switched Path to 10.6.7.4/32, timeout is 2 seconds
Codes: '!' - success, 'Q' - request not transmitted,
       '.' - timeout, 'U' - unreachable,
       'R' - downstream router but not target
Type escape sequence to abort.
  0 10.6.1.14 MRU 4470 [Labels: 22 Exp: 0]
R 1 10.6.1.5 MRU 4470 [Labels: 21 Exp: 0] 2 ms
R 2 10.6.1.6 4 ms                  <------ Router address repeated for 2nd to 30th TTL.
R 3 10.6.1.6 1 ms
R 4 10.6.1.6 1 ms
R 5 10.6.1.6 3 ms
R 6 10.6.1.6 4 ms
R 7 10.6.1.6 1 ms
R 8 10.6.1.6 2 ms
R 9 10.6.1.6 3 ms
R 10 10.6.1.6 4 ms
R 11 10.6.1.6 1 ms
R 12 10.6.1.6 2 ms
R 13 10.6.1.6 4 ms
R 14 10.6.1.6 5 ms
R 15 10.6.1.6 2 ms
R 16 10.6.1.6 3 ms
R 17 10.6.1.6 4 ms
R 18 10.6.1.6 2 ms
R 19 10.6.1.6 3 ms
```

```
R 20 10.6.1.6 4 ms
R 21 10.6.1.6 1 ms
R 22 10.6.1.6 2 ms
R 23 10.6.1.6 3 ms
R 24 10.6.1.6 4 ms
R 25 10.6.1.6 1 ms
R 26 10.6.1.6 3 ms
R 27 10.6.1.6 4 ms
R 28 10.6.1.6 1 ms
R 29 10.6.1.6 2 ms
R 30 10.6.1.6 3 ms                    <------ TTL 30.
```

If you know the maximum number of hops in your network, you can set the TTL to a smaller value with the **trace mpls ttl** *maximum-time-to-live* command. The following example shows the same **traceroute** command as the previous example, except that this time the TTL is set to 5.

```
Device# traceroute mpls ipv4 10.6.7.4/32 ttl 5
Tracing MPLS Label Switched Path to 10.6.7.4/32, timeout is 2 seconds
Codes: '!' - success, 'Q' - request not transmitted,
       '.' - timeout, 'U' - unreachable,
       'R' - downstream router but not target
Type escape sequence to abort.
  0 10.6.1.14 MRU 4470 [Labels: 22 Exp: 0]
R 1 10.6.1.5 MRU 4474 [No Label] 3 ms
R 2 10.6.1.6 4 ms                  <------ Router address repeated for 2nd to 5th TTL.
R 3 10.6.1.6 1 ms
R 4 10.6.1.6 3 ms
R 5 10.6.1.6 4 ms
```

Any Transport over MPLS Virtual Circuit Connection Verification

AToM Virtual Circuit Connection Verification (AToM VCCV) allows the sending of control packets inband of an AToM pseudowire (PW) from the originating provider edge (PE) device. The transmission is intercepted at the destination PE device, instead of being forwarded to the customer edge (CE) device. This capability allows you to use MPLS LSP Ping to test the PW section of AToM virtual circuits (VCs).

AToM VCCV consists of the following:

- A signaled component in which the AToM VCCV capabilities are advertised during VC label signaling
- A switching component that causes the AToM VC payload to be treated as a control packet

AToM VCCV Signaling

One of the steps involved in Any Transport over Multiprotocol Label Switching (AToM) virtual circuit (VC) setup is the signaling of VC labels and AToM Virtual Circuit Connection Verification (VCCV) capabilities between AToM VC endpoints. The device uses an optional parameter, defined in the Internet Draft *draft-ieft-pwe3-vccv-01.txt,* to communicate the AToM VCCV disposition capabilities of each endpoint.

The AToM VCCV disposition capabilities are categorized as follows:

- Applications—MPLS LSP Ping and Internet Control Message Protocol (ICMP) Ping are applications that AToM VCCV supports to send packets inband of an AToM PW for control purposes.
- Switching modes—Type 1 and Type 2 are switching modes that AToM VCCV uses for differentiating between control and data traffic.

The table below describes AToM VCCV Type 1 and Type 2 switching modes.

Table 3: Type 1 and Type 2 AToM VCCV Switching Modes

Switching Mode	Description
Type 1	Uses a Protocol ID (PID) field in the AToM control word to identify an AToM VCCV packet.
Type 2	Uses an MPLS Router Alert Label above the VC label to identify an AToM VCCV packet.

Selection of AToM VCCV Switching Types

Cisco devices always use Type 1 switching, if available, when they send MPLS LSP Ping packets over an Any Transport over Multiprotocol Label Switching (AToM) virtual circuit (VC) control channel. Type 2 switching accommodates those VC types and implementations that do not support or interpret the AToM control word.

The table below shows the AToM Virtual Circuit Connection Verification (VCCV) switching mode advertised and the switching mode selected by the AToM VC.

Table 4: AToM VCCV Switching Mode Advertised and Selected by AToM Virtual Circuit

Type Advertised	Type Selected
AToM VCCV not supported	–
Type 1 AToM VCCV switching	Type 1 AToM VCCV switching
Type 2 AToM VCCV switching	Type 2 AToM VCCV switching
Type 1 and Type 2 AToM VCCV switching	Type 1 AToM VCCV switching

An AToM VC advertises its AToM VCCV disposition capabilities in both directions: that is, from the originating device (PE1) to the destination device (PE2), and from PE2 to PE1.

In some instances, AToM VCs might use different switching types if the two endpoints have different AToM VCCV capabilities. If PE1 supports Type 1 and Type 2 AToM VCCV switching and PE2 supports only Type 2 AToM VCCV switching, there are two consequences:

- LSP ping packets sent from PE1 to PE2 are encapsulated with Type 2 switching.
- LSP ping packets sent from PE2 to PE1 use Type 1 switching.

You can determine the AToM VCCV capabilities advertised to and received from the peer by entering the **show mpls l2transport binding** command at the PE device. For example:

```
Device# show mpls l2transport binding

  Destination Address: 10.131.191.252,  VC ID: 333
    Local Label:  16
        Cbit: 1,    VC Type: FastEthernet,    GroupID: 0
        MTU: 1500,   Interface Desc: n/a
        VCCV Capabilities: Type 1, Type 2
    Remote Label: 19
        Cbit: 1,    VC Type: FastEthernet,    GroupID: 0
        MTU: 1500,   Interface Desc: n/a
        VCCV Capabilities: Type 1
```

Command Options for ping mpls and trace mpls

MPLS LSP Ping and Traceroute command options are specified as keywords and arguments on the **ping mpls** and **trace mpls** commands.

The **ping mpls** command provides the options displayed in the command syntax below:

The **trace mpls** command provides the options displayed in the command syntax below:

Selection of FECs for Validation

A label switched path (LSP) is formed by labels. Devices learn labels through the Label Distribution Protocol (LDP), traffic engineering (TE), Any Transport over Multiprotocol Label Switching (AToM), or other MPLS applications. You can use MPLS LSP Ping and Traceroute to validate an LSP used for forwarding traffic for a given Forwarding Equivalence Class (FEC). The table below lists the keywords and arguments for the **ping mpls** and **traceroute mpls** commands that allow the selection of an LSP for validation.

Table 5: Selection of LSPs for Validation

FEC Type	ping mpls Keyword and Argument	traceroute mpls Keyword and Argument
LDP IPv4 prefix	**ipv4** *destination-address destination-mask*	**ipv4** *destination-address destination-mask*
MPLS TE tunnel	**traffic-eng** *tunnel-interface tunnel-number*	**traffic-eng** *tunnel-interface tunnel-number*
AToM VC	**pseudowire** *ipv4-address* **vc-id** *vc-id*	MPLS LSP Traceroute does not support the AToM tunnel LSP type for this release.

Reply Mode Options for MPLS LSP Ping and Traceroute

The reply mode is used to control how the responding device replies to a Multiprotocol Label Switching (MPLS) echo request sent by an MPLS LSP Ping or MPLS LSP Traceroute command. The table below describes the reply mode options.

Table 6: Reply Mode Options for a Responding Device

Option	Description
ipv4	Reply with an IPv4 User Datagram Protocol (UDP) packet (default). This is the most common reply mode selected for use with an MPLS LSP Ping and Traceroute command when you want to periodically poll the integrity of a label switched path (LSP). With this option, you do not have explicit control over whether the packet traverses IP or MPLS hops to reach the originator of the MPLS echo request. If the headend device fails to receive a reply, select the router-alert option, “Reply with an IPv4 UDP packet with a router alert.” The responding device sets the IP precedence of the reply packet to 6. You implement this option using the **reply mode ipv4** keywords.

Option	Description
router-alert	Reply with an IPv4 UDP packet with a device alert. This reply mode adds the router alert option to the IP header. This forces the packet to be special handled by the Cisco device at each intermediate hop as it moves back to the destination. This reply mode is more expensive, so use the router-alert option only if you are unable to get a reply with the ipv4 option, "Reply with an IPv4 UDP packet." You implement this option using the **reply mode router-alert** keywords

The reply with an IPv4 UDP packet implies that the device should send an IPv4 UDP packet in reply to an MPLS echo request. If you select the ipv4 reply mode, you do not have explicit control over whether the packet uses IP or MPLS hops to reach the originator of the MPLS echo request. This is the mode that you would normally use to test and verify LSPs.

The reply with an IPv4 UDP packet that contains a device alert forces the packet to go back to the destination and be processed by the Route Processor (RP) process switching at each intermediate hop. This bypasses hardware/line card forwarding table inconsistencies. You should select this option when the originating (headend) devices fail to receive a reply to the MPLS echo request.

You can instruct the replying device to send an echo reply with the IP router alert option by using one of the following commands:

or

However, the reply with a router alert adds overhead to the process of getting a reply back to the originating device. This method is more expensive to process than a reply without a router alert and should be used only if there are reply failures. That is, the reply with a router alert label should only be used for MPLS LSP Ping or MPLS LSP Traceroute when the originating (headend) device fails to receive a reply to an MPLS echo request.

Reply Mode Options for MPLS LSP Ping and Traceroute

The reply mode is used to control how the responding device replies to a Multiprotocol Label Switching (MPLS) echo request sent by an MPLS LSP Ping or MPLS LSP Traceroute command. The table below describes the reply mode options.

Table 7: Reply Mode Options for a Responding Device

Option	Description
ipv4	Reply with an IPv4 User Datagram Protocol (UDP) packet (default). This is the most common reply mode selected for use with an MPLS LSP Ping and Traceroute command when you want to periodically poll the integrity of a label switched path (LSP). With this option, you do not have explicit control over whether the packet traverses IP or MPLS hops to reach the originator of the MPLS echo request. If the headend device fails to receive a reply, select the router-alert option, "Reply with an IPv4 UDP packet with a router alert." The responding device sets the IP precedence of the reply packet to 6. You implement this option using the **reply mode ipv4** keywords.

Option	Description
router-alert	Reply with an IPv4 UDP packet with a device alert. This reply mode adds the router alert option to the IP header. This forces the packet to be special handled by the Cisco device at each intermediate hop as it moves back to the destination. This reply mode is more expensive, so use the router-alert option only if you are unable to get a reply with the ipv4 option, “Reply with an IPv4 UDP packet.” You implement this option using the **reply mode router-alert** keywords

The reply with an IPv4 UDP packet implies that the device should send an IPv4 UDP packet in reply to an MPLS echo request. If you select the ipv4 reply mode, you do not have explicit control over whether the packet uses IP or MPLS hops to reach the originator of the MPLS echo request. This is the mode that you would normally use to test and verify LSPs.

The reply with an IPv4 UDP packet that contains a device alert forces the packet to go back to the destination and be processed by the Route Processor (RP) process switching at each intermediate hop. This bypasses hardware/line card forwarding table inconsistencies. You should select this option when the originating (headend) devices fail to receive a reply to the MPLS echo request.

You can instruct the replying device to send an echo reply with the IP router alert option by using one of the following commands:

or

However, the reply with a router alert adds overhead to the process of getting a reply back to the originating device. This method is more expensive to process than a reply without a router alert and should be used only if there are reply failures. That is, the reply with a router alert label should only be used for MPLS LSP Ping or MPLS LSP Traceroute when the originating (headend) device fails to receive a reply to an MPLS echo request.

Packet Handling Along Return Path with an IP MPLS Router Alert

When an IP packet that contains an IP router alert option in its IP header or a Multiprotocol Label Switching (MPLS) packet with a router alert label as its outermost label arrives at a device, the device punts (redirects) the packet to the Route Processor (RP) process level for handling. This allows these packets to bypass the forwarding failures in hardware routing tables. The table below describes how IP and MPLS packets with an IP router alert option are handled by the device switching path processes.

Table 8: Switching Path Process Handling of IP and MPLS Router Alert Packets

Incoming Packet	Normal Switching Action	Process Switching Action	Outgoing Packet
IP packet—Router alert option in IP header	A rRouter alert option in the IP header causes the packet to be punted to the process switching path.	Forwards the packet as is.	IP packet—Router alert option in IP header.
	A router alert option in theIP header causes the packet to be punted to the process switching path.	Adds a router alert as the outermost label and forwards as an MPLS packet.	MPLS packet— Outermost label contains a router alert.

Incoming Packet	Normal Switching Action	Process Switching Action	Outgoing Packet
MPLS packet—Outermost label contains a router alert	If the router alert label is the outermost label, the packet is punted to the process switching path.	Removes the outermost router alert label, adds an IP router alert option to the IP header, and forwards as an IP packet.	IP packet—Router alert option in IP header.
	If the router alert label is the outermost label, the packet is punted to the process switching path.	Preserves the outermost router alert label and forwards the MPLS packet.	MPLS packet— Outermost label contains a router alert.

Other MPLS LSP Ping and Traceroute Command Options

The table below describes other MPLS LSP Ping and Traceroute command options that can be specified as keywords or arguments with the **ping mpls** command, or with both the **ping mpls** and **trace mpls** commands. Options available to use only on the **ping mpls** command are indicated as such.

Table 9: Other MPLS LSP Ping and Traceroute and AToM VCCV Options

Option	Description
Datagram size	Size of the packet with the label stack imposed. Specified with the **size** *packet-size* keyword and argument. The default size is 100. For use with the MPLS LSP Ping feature only.
Padding	Padding (the pad time-length-value [TLV]) is used as required to fill the datagram so that the MPLS echo request (User Datagram Protocol [UDP] packet with a label stack) is the size specified. Specify with the **pad** *pattern* keyword and argument. For use with the MPLS LSP Ping feature only.
Sweep size range	Parameter that enables you to send a number of packets of different sizes, ranging from a start size to an end size. This parameter is similar to the Internet Control Message Protocol (ICMP) ping sweep parameter. The lower boundary on the sweep range varies depending on the label switched path (LSP) type. You can specify a sweep size range when you use the **ping mpls** command. Use the **sweep** *minimum maximum size-increment* keyword and arguments. For use with the MPLS LSP Ping feature only.
Repeat count	Number of times to resend the same packet. The default is 5 times. You can specify a repeat count when you use the **ping mpls** command. Use the **repeat** *count* keyword and argument. For use with the MPLS LSP Ping feature only.
MPLS echo request source address	Routable address of the sender. The default address is loopback0. This address is used as the destination address in the Multiprotocol Label Switching (MPLS) echo response. Use the **source** *source-address* keyword and argument. For use with the MPLS LSP Ping and Traceroute features.

Option	Description
UDP destination address	A valid 127/8 address. You have the option to specify a single *x.y.z* or a range of numbers between 0.0.0 and *x.y.z* , where *x.y.z* are numbers between 0 and 255 and correspond to 127.*x.y.z*. Use the **destination** {*address* | *address-start address-end increment*} keyword and arguments. The MPLS echo request destination address in the UDP packet is not used to forward the MPLS packet to the destination device. The label stack that is used to forward the echo request routes the MPLS packet to the destination device. The 127/8 address guarantees that the packets are routed to the localhost (the default loopback address of the device processing the address) if the UDP packet destination address is used for forwarding. In addition, the destination address is used to affect load balancing when the destination address of the IP payload is used for load balancing. For use with IPv4 and Any Transport over MPLS (AToM) Forwarding Equivalence Classes (FECs) with the MPLS LSP Ping feature and with IPv4 FECs with the MPLS LSP Traceroute feature.
Time-to-live (TTL)	A parameter you can set that indicates the maximum number of hops a packet should take to reach its destination. The time-to-live (TTL) field in a packet is decremented by 1 each time it travels through a device. For MPLS LSP Ping, the TTL is a value after which the packet is discarded and an MPLS echo reply is sent back to the originating device. Use the **ttl** *time-to-live* keyword and argument. For MPLS LSP Traceroute, the TTL is a maximum time to live and is used to discover the number of downstream hops to the destination device. MPLS LSP Traceroute incrementally increases the TTL value in its MPLS echo requests (TTL = 1, 2, 3, 4, ...) to accomplish this. Use the **ttl** *time-to-live* keyword and argument.
Timeouts	A parameter you can specify to control the timeout in seconds for an MPLS request packet. The range is from 0 to 3600 seconds. The default is 2. Set with the **timeout** *seconds* keyword and argument. For use with the MPLS LSP Ping and Traceroute features.
Intervals	A parameter you can specify to set the time in milliseconds between successive MPLS echo requests. The default is 0. Set with the **interval** *msec* keyword and argument.
Experimental bits	Three experimental bits in an MPLS header used to specify precedence for the MPLS echo reply. (The bits are commonly called EXP bits.) The range is from 0 to 7, and the default is 0. Specify with the **exp** *exp-bits* keyword and argument. For use with the MPLS LSP Ping and Traceroute features.

Option	Description
Verbose	Option that provides additional information for the MPLS echo reply--source address and return codes. For the MPLS LSP Ping feature, this option is implemented with the **verbose** keyword. For use with the MPLS LSP Ping feature only.

MPLS LSP Ping options described in the table above can be implemented by using the following syntax:

```
ping mpls
{ipv4 destination-address destination-mask [destination address-start address-end increment]

 [ttl time-to-live] | pseudowire ipv4-address
 vc-id vc-id
[destination address-start address-end increment] | traffic-eng tunnel-interface
tunnel-number
[ttl time-to-live]}
[source source-address] [repeat count]
[{size packet-size} | {sweep minimum maximum size-Increment}]
[pad pattern]
[timeout seconds] [intervalmsec]
[exp exp-bits] [verbose
```

MPLS LSP Traceroute options described in the table below can be implemented by the use of the following syntax:

```
trace mpls
 {ipv4 destination-address destination-mask
 [destination address-start address-end address-increment] | traffic-eng tunnel-interface
tunnel-number}
[source source-address] [timeout seconds]
[ttl maximum-time-to-live]
[exp exp-bits]
```

Option Interactions and Loops

Usage examples for the MPLS LSP Ping and Traceroute and AToM VCCV feature in this and subsequent sections are based on the sample topology shown in the figure below.

Figure 3: Sample Topology for Configuration Examples

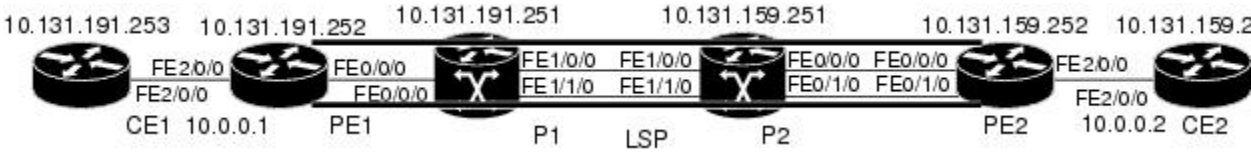

The interaction of some MPLS LSP Ping and Traceroute and AToM VCCV options can cause loops. See the following topic for a description of the loops you might encounter with the **ping mpls** and **trace mpls** commands:

Possible Loops with MPLS LSP Ping

With the MPLS LSP Ping feature, loops can occur if you use the repeat count option, the sweep size range option, or the User Datagram Protocol (UDP) destination address range option.

```
ping mpls
 {ipv4 destination-address/destination-mask
[destination address-start address-end increment] | pseudowire ipv4-address
 vc-id vc-id
 [destination address-start address-end increment] |
traffic-eng tunnel-interface tunnel-number}
[repeat count]
[sweep minimum maximum size-increment]
```

Following is an example of how a loop operates if you use the following keywords and arguments on the **ping mpls** command:

```
Device# ping mpls
 ipv4
 10.131.159.251/32 destination 127.0.0.1 127.0.0.1 0.0.0.1 repeat 2
sweep 1450 1475 25
Sending 2, [1450..1500]-byte MPLS Echos to 10.131.159.251/32,
     timeout is 2 seconds, send interval is 0 msec:
Codes: '!' - success, 'Q' - request not transmitted,
       '.' - timeout, 'U' - unreachable,
       'R' - downstream router but not target
Type escape sequence to abort.
Destination address 127.0.0.1
!
!
Destination address 127.0.0.1
!
!
Destination address 127.0.0.1
!
!
Destination address 127.0.0.1
!
!
```

An **mpls ping** command is sent for each packet size range for each destination address until the end address is reached. For this example, the loop continues in the same manner until the destination address, 127.0.0.1, is reached. The sequence continues until the number is reached that you specified with the **repeat** *count* keyword and argument. For this example, the repeat count is 2. The MPLS LSP Ping loop sequence is as follows:

```
repeat  = 1
  destination address 1 (address-start
)
    for (size from sweep
 minimum
 to maximum
, counting by size-increment
)
      send an lsp ping
  destination address 2 (address-start
 +
address-
increment
)
    for (size from sweep
 minimum
 to maximum
, counting by size-increment
)
      send an lsp ping
```

```
  destination address 3 (address-start
 +
address-
increment
 +
address-
increment
)
     for (size from sweep
 minimum
 to maximum
, counting by size-increment
)
      send an lsp ping
.
.
.
  until destination address = address-end
.
.
.
until repeat = count
```

Possible Loop with MPLS LSP Traceroute

With the MPLS LSP Traceroute feature, loops can occur if you use the User Datagram Protocol (UDP) destination address range option and the time-to-live option.

Here is an example of how a loop operates if you use the following keywords and arguments on the **trace mpls** command:

```
Device# trace mpls
ipv4
 10.131.159.251/32 destination 127.0.0.1 127.0.0.1 1 ttl 5
Tracing MPLS Label Switched Path to 10.131.159.251/32, timeout is 2 seconds
Codes: '!' - success, 'Q' - request not transmitted,
       '.' - timeout, 'U' - unreachable,
       'R' - downstream router but not target
Type escape sequence to abort.
Destination address 127.0.0.1
  0 10.131.191.230 MRU 1500 [Labels: 19 Exp: 0]
R 1 10.131.159.226 MRU 1504 [implicit-null] 40 ms
! 2 10.131.159.225 40 ms
Destination address 127.0.0.2
  0 10.131.191.230 MRU 1500 [Labels: 19 Exp: 0]
R 1 10.131.159.226 MRU 1504 [implicit-null] 40 ms
! 2 10.131.159.225 40 ms
Destination address 127.0.0.3
  0 10.131.191.230 MRU 1500 [Labels: 19 Exp: 0]
R 1 10.131.159.226 MRU 1504 [implicit-null] 40 ms
! 2 10.131.159.225 48 ms
```

An **mpls trace** command is sent for each TTL from 1 to the maximum TTL (**ttl** *maximum-time-to-live* keyword and argument) for each destination address until the address specified with the destination *end-address* argument is reached. For this example, the maximum TTL is 5 and the end destination address is 127.0.0.1. The MPLS LSP Traceroute loop sequence is as follows:

```
destination address 1 (address-start
)
  for (ttl
 from 1 to maximum-time-to-live
```

```
)
    send an lsp trace
destination address 2 (address-start
 + address-increment
)
  for (ttl
 from 1 to maximum-time-to-live
)
    send an lsp trace
destination address 3 (address-start
 + address-increment
 + address-increment
)
  for (ttl
 from 1 to
maximum-time-to-live)
    send an lsp trace
.
.
.
until destination address = address-end
```

MPLS Echo Request Packets Not Forwarded by IP

Multiprotocol Label Switching (MPLS) echo request packets sent during a label switched path (LSP) ping are never forwarded by IP. The IP header destination address field in an MPLS echo request packet is a 127.*x.y.z* /8 address. Devices should not forward packets using a 127.*x.y.z* /8 address. The 127.*x.y.z* /8 address corresponds to an address for the local host.

The use of a 127.*x .y .z* address as a destination address of the User Datagram Protocol (UDP) packet is significant in that the MPLS echo request packet fails to make it to the target device if a transit device does not label switch the LSP. This allows for the detection of LSP breakages.

- If an LSP breakage occurs at a transit device, the MPLS echo packet is not forwarded, but consumed by the device.
- If the LSP is intact, the MPLS echo packet reaches the target device and is processed by the terminal point of the LSP.

The figure below shows the path of the MPLS echo request and reply when a transit device fails to label switch a packet in an LSP.

Figure 4: Path When Transit Device Fails to Label Switch a Packet

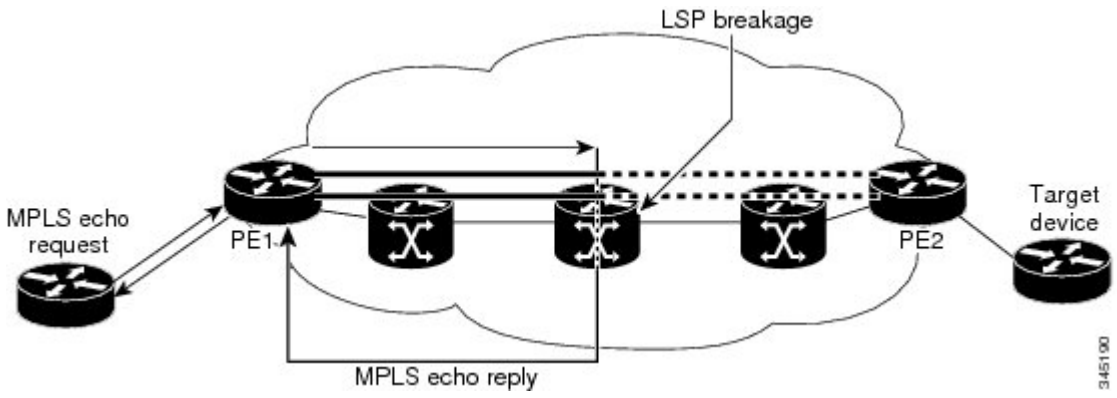

Note An Any Transport over MPLS (AToM) payload does not contain usable forwarding information at a transit device because the payload might not be an IP packet. An MPLS virtual private network (VPN) packet, although an IP packet, does not contain usable forwarding information at a transit device because the destination IP address is only significant to the virtual routing and forwarding (VRF) instances at the endpoints of the MPLS network.

Information Provided by the Device Processing LSP Ping or LSP Traceroute

The table below describes the characters that the device processing an LSP ping or LSP traceroute packet returns to the sender about the failure or success of the request.

You can also view the return code for an MPLS LSP Ping operation if you enter the **ping mpls verbose** command.

Table 10: LSP Ping and Traceroute Reply Characters

Character	Meaning
Period "."	A timeout occurs before the target device can reply.
U	The target device is unreachable.
R	The device processing the Multiprotocol Label Switching (MPLS) echo request is a downstream device but is not the destination.
Exclamation mark "!"	Replying device is an egress for the destination.
Q	Echo request was not successfully transmitted. This could be returned because of insufficient memory or more probably because no label switched path (LSP) exists that matches the Forwarding Equivalence Class (FEC) information.
C	Replying device rejected the echo request because it was malformed.

MTU Discovery in an LSP

During an MPLS LSP Ping, Multiprotocol Label Switching (MPLS) echo request packets are sent with the IP packet attribute set to do not fragment. That is, the DF bit is set in the IP header of the packet. This allows you to use the MPLS echo request to test for the MTU that can be supported for the packet through the label switched path (LSP) without fragmentation.

The figure below shows a sample network with a single LSP from PE1 to PE2 formed with labels advertised by means of LDP.

Figure 5: Sample Network with LSP—Labels Advertised by LDP

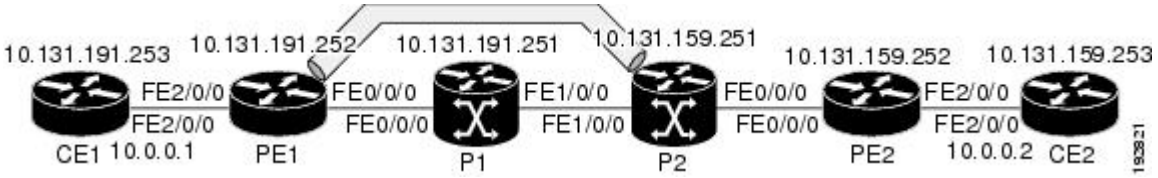

You can determine the maximum receive unit (MRU) at each hop by tracing the LSP using the MPLS Traceroute feature. The MRU is the maximum size of a labeled packet that can be forwarded through an LSP. The following example shows the results of a **trace mpls** command when the LSP is formed with labels created by the Label Distribution Protocol (LDP):

```
Device# trace mpls ipv4 10.131.159.252/32
Tracing MPLS Label Switched Path to 10.131.159.252/32, timeout is 2 seconds
Codes: '!' - success, 'Q' - request not transmitted,
       '.' - timeout, 'U' - unreachable,
       'R' - downstream router but not target
Type escape sequence to abort.
  0 10.131.191.230 MRU 1496 [Labels: 22/19 Exp: 0/0]
R 1 10.131.159.226 MRU 1500 [Labels: 19 Exp: 0] 40 ms
R 2 10.131.159.229 MRU 1504 [implicit-null] 28 ms
! 3 10.131.159.230 40 ms
```

You can determine the MRU for the LSP at each hop through the use of the **show forwarding detail** command:

```
Device# show mpls forwarding 10.131.159.252 detail

Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
22     19          10.131.159.252/32 0          Tu1        point2point
        MAC/Encaps=14/22, MRU=1496, Tag Stack{22 19}, via Et0/0
        AABBCC009700AABBCC0098008847 0001600000013000
        No output feature configured
```

To determine the maximum sized echo request that will fit on the LSP, you can find the IP MTU by using the **show interface** *type number* command.

```
Device# show interface e0/0

FastEthernet0/0/0 is up, line protocol is up
  Hardware is Lance, address is aabb.cc00.9800 (bia aabb.cc00.9800)
  Internet address is 10.131.191.230/30
  MTU 1500 bytes, BW 10000 Kbit, DLY 1000 usec, rely 255/255, load ½55
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:01, output 00:00:01, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     377795 packets input, 33969220 bytes, 0 no buffer
     Received 231137 broadcasts, 0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
     0 input packets with dribble condition detected
     441772 packets output, 40401350 bytes, 0 underruns
```

```
     0 output errors, 0 collisions, 10 interface resets
     0 babbles, 0 late collision, 0 deferred
     0 lost carrier, 0 no carrier
     0 output buffer failures, 0 output buffers swapped out
```

The IP MTU in the **show interface** *type number* example is 1500 bytes. Subtract the number of bytes corresponding to the label stack from the MTU number. From the output of the **show mpls forwarding** command, the Tag stack consists of one label (21). Therefore, the largest MPLS echo request packet that can be sent in the LSP, shown in the figure above, is 1500 - (2 x 4) = 1492.

You can validate this by using the following **ping mpls** command:

```
Device# ping mpls ipv4 10.131.159.252/32 sweep 1492 1500 1 repeat 1
Sending 1, [1492..1500]-byte MPLS Echos to 10.131.159.252/32,
     timeout is 2 seconds, send interval is 0 msec:
Codes: '!' - success, 'Q' - request not transmitted,
       '.' - timeout, 'U' - unreachable,
       'R' - downstream router but not target
Type escape sequence to abort.
!QQQQQQQQ
Success rate is 11 percent (1/9), round-trip min/avg/max = 40/40/40 ms
```

In this command, only packets of 1492 bytes are sent successfully, as indicated by the exclamation point (!). Packets of byte sizes 1493 to 1500 are source-quenched, as indicated by the Q.

You can pad an MPLS echo request so that a payload of a given size can be tested. The pad TLV is useful when you use the MPLS echo request to discover the MTU supportable by an LSP. MTU discovery is extremely important for applications like AToM that contain non-IP payloads that cannot be fragmented.

LSP Network Management

To manage a Multiprotocol Label Switching (MPLS) network you must have the ability to monitor label switched paths (LSPs) and quickly isolate MPLS forwarding problems. You need ways to characterize the liveliness of an LSP and reliably detect when a label switched path fails to deliver user traffic.

You can use MPLS LSP Ping to verify the LSP that is used to transport packets destined for IPv4 Label Distribution Protocol (LDP) prefixes, traffic engineering (TE) tunnels, and Any Transport over MPLS pseudowire Forwarding Equivalence Classes (AToM PW FECs). You can use MPLS LSP Traceroute to trace LSPs that are used to carry packets destined for IPv4 LDP prefixes and TE tunnel FECs.

An MPLS echo request is sent through an LSP to validate it. A TTL expiration or LSP breakage causes the transit device to process the echo request before it gets to the intended destination and returns an MPLS echo reply that contains an explanatory reply code to the originator of the echo request.

The successful echo request is processed at the egress of the LSP. The echo reply is sent via an IP path, an MPLS path, or a combination of both back to the originator of the echo request.

ICMP ping and trace Commands and Troubleshooting

Internet Control Message Protocol (ICMP) **ping** and **trace** commands are often used to help diagnose the root cause of a failure. When a label switched path (LSP) is broken, the packet might make its way to the target device by way of IP forwarding, thus making ICMP ping and traceroute unreliable for detecting Multiprotocol Label Switching (MPLS) forwarding problems. The MPLS LSP Ping, Traceroute and AToM VCCV feature extends this diagnostic and troubleshooting ability to the MPLS network and handles inconsistencies between the IP and MPLS forwarding tables, inconsistencies in the MPLS control and data plane, and problems with the reply path.

The figure below shows a sample topology with a Label Distribution Protocol (LDP) LSP and traffic engineering (TE) tunnel LSP.

Figure 6: Sample Topology with LDP and TE Tunnel LSPs

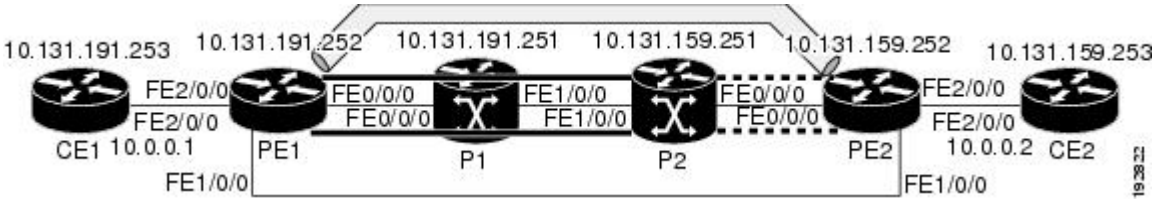

This section contains the following topics:

MPLS LSP Ping and Traceroute Discovers LSP Breakage

Configuration for Sample Topology

These are sample topology configurations for the troubleshooting examples in the following sections (see the figure above). There are the six sample device configurations.

Device CE1 Configuration

```
version 12.0
!
hostname ce1
!
enable password lab
!
interface Loopback0
 ip address 10.131.191.253 255.255.255.255
 no ip directed-broadcast
!
interface
 ip address 10.0.0.1 255.255.255.255
 no ip directed-broadcast
 no keepalive
 no cdp enable
!
end
```

Device PE1 Configuration

```
version 12.0
!
hostname pe1
!
ip cef
mpls label protocol ldp
mpls traffic-eng tunnels
no mpls traffic-eng auto-bw timers frequency 0
mpls ldp discovery targeted-hello accept
!
interface Loopback0
 ip address 10.131.191.252 255.255.255.255
 no ip directed-broadcast
!
interface Tunnel1
```

```
 ip unnumbered Loopback0
 no ip directed-broadcast
 mpls label protocol ldp
 mpls ip
 tunnel destination 10.131.159.255
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng autoroute announce
 tunnel mpls traffic-eng priority 2 2
 tunnel mpls traffic-eng bandwidth  512
 tunnel mpls traffic-eng path-option 1 dynamic
!
interface Tunnel2
 ip unnumbered Loopback0
 no ip directed-broadcast
 shutdown
 mpls label protocol ldp
 mpls ip
 tunnel destination 10.131.159.255
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng autoroute announce
 tunnel mpls traffic-eng priority 1 1
 tunnel mpls traffic-eng bandwidth  100
 tunnel mpls traffic-eng path-option 1 dynamic
!
interface
 ip address 10.131.191.230 255.255.255.255
 no ip directed-broadcast
 mpls traffic-eng tunnels
 mpls ip
 ip rsvp bandwidth 1500 1500
 ip rsvp signalling dscp 0
!
interface
 ip address 10.131.159.246 255.255.255.255
 no ip directed-broadcast
 no shutdown
 mpls ip
 ip rsvp bandwidth 1500 1500
 ip rsvp signalling dscp 0
!
interface
 no ip address
 no ip directed-broadcast
 no cdp enable
 xconnect 10.131.159.252 333 encapsulation mpls
!
interface
 no ip address
 no ip directed-broadcast
 shutdown
!
router ospf 1
 log-adjacency-changes
 passive-interface Loopback0
 network 10.131.159.244 0.0.0.3 area 0
 network 10.131.191.228 0.0.0.3 area 0
 network 10.131.191.232 0.0.0.3 area 0
 network 10.131.191.252 0.0.0.0 area 0
 mpls traffic-eng router-id Loopback0
 mpls traffic-eng area 0
!
ip classless

end
```

Device P1 Configuration

```
version 12.0
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname p1
!
enable password lab
!
ip cef
mpls label protocol ldp
mpls ldp logging neighbor-changes
mpls traffic-eng tunnels
no mpls traffic-eng auto-bw timers frequency 0
mpls ldp discovery targeted-hello accept
!
interface Loopback0
 ip address 10.131.191.251 255.255.255.255
 no ip directed-broadcast
!
interface
 ip address 10.131.191.229 255.255.255.255
 no ip directed-broadcast
 mpls traffic-eng tunnels
 mpls ip
 ip rsvp bandwidth 1500 1500
 ip rsvp signalling dscp 0
!
interface
 ip address 10.131.159.226 255.255.255.255
 no ip directed-broadcast
 mpls traffic-eng tunnels
 mpls ip
 ip rsvp bandwidth 1500 1500
 ip rsvp signalling dscp 0
!
router ospf 1
 log-adjacency-changes
 passive-interface Loopback0
 network 10.131.159.224 0.0.0.3 area 0
 network 10.131.191.228 0.0.0.3 area 0
 network 10.131.191.251 0.0.0.0 area 0
 mpls traffic-eng router-id Loopback0
 mpls traffic-eng area 0
!
end
```

Device P2 Configuration

```
version 12.0
hostname p2
!
ip cef
mpls label protocol ldp
mpls ldp logging neighbor-changes
mpls traffic-eng tunnels
no mpls traffic-eng auto-bw timers frequency 0
mpls ldp discovery directed-hello accept
!
!
```

```
interface Loopback0
 ip address 10.131.159.251 255.255.255.255
 no ip directed-broadcast
!
interface
 ip address 10.131.159.229 255.255.255.255
 no ip directed-broadcast
 mpls traffic-eng tunnels
 mpls ip
 ip rsvp bandwidth 1500 1500
 ip rsvp signalling dscp 0
!
interface
 ip address 10.131.159.225 255.255.255.255
 no ip directed-broadcast
 mpls traffic-eng tunnels
 mpls ip
 ip rsvp bandwidth 1500 1500
 ip rsvp signalling dscp 0
!
router ospf 1
 log-adjacency-changes
 passive-interface Loopback0
 network 10.131.159.224 0.0.0.3 area 0
 network 10.131.159.228 0.0.0.3 area 0
 network 10.131.159.251 0.0.0.0 area 0
 mpls traffic-eng router-id Loopback0
 mpls traffic-eng area 0
!
end
```

Device PE2 Configuration

```
version 12.0
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname pe2
!
logging snmp-authfail
enable password lab
!
clock timezone EST -5
ip subnet-zero
ip cef
no ip domain-lookup
mpls label protocol ldp
mpls ldp logging neighbor-changes
mpls ldp explicit-null
mpls traffic-eng tunnels
no mpls traffic-eng auto-bw timers frequency 0
tag-switching tdp discovery directed-hello accept
frame-relay switching
!
!
interface Loopback0
 ip address 10.131.159.252 255.255.255.255
 no ip directed-broadcast
!
interface Tunnel0
 ip unnumbered Loopback0
 no ip directed-broadcast
```

```
 tunnel destination 10.131.191.252
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng path-option 5 explicit name as1pe-long-path
!
interface
 ip address 10.131.159.230 255.255.255.255
 no ip directed-broadcast
 mpls traffic-eng tunnels
 tag-switching ip
 ip rsvp bandwidth 1500 1500
 ip rsvp signalling dscp 0
!
interface
 ip address 10.131.159.245 255.255.255.255
 no ip directed-broadcast
 mpls traffic-eng tunnels
 tag-switching ip
 ip rsvp bandwidth 1500 1500
 ip rsvp signalling dscp 0
!
interface
 no ip address
 no ip directed-broadcast
 no cdp enable
 xconnect 10.131.191.252 333 encapsulation mpls
!
interface
 no ip address
 no ip directed-broadcast
!
interface
 no ip address
 no ip directed-broadcast
 shutdown
!
interface
 no ip address
 no ip directed-broadcast
 shutdown
!
router ospf 1
 mpls traffic-eng router-id Loopback0
 mpls traffic-eng area 0
 log-adjacency-changes
 passive-interface Loopback0
 network 10.131.122.0 0.0.0.3 area 0
 network 10.131.159.228 0.0.0.3 area 0
 network 10.131.159.232 0.0.0.3 area 0
 network 10.131.159.244 0.0.0.3 area 0
 network 10.131.159.252 0.0.0.0 area 0
!
ip classless
!
!
ip explicit-path name as1pe-long-path enable
 next-address 10.131.159.229
 next-address 10.131.159.226
 next-address 10.131.191.230
!
!
line con 0
 exec-timeout 0 0
line aux 0
line vty 0 4
```

```
 exec-timeout 0 0
 password lab
 login
!
end
```

Device CE2 Configuration

```
version 12.0
!
hostname ce2
!
enable password lab
!
interface Loopback0
 ip address 10.131.159.253 255.255.255.255
 no ip directed-broadcast
!
interface
 ip address 10.0.0.2 255.255.255.255
 no ip directed-broadcast
 no keepalive
 no cdp enable
!
end
```

Verifying That the LSP Is Set Up Correctly

A **show mpls forwarding-table** command shows that tunnel 1 is in the Multiprotocol Label Switching (MPLS) forwarding table.

```
Device# show mpls forwarding-table 10.131.159.252

Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
22     19
     [T] 10.131.159.252/32 0          Tu1
        point2point
[T]     Forwarding through a TSP tunnel.
        View additional tagging info with the 'detail' option
```

A **show mpls traffic-eng tunnels tunnel 1** command entered at PE1 displays information about tunnel 1 and verifies that it is forwarding packets with an out label of 22.

```
Device# show mpls traffic-eng tunnels tunnel 1

Name: PE1_t1                           (Tunnel1) Destination: 10.131.159.251
  Status:
    Admin: up         Oper: up     Path: valid       Signalling: connected
    path option 1, type dynamic (Basis for Setup, path weight 20)
  Config Parameters:
    Bandwidth: 512      kbps (Global)  Priority: 2  2   Affinity: 0x0/0xFFFF
    Metric Type: TE (default)
    AutoRoute:  enabled   LockDown: disabled  Loadshare: 512      bw-based
    auto-bw: disabled
  Active Path Option Parameters:
    State: dynamic path option 1 is active
    BandwidthOverride: disabled  LockDown: disabled  Verbatim: disabled
  InLabel  :  -
  OutLabel : FastEthernet0/0/0, 22
  RSVP Signalling Info:
```

```
      Src 10.131.191.252, Dst 10.131.159.251, Tun_Id 1, Tun_Instance 28
    RSVP Path Info:
      My Address: 10.131.191.230
      Explicit Route: 10.131.191.229 10.131.159.226 10.131.159.225 10.131.159.251
      Record   Route:   NONE
      Tspec: ave rate=512 kbits, burst=1000 bytes, peak rate=512 kbits
    RSVP Resv Info:
      Record   Route:   NONE
      Fspec: ave rate=512 kbits, burst=1000 bytes, peak rate=512 kbits
  Shortest Unconstrained Path Info:
    Path Weight: 20 (TE)
    Explicit Route: 10.131.191.230 10.131.191.229 10.131.159.226 10.131.159.225
                    10.131.159.251
  History:
    Tunnel:
      Time since created: 9 days, 14 hours, 12 minutes
      Time since path change: 2 minutes, 18 seconds
    Current LSP:
      Uptime: 2 minutes, 18 seconds
    Prior LSP:
      ID: path option 1 [3]
      Removal Trigger: tunnel shutdown
```

A **trace mpls** command issued at PE1 verifies that packets with 22 as the outermost label and 19 as the end of stack label are forwarded from PE1 to PE2.

```
Device# trace mpls ipv4 10.131.159.252/32
Tracing MPLS Label Switched Path to 10.131.159.252/32, timeout is 2 seconds
Codes: '!' - success, 'Q' - request not transmitted,
       '.' - timeout, 'U' - unreachable,
       'R' - downstream router but not target
Type escape sequence to abort.
  0 10.131.191.230 MRU 1496 [Labels: 22/19
 Exp: 0/0]
R 1 10.131.159.226 MRU 1504 [Labels: 19 Exp: 0] 40 ms
R 2 10.131.159.229 MRU 1504 [implicit-null] 28 ms
! 3 10.131.159.230 40 ms
```

The MPLS LSP Traceroute to PE2 is successful, as indicated by the exclamation point (!).

Discovering LSP Breakage

A Label Distribution Protocol (LDP) target-session is established between devices PE1 and P2, as shown in the output of the following **show mpls ldp discovery** command:

```
Device# show mpls ldp discovery

 Local LDP Identifier:
    10.131.191.252:0
    Discovery Sources:
    Interfaces:
         (ldp): xmit/recv
            LDP Id: 10.131.191.251:0
        Tunnel1 (ldp): Targeted -> 10.131.159.251
    Targeted Hellos:
        10.131.191.252 -> 10.131.159.252 (ldp): active/passive, xmit/recv
            LDP Id: 10.131.159.252:0
        10.131.191.252 -> 10.131.159.251 (ldp): active, xmit/recv
            LDP Id: 10.131.159.251:0
```

Enter the following command on the P2 device in global configuration mode:

```
Device# no mpls ldp discovery targeted-hello accept
```

The LDP configuration change causes the targeted LDP session between the headend and tailend of the traffic engineering (TE) tunnel to go down. Labels for IPv4 prefixes learned by P2 are not advertised to PE1. Thus, all IP prefixes reachable by P2 are reachable by PE1 only through IP (not MPLS). In other words, packets destined for those prefixes through Tunnel 1 at PE1 will be IP switched at P2 (which is undesirable).

The following **show mpls ldp discovery** command shows that the LDP targeted-session is down:

```
Device# show mpls ldp discovery

 Local LDP Identifier:
    10.131.191.252:0
    Discovery Sources:
    Interfaces:
        (ldp): xmit/recv
            LDP Id: 10.131.191.251:0
        Tunnel1 (ldp): Targeted -> 10.131.159.251
    Targeted Hellos:
        10.131.191.252 -> 10.131.159.252 (ldp): active/passive, xmit/recv
            LDP Id: 10.131.159.252:0
        10.131.191.252 -> 10.131.159.251 (ldp): active, xmit
```

Enter the **show mpls forwarding-table** command at the PE1 device. The display shows that the outgoing packets are untagged as a result of the LDP configuration changes.

```
Device# show mpls forwarding-table 10.131.159.252

Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
22     Untagged[T]
 10.131.159.252/32 0          Tu1        point2point
[T]     Forwarding through a TSP tunnel.
        View additional tagging info with the 'detail' option
```

A **ping mpls** command entered at the PE1 device displays the following:

```
Device# ping mpls ipv4 10.131.159.252/32 repeat 1
Sending 1, 100-byte MPLS Echos to 10.131.159.252/32,
      timeout is 2 seconds, send interval is 0 msec:
Codes: '!' - success, 'Q' - request not transmitted,
       '.' - timeout, 'U' - unreachable,
       'R' - downstream router but not target
Type escape sequence to abort.
R
Success rate is 0 percent (0/1)
```

The **ping mpls** command fails. The R indicates that the sender of the Multiprotocol Label Switching (MPLS) echo reply had a routing entry but no MPLS Forwarding Equivalence Class (FEC) . Entering the **ping mpls verbose** command displays the MPLS label switched path (LSP) echo reply sender address and the return code. You should be able to solve the problem by Telneting to the replying device and inspecting its forwarding and label tables. You might need to look at the neighboring upstream device as well, because the breakage might be on the upstream device.

```
Device# ping mpls ipv4 10.131.159.252/32 repeat 1 verbose
Sending 1, 100-byte MPLS Echos to 10.131.159.252/32,
      timeout is 2 seconds, send interval is 0 msec:
Codes: '!' - success, 'Q' - request not transmitted,
       '.' - timeout, 'U' - unreachable,
```

```
       'R' - downstream router but not target
Type escape sequence to abort.
R   10.131.159.225, return code 6
Success rate is 0 percent (0/1)
```

Alternatively, use the LSP **traceroute** command to figure out which device caused the breakage. In the following example, for subsequent values of TTL greater than 2, the same device keeps responding (10.131.159.225). This suggests that the MPLS echo request keeps getting processed by the device regardless of the TTL. Inspection of the label stack shows that P1 pops the last label and forwards the packet to P2 as an IP packet. This explains why the packet keeps getting processed by P2. MPLS echo request packets cannot be forwarded by use of the destination address in the IP header because the address is set to a 127/8 address.

```
Device# trace mpls ipv4 10.131.159.252/32 ttl 5
Tracing MPLS Label Switched Path to 10.131.159.252/32, timeout is 2 seconds
Codes: '!' - success, 'Q' - request not transmitted,
       '.' - timeout, 'U' - unreachable,
       'R' - downstream router but not target
Type escape sequence to abort.
  0 10.131.191.230 MRU 1500 [Labels: 22 Exp: 0]
R 1 10.131.159.226 MRU 1504 [implicit-null] 40 ms
R 2 10.131.159.225 40 ms
R 3 10.131.159.225 40 ms
R 4 10.131.159.225 40 ms
R 5 10.131.159.225 40 ms
```

MPLS LSP Traceroute Tracks Untagged Cases

This troubleshooting section contains examples of how to use MPLS LSP Traceroute to determine potential issues with packets that are tagged as implicit null and packets that are untagged.

Untagged output interfaces at a penultimate hop do not impact the forwarding of IP packets through a label switched path (LSP) because the forwarding decision is made at the penultimate hop through use of the incoming label. The untagged case causes Any Transport over Multiprotocol Label Switching (AToM) and MPLS virtual private network (VPN) traffic to be dropped at the penultimate hop.

Troubleshooting Implicit Null Cases

In the following example, Tunnel 1 is shut down, and only a label switched path (LSP) formed with Label Distribution Protocol (LDP) labels is established. An implicit null is advertised between the P2 and PE2 devices. Entering an MPLS LSP Traceroute at the PE1 device results in the following display:

```
Device# trace mpls ipv4 10.131.159.252/32
Tracing MPLS Label Switched Path to 10.131.159.252/32, timeout is 2 seconds
Codes: '!' - success, 'Q' - request not transmitted,
       '.' - timeout, 'U' - unreachable,
       'R' - downstream router but not target
Type escape sequence to abort.
  0 10.131.191.230 MRU 1500 [Labels: 20 Exp: 0]
R 1 10.131.159.226 MRU 1500 [Labels: 19 Exp: 0] 80 ms
R 2 10.131.159.229 MRU 1504 [implicit-null] 28 ms
! 3 10.131.159.230 40 ms
```

This output shows that packets are forwarded from P2 to PE2 with an implicit-null label. Address 10.131.159.229 is configured for the P2 Fast Ethernet 0/0/0 out interface for the PE2 device.

Troubleshooting Untagged Cases

Untagged cases are valid configurations for Interior Gateway Protocol (IGP) label switched paths (LSPs) that could cause problems for Multiprotocol Label Switching (MPLS) virtual private networks (VPNs).

A **show mpls forwarding-table** command and a **show mpls ldp discovery** command issued at the P2 device show that the Label Distribution Protocol (LDP) is properly set up:

```
Device# show mpls forwarding-table 10.131.159.252

Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
19     Pop tag     10.131.159.252/32 0          Et0/0      10.131.159.230
Device# show mpls ldp discovery
 Local LDP Identifier:
    10.131.159.251:0
    Discovery Sources:
    Interfaces:
            (ldp): xmit/recv
             LDP Id: 10.131.159.252:0
         FastEthernet1/0/0 (ldp): xmit/recv
             LDP Id: 10.131.191.251:0
```

The **show mpls ldp discovery** command output shows that, which connects PE2 to P2, is sending and receiving packets.

If a **no mpls ip** command is entered on , this could prevent an LDP session between the P2 and PE2 devices from being established. A **show mpls ldp discovery** command entered on the PE device shows that the MPLS LDP session with the PE2 device is down:

```
Device# show mpls ldp discovery

 Local LDP Identifier:
    10.131.159.251:0
    Discovery Sources:
    Interfaces:
        (ldp): xmit
        FastEthernet1/0/0 (ldp): xmit/recv
            LDP Id: 10.131.191.251:0
```

If the MPLS LDP session to PE2 goes down, the LSP to 10.131.159.252 becomes untagged, as shown by the **show mpls forwarding-table** command:

```
Device# show mpls forwarding-table 10.131.159.252

Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
19     Untagged
    10.131.159.252/32 864        Et0/0      10.131.159.230
```

Untagged cases would provide an MPLS LSP Traceroute reply with packets tagged with No Label, as shown in the following display:

```
Device# trace mpls ipv4 10.131.159.252/32
Tracing MPLS Label Switched Path to 10.131.159.252/32, timeout is 2 seconds
Codes: '!' - success, 'Q' - request not transmitted,
       '.' - timeout, 'U' - unreachable,
       'R' - downstream router but not target
Type escape sequence to abort.
  0 10.131.191.230 MRU 1500 [Labels: 20 Exp: 0]
```

```
R 1 10.131.159.226 MRU 1500 [Labels: 19 Exp: 0] 80 ms
R 2 10.131.159.229 MRU 1504 [No Label] 28 ms
! 3 10.131.159.230 40 ms
```

MPLS LSP Ping and Traceroute Returns a Q

The Q return code always means that the packet could not be transmitted. The problem can be caused by insufficient memory, but it probably results because a label switched path (LSP) could not be found that matches the Forwarding Equivalence Class (FEC), information that was entered on the command line.

The reason that the packet was not forwarded needs to be determined. To do so, look at the Routing Information Base (RIB), the Forwarding Information Base (FIB), the Label Information Base (LIB), and the MPLS Label Forwarding Information Base (LFIB). Lack of an entry for the FEC in any one of these routing/forwarding bases would return a Q.

The table below lists commands that you can use for troubleshooting when the MPLS echo reply returns a Q.

Table 11: Troubleshooting a Q

Database	Command to View Contents
Routing Information Base	**show ip route**
Label Information Base and MPLS Forwarding Information Base	**show mpls forwarding-table detail**

The following example shows a **ping mpls** command where the MPLS echo request is not transmitted, as shown by the returned Qs:

```
Device# ping mpls ipv4 10.0.0.1/32
Sending 5, 100-byte MPLS Echos to 10.0.0.1/32,
      timeout is 2 seconds, send interval is 0 msec:
Codes: '!' - success, 'Q' - request not transmitted,
       '.' - timeout, 'U' - unreachable,
       'R' - downstream router but not target
Type escape sequence to abort.
QQQQQ
Success rate is 0 percent (0/5)
```

A **show mpls forwarding-table** command and a **show ip route** command demonstrate that the address is not in either routing table:

```
Device# show mpls forwarding-table 10.0.0.1

Local  Outgoing    Prefix           Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id     switched   interface
Device# show ip route 10.0.0.1

% Subnet not in table
```

The MPLS echo request is not transmitted because the IPv4 address (10.0.0.1) is not found in either the LFIB or the RIB routing table.

Load Balancing for IPv4 LDP LSPs

An Internet Control Message Protocol (ICMP) ping or trace follows one path from the originating device to the target device. Round robin load balancing of IP packets from a source device is used to discover the various output paths to the target IP address.

For MPLS LSP Ping and Traceroute, Cisco devices use the source and destination addresses in the IP header for load balancing when multiple paths exist through the network to a target device. The Cisco implementation of MPLS might check the destination address of an IP payload to accomplish load balancing (this checking depends on the platform).

To check for load balancing paths, you use the 127.*z.y.x* /8 destination address in the **ping mpls ipvr** *ip-address address-mask* **destination** *address-start address-end address-increment* command. The following examples show that different paths are followed to the same destination. This demonstrates that load balancing occurs between the originating device and the target device.

To ensure that the Fast Ethernet interface 1/0/0 on the PE1 device is operational, you enter the following commands on the PE1 device:

```
Device# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Device(config)# interface fastethernet 1/0/0
Device(config-if)# no shutdown
Device(config-if)# end
*Dec 31 19:14:10.034: %LINK-3-UPDOWN: Interface FastEthernet1/0/0, changed state to up
*Dec 31 19:14:11.054: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet1/0/0,
changed state to upend
PE1#
*Dec 31 19:14:12.574: %SYS-5-CONFIG_I: Configured from console by console
*Dec 31 19:14:19.334: %OSPF-5-ADJCHG: Process 1, Nbr 10.131.159.252 on FastEthernet1/0/0
from LOADING to FULL, Loading Done
PE1#
```

The following **show mpls forwarding-table** command displays the possible outgoing interfaces and next hops for the prefix 10.131.159.251/32:

```
Device# show mpls forwarding-table 10.131.159.251

Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
21     19          10.131.159.251/32 0          FE0/0/0 10.131.191.229
       20          10.131.159.251/32 0          FE1/0/0 10.131.159.245
```

The following **ping mpls** command to 10.131.159.251/32 with a destination UDP address of 127.0.0.1 shows that the path selected has a path index of 0:

```
Device# ping mpls ipv4
 10.131.159.251/32 destination
 127.0.0.1 repeat 1
Sending 1, 100-byte MPLS Echos to 10.131.159.251/32,
      timeout is 2 seconds, send interval is 0 msec:
Codes: '!' - success, 'Q' - request not transmitted,
       '.' - timeout, 'U' - unreachable,
       'R' - downstream router but not target
Type escape sequence to abort.
!
Success rate is 100 percent (1/1), round-trip min/avg/max = 40/40/40 ms
PE1#
*Dec 29 20:42:40.638: LSPV: Echo Request sent on IPV4 LSP, load_index 2,
pathindex 0
, size 100
*Dec 29 20:42:40.638: 46 00 00 64 00 00 40 00 FF 11 9D 03 0A 83 BF FC
*Dec 29 20:42:40.638: 7F 00 00 01 94 04 00 00 0D AF 0D AF 00 4C 14 70
*Dec 29 20:42:40.638: 00 01 00 00 01 02 00 00 1A 00 00 1C 00 00 00 01
*Dec 29 20:42:40.638: C3 9B 10 40 A3 6C 08 D4 00 00 00 00 00 00 00 00
*Dec 29 20:42:40.638: 00 01 00 09 00 01 00 05 0A 83 9F FB 20 00 03 00
```

```
*Dec 29 20:42:40.638: 13 01 AB CD AB CD AB CD AB CD AB CD AB CD
*Dec 29 20:42:40.638: AB CD AB CD
*Dec 29 20:42:40.678: LSPV: Echo packet received: src 10.131.159.225,
dst 10.131.191.252, size 74
*Dec 29 20:42:40.678: AA BB CC 00 98 01 AA BB CC 00 FC 01 08 00 45 C0
*Dec 29 20:42:40.678: 00 3C 32 D6 00 00 FD 11 15 37 0A 83 9F E1 0A 83
*Dec 29 20:42:40.678: BF FC 0D AF 0D AF 00 28 D1 85 00 01 00 00 02 02
*Dec 29 20:42:40.678: 03 00 1A 00 00 1C 00 00 00 01 C3 9B 10 40 A3 6C
*Dec 29 20:42:40.678: 08 D4 C3 9B 10 40 66 F5 C3 C8
```

The following **ping mpls** command to 10.131.159.251/32 with a destination UDP address of 127.0.0.1 shows that the path selected has a path index of 1:

```
Device# ping mpls ipv4 10.131.159.251/32 dest 127.0.0.1 repeat 1
Sending 1, 100-byte MPLS Echos to 10.131.159.251/32,
      timeout is 2 seconds, send interval is 0 msec:
Codes: '!' - success, 'Q' - request not transmitted,
       '.' - timeout, 'U' - unreachable,
       'R' - downstream router but not target
Type escape sequence to abort.
!
Success rate is 100 percent (1/1), round-trip min/avg/max = 40/40/40 ms
*Dec 29 20:43:09.518: LSPV: Echo Request sent on IPV4 LSP, load_index 13,
pathindex 1
, size 100
*Dec 29 20:43:09.518: 46 00 00 64 00 00 40 00 FF 11 9D 01 0A 83 BF FC
*Dec 29 20:43:09.518: 7F 00 00 03 94 04 00 00 0D AF 0D AF 00 4C 88 58
*Dec 29 20:43:09.518: 00 01 00 00 01 02 00 00 38 00 00 1D 00 00 00 01
*Dec 29 20:43:09.518: C3 9B 10 5D 84 B3 95 84 00 00 00 00 00 00 00 00
*Dec 29 20:43:09.518: 00 01 00 09 00 01 00 05 0A 83 9F FB 20 00 03 00
*Dec 29 20:43:09.518: 13 01 AB CD AB CD AB CD AB CD AB CD AB CD
*Dec 29 20:43:09.518: AB CD AB CD
*Dec 29 20:43:09.558: LSPV: Echo packet received: src 10.131.159.229,
dst 10.131.191.252, size 74
*Dec 29 20:43:09.558: AA BB CC 00 98 01 AA BB CC 00 FC 01 08 00 45 C0
*Dec 29 20:43:09.558: 00 3C 32 E9 00 00 FD 11 15 20 0A 83 9F E5 0A 83
*Dec 29 20:43:09.558: BF FC 0D AF 0D AF 00 28 D7 57 00 01 00 00 02 02
*Dec 29 20:43:09.558: 03 00 38 00 00 1D 00 00 00 01 C3 9B 10 5D 84 B3
*Dec 29 20:43:09.558: 95 84 C3 9B 10 5D 48 3D 50 78
```

To see the actual path chosen, you use the **debug mpls lspv packet data** command.

Note The hashing algorithm is nondeterministic. Therefore, the selection of the *address-start* , *address-end* , and *address-increment* arguments for the **destination** keyword might not provide the expected results.

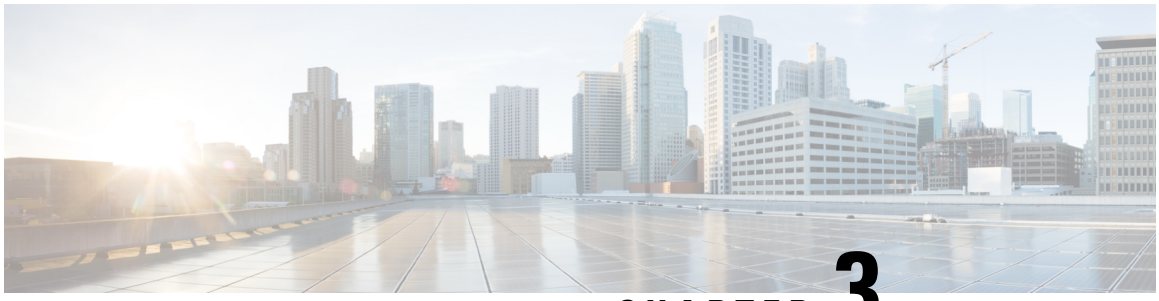

CHAPTER 3

NSR LDP Support

The NSR LDP Support feature allows the Label Distribution Protocol (LDP) to continue to operate across a Router Processor (RP) failure in redundant systems, without losing peer sessions. Before the introduction of nonstop routing (NSR), LDP sessions with peers reset if an RP failover (in a redundant system) or a Cisco In-Service Software Upgrade (ISSU) occurred. When peers reset, traffic is lost while the session is down. Protocol reconvergence occurs after the session is reestablished.

When NSR is enabled, RP failover and Cisco ISSU events are not visible to the peer device, and the LDP sessions that were established prior to failover do not flap. The protocol state learned from the peers persists across an RP failover or Cisco ISSU event and does not need to be relearned.

Prerequisites for NSR LDP Support

The Label Distribution Protocol (LDP) must be up and running on the standby Route Processor (RP) for NSR LDP Support to work.

Information About NSR LDP Support

Roles of the Standby Route Processor and Standby LDP

For the NSR LDP Support feature to work, the Label Distribution Protocol (LDP) must be up and running on the standby Route Processor (RP). The LDP component running on the active RP is called the active LDP, and the LDP component running on the standby RP is called the standby LDP.

When nonstop routing (NSR) is enabled, the standby LDP runs independently from the active LDP, but with the assistance of some software components. The standby LDP maintains LDP session states and database information, ready to take over for the active LDP if the failover occurs.

Standby LDP maintains its local database by querying or receiving notifications of interface status change, configuration changes from the CLI, and checkpoints from the active LDP for other information that is not directly available on the standby RP.

To keep the protocol and session-state information synchronized with the active LDP, the standby LDP depends on TCP to replicate all LDP messages on the active RP to the standby RP. The standby LDP processes all received messages, updates its state, but does not send any responses to its neighbors.

The standby LDP performs the following tasks:

- Processes LDP configuration on startup and during steady state
- Processes active LDP checkpoints of state and session information such as LDP adjacencies, remote addresses, remote bindings, and so forth
- Builds its database of local interfaces
- Processes interface change events
- Receives and processes all LDP messages replicated by TCP
- Updates remote address and label databases

After a switchover and notification that the RP has become active, the standby LDP takes over the role of the active LDP and performs the following tasks:

- Sends hello messages immediately to prevent neighbors from reaching the discovery timeout and bringing down the session
- Retransmits any protocol-level response that has not been sent by the previous active LDP
- Readvertises label bindings
- Refreshes all forwarding entries
- Processes and responds to any LDP message from its neighbor

When the NSR LDP Support feature is disabled, the active LDP performs the following tasks:

- Stops checkpointing to the standby LDP
- Continues to manage all existing sessions

The standby LDP performs the following tasks:

- Cleans up all session-state information
- Reverses to the behavior before NSR is enabled

LDP Operating States

When the NSR LDP Support feature is enabled, the Label Distribution Protocol (LDP) operates in the following states:

Initial State

In the initial state, the active Label Distribution Protocol (LDP) process sets up the standby LDP to be ready to support nonstop routing (NSR). The active LDP performs the following tasks:

- Replicates all TCP sessions used by LDP with the standby LDP
- Synchronizes all existing session-state information with the standby LDP
- Synchronizes the LDP database with the standby LDP

LDP could be in the initial state because of one of these conditions:

- NSR is enabled
- NSR was enabled and the standby Route Processor (RP) starts up (asymmetric startup)
- System boots up and NSR is configured (symmetric startup)

Steady State

In the steady state, the active and standby Label Distribution Protocol (LDP) databases are synchronized. The active and standby LDP process the same LDP messages and update their states independently. The standby LDP is ready to take over the active LDP role in a switchover event.

On the active Route Processor (RP), the active LDP performs the following tasks:

- Continues to manage all existing sessions and checkpoints any significant session event to the standby LDP (such as adjacency delete, session shutdown, timers)
- Notifies the standby LDP of new adjacencies and neighbors

On the standby RP, the standby LDP performs these tasks:

- Processes all received messages but does not send any messages to its neighbor
- Processes checkpoint information from the active LDP
- Manages session keepalive timers but does not bring down the session if a keepalive timer times out

Post Switchover

In the post switchover state, the standby Label Distribution Protocol (LDP) process takes over the active LDP role while the active Route Processor (RP) is reloading.

Supported NSR Scenarios

The NSR LDP Support feature is supported under the following scenarios:

- Route Processor (RP) failover or node failure

  The Label Distribution Protocol (LDP) keeps the session up during an RP or node failover because the LDP adjacency and session-state information between LDP on the active and standby RPs are synchronized. As sessions are created on the active RP, new adjacencies are synchronized to the standby RP. If a standby RP is brought online after sessions are already up (asymmetric startup), LDP synchronizes the existing session-state information from the active to the standby RP.

- Cisco In-Service Software Upgrade (ISSU)

  LDP supports Cisco ISSU negotiation between RPs when a standby RP comes online for the MPLS LDP IGP Synchronization feature. Current Cisco ISSU negotiation is not impacted by NSR. For NSR, LDP negotiates messages specific to NSR, which are checkpointed during initial synchronization (adjacency and session-state information).

How to Configure NSR LDP Support

Enabling NSR LDP Support

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **mpls ldp nsr**
4. **exit**
5. **show mpls ldp nsr**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**mpls ldp nsr** **Example:** `Device(config)# mpls ldp nsr`	Enables nonstop routing (NSR) for all Label Distribution Protocol (LDP) sessions for both link and targeted.
Step 4	**exit** **Example:** `Device(config)# exit`	Returns to privileged EXEC mode.
Step 5	**show mpls ldp nsr** **Example:** `Device# show mpls ldp nsr`	Displays whether NSR is enabled.

Troubleshooting Tips for NSR LDP Support

Use the **debug mpls ldp nsr** command to enable the display of Multiprotocol Label Switching (MPLS) Label Distribution Protocol (LDP) nonstop routing (NSR) debugging events for all NSR sessions or for the specified peer.

Configuration Examples for NSR LDP Support

Example: NSR LDP Configuration

Additional References for NSR LDP Support

Related Documents

Related Topic	Document Title
Cisco IOS commands	Cisco IOS Master Command List, All Releases
MPLS commands	Cisco IOS Multiprotocol Label Switching Command Reference
LDP configuration tasks	*MPLS Label Distribution Protocol Configuration Guide*

Technical Assistance

Description	Link
The Cisco Support and Documentation website provides online resources to download documentation, software, and tools. Use these resources to install and configure the software and to troubleshoot and resolve technical issues with Cisco products and technologies. Access to most tools on the Cisco Support and Documentation website requires a Cisco.com user ID and password	http://www.cisco.com/cisco/web/support/index.html

Feature Information for NSR LDP Support

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Table 12: Feature Information for NSR LDP Support

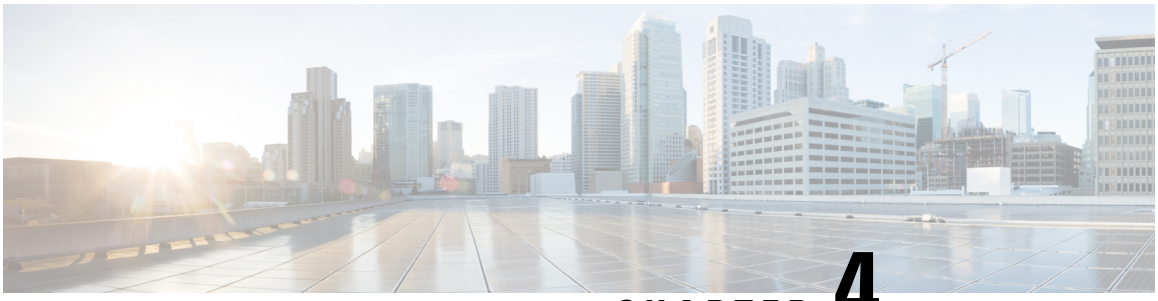

CHAPTER 4

Flex LSP Overview

Flex LSP also known as Associated Bidirectional LSPs is the combination of static bidirectional MPLS-TP and dynamic MPLS-TE. Flex LSP provides bidirectional label switched paths (LSPs) set up dynamically through Resource Reservation Protocol–Traffic Engineering (RSVP-TE). It does not support non-co routed LSPs.

Flex Label Switched Paths are LSP instances where the forward and the reverse direction paths are setup, monitored and protected independently and associated together during signaling. You use a RSVP Association object to bind the two forward and reverse LSPs together to form either a co-routed or non co-routed associated bidirectional TE tunnel.

You can associate a protecting MPLS-TE tunnel with either a working MPLS-TE LSP, protecting MPLS-TE LSP, or both. The working LSP is the primary LSP backed up by the protecting LSP. When a working LSP goes down, the protecting LSP is automatically activated. You can configure a MPLS-TE tunnel to operate without protection as well.

Effective Cisco IOS XE Release 3.18.1SP, Flex LSP supports inter-area tunnels with non co-routed mode.

Signaling Methods and Object Association for Flex LSPs

This section provides an overview of the association signaling methods for the bidirectional LSPs. Two unidirectional LSPs can be bound to form an associated bidirectional LSP in the following scenarios:

- No unidirectional LSP exists, and both must be established.
- Both unidirectional LSPs exist, but the association must be established.
- One unidirectional LSP exists, but the reverse associated LSP must be established.

Associated Bidirectional Non Co-routed and Co-routed LSPs

This section provides an overview of associated bidirectional non co-routed and co-routed LSPs. Establishment of MPLS TE-LSP involves computation of a path between a head-end node to a tail-end node, signaling along the path, and modification of intermediate nodes along the path. The signaling process ensures bandwidth reservation (if signaled bandwidth is lesser than 0 and programming of forwarding entries).

Path computation is performed by the head-end nodes of both the participating LSPs using Constrained Shortest Path First (CSPF). CSPF is the 'shortest path (measured in terms of cost) that satisfies all relevant LSP TE constraints or attributes, such as required bandwidth, priority and so on.

Associated Bidirectional Non Co-routed LSPs: A non co-routed bidirectional TE LSP follows two different paths, that is, the forward direction LSP path is different than the reverse direction LSP path. Here is an illustration.

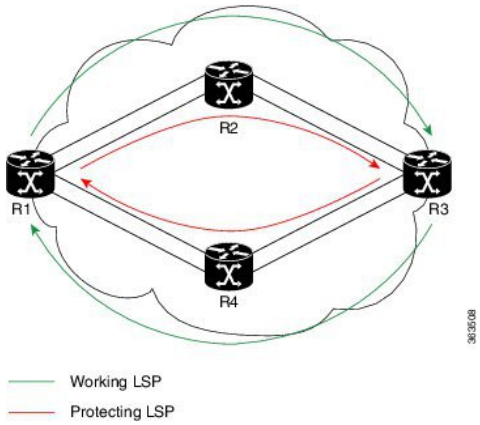

In the above topology:

- The outer paths (in green) are working LSP pairs.
- The inner paths (in red) are protecting LSP pairs.
- Router 1 sets up working LSP to Router 3 and protecting LSP to Router 3 independently.
- Router 3 sets up working LSP to Router 1 and protecting LSP to Router 1 independently.

Non co-routed bidirectional TE LSP is available by default, and no configuration is required.

Note In case of non co-routed LSPs, the head-end nodes relax the constraint on having identical forward and reverse paths. Hence, depending on network state you can have identical forward and reverse paths, though the bidirectional LSP is co-routed.

Associated Bidirectional Co-routed LSPs: A co-routed bidirectional TE LSP denotes a bidirectional tunnel where the forward direction LSP and reverse direction LSP must follow the same path, for example, the same nodes and paths. Here is an illustration.

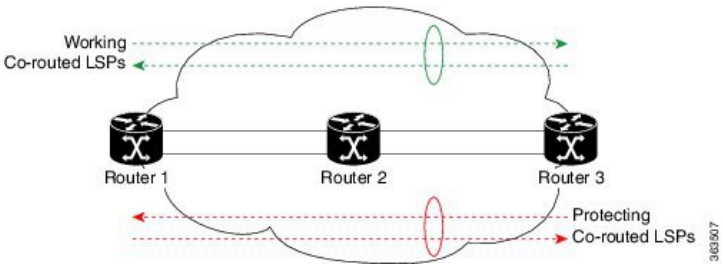

In the above topology:

- Paths at the top of the figure (in green) indicate working co-routed LSP pairs.
- Paths at the bottom of the figure (in red) indicate protecting co-routed LSP pairs.
- Router 1 sets up working LSP to Router 3 (in red) after performing bidirectional CSPF and sends reverse explicit route object (ERO) to Router 3. Node Router 3 uses the received reverse ERO to set up reverse red working LSP to Router 1.
- Router 3 sets up protecting LSP to Router 1 (in green) after performing bidirectional CSPF and sends reverse ERO to Router 1. Node Router 1 uses the received reverse ERO to set up reverse green protecting LSP to Router 3.

Restrictions for Flex LSP

- Exp-null over Flex-LSP is not supported.
- Flex-LSP does not support tunnel statistics.
- VC (layer 2 VPN ckts) statistics are not supported.
- It is recommended to configure for the following timers for Flex-LSP deployments:

  mpls traffic-eng reoptimize timers frequency 120

  mpls traffic-eng reoptimize timers delay installation 30

  mpls traffic-eng reoptimize timers delay cleanup 90
- The **no mpls ip propagate-tcl** command is not recommended with Flex-LSP. The PREC value of BFD control packet is set to “0”. Therefore, packet prioritization cannot be done at midpoints and BFD flap can occur with traffic congestions.
- It is recommended to configure BFD timers as 10x3 during cable pull testing or in Flex LSP feature deployments.
- 50-msec convergence is not guaranteed for local shut.
- 50-msec convergence is not guaranteed without WRAP protection. WRAP protection is mandatory to achieve 50-msec convergence for remote failures.
- 50-msec convergence is expected with a maximum of 30 Flex-LSP tunnels that use the same egress interface.
- With scale and multiple other feature mix-ups, it is possible to see higher convergence.
- TE NSR and IGP NSR are mandatory for RSP switchover.
- Flex LSP is supported with the IPv4 template.

- The **ip rsvp signalling hello** command is not mandatory and it can cause a large punt during the cutover and can lead to unexpected results like BFD flapping.
- VPLS over Flex-LSP is not supported.
- Both IGP and FRR must be configured as clients for single-hop BFD when the WRAP protection is enabled; only FRR cannot be the only client that is configured at midpoint.
- Layer 3 VPN over Flex-LSP is not supported.
- It is recommended to configure 10x3 BFD timers for cable failures, to achieve 50 msec of convergence.
- Dynamic diverse paths are not supported for Flex LSP Tunnel.
- The Diverse node SRLG path option is not supported.
- The protect dynamic SRLG path is diverse from the primary path and thus the shortest path is not always chosen.
- When the constraint for the protect path of Flex-LSP tunnel does not meet, it will wait in the REQUESTED state.

Restrictions for Non Co-routed Inter-Area Flex LSP Tunnels

- The dynamic path option feature for TE tunnels (**tunnel mpls traffic-eng path-option number dynamic**) is not supported for inter-area tunnels. An explicit path identifying the area border routers (ABRs) is required.
- The MPLS TE AutoRoute feature (**tunnel mpls traffic-eng autoroute announce**) is not supported for inter-area tunnels.
- Tunnel affinity (**tunnel mpls traffic-eng affinity**) is not supported for inter-area tunnels.
- Tunnel metric (**tunnel mpls traffic-eng path-selection metric**) is not supported for inter-area tunnels.
- BFD is not supported with non co-routed inter-area flex LSP tunnels.

How to Configure Co-routed Flex LSPs

A co-routed bidirectional packet LSP is a combination of two LSPs (one in the forward direction and the other in reverse direction) sharing the same path between a pair of ingress and egress nodes. It is established using the extensions to RSVP-TE. This type of LSP can be used to carry any of the standard types of MPLS-based traffic, including Layer 2 VPNs and Layer 2 circuits. You can configure a single BFD session for the bidirectional LSP (that is, you do not need to configure a BFD session for each LSP in each direction). You can also configure a single standby bidirectional LSP to provide a backup for the primary bidirectional LSP.

The configuration includes the following steps:

1. Enable basic MPLS Traffic Engineering on hostname PE1.
2. Map L2VPN pseudowire to a specific FLEX LSP tunnel.
3. Configure Flex LSP.
4. Enable BFD.

5. Enable Wrap and Fault OAM.
6. Enable BDIs on a core-facing interface.

Configuring Co-routed Flex LSPs

Before you begin

- You must have symmetric source and destination TE router IDs in order for bidirectional LSPs to be associated.
- Tunnels attributes must be configured identically on both sides of co-routed bidirectional LSP.

Note Up to 250 Flex LSP tunnels are supported.

Procedure

1. Enable basic MPLS Traffic Engineering on hostname PE1:

```
mpls traffic-eng tunnels
mpls traffic-eng fault-oam
mpls traffic-eng nsr
router ospf 100
 router-id 1.1.1.1
 nsr
 mpls traffic-eng router-id Loopback0
 mpls traffic-eng area 0
```

2. Map L2VPN pseudowire to a specific Flex LSP tunnel:

```
template type pseudowire mpls-te1 (mpls-te1 can be any name)
encapsulation mpls
preferred-path interface Tunnel1 disable-fallback
bandwidth 100

template type pseudowire mpls-te4
encapsulation mpls
preferred-path interface Tunnel4 disable-fallback
bandwidth 100
```

3. Configure Flex LSP:

```
interface Tunnel1
bandwidth 1000
ip unnumbered Loopback0
tunnel mode mpls traffic-eng
tunnel destination 22.22.22.22
tunnel mpls traffic-eng autoroute announce
tunnel mpls traffic-eng priority 7 7
tunnel mpls traffic-eng bandwidth 1000
tunnel mpls traffic-eng path-option 1 explicit name BDI1 bandwidth 1000
tunnel mpls traffic-eng path-option protect 1 explicit name BACKUP1 bandwidth 1000
tunnel mpls traffic-eng bidirectional association id 1 source-address 11.11.11.11 global-id
1
tunnel mpls traffic-eng bidirectional association type co-routed
ip explicit-path name BDI1 enable
```

```
next-address 1.11.1.1
next-address 10.1.2.2
next-address 2.22.1.22
ip explicit-path name BACKUP1 enable
next-address 10.3.11.1.10
next-address 10.4.22.22
```

Note To bring up the bi-directional tunnels, association ID, source address and global ID must match on both sides of the tunnel.

4. Enable BFD

```
bfd-template single-hop BFD_FLEX
interval min-tx 50 min-rx 50 multiplier 3
interface Tunnel1
tunnel mpls traffic-eng bfd encap-mode gal BFD_FLEX
```

5. Enable Wrap and Fault OAM

```
interface Tunnel1
tunnel mpls traffic-eng bidirectional association type co-routed fault-oam wrap-protection
```

6. Enable BDIs on core-facing interface:

```
interface BDI1
ip address 1.11.1.11 255.255.255.0
ip ospf 1 area 0
mpls traffic-eng tunnels


interface BDI4
ip address 1.11.4.11 255.255.255.0
ip ospf 1 area 0
mpls traffic-eng tunnels

interface GigabitEthernet0/3/1
ip address 10.3.11.11 255.255.255.0
ip ospf 1 area 0
mpls traffic-eng tunnels

interface GigabitEthernet0/3/0
service instance 1 ethernet
encapsulation dot1q 1
rewrite ingress tag pop 1 symmetric
bridge-domain 1
service instance 4 ethernet
encapsulation dot1q 4
rewrite ingress tag pop 1 symmetric
bridge-domain 4
End
```

Note NOTE: Since VLANs are not supported, to represent a VLAN interface, BDI must be used towards core-facing interface.

Verifying the Co-routed Flex LSP Configuration

To verify the FLEX LSP tunnel summary, use the **show mpls traffic-eng tunnels bidirectional-associated concise** command in MPLS tunnel-te interface.

```
Router# show mpls traffic-eng tunnels summary
Signalling Summary:
    LSP Tunnels Process:            running
    Passive LSP Listener:           running
    RSVP Process:                   running
    Forwarding:                     enabled
    auto-tunnel:
    p2p    Disabled (0), id-range:62336-64335

    Periodic reoptimization:        every 3600 seconds, next in 2942 seconds
    Periodic FRR Promotion:         Not Running
    Periodic auto-bw collection:    every 300 seconds, next in 243 seconds
    SR tunnel max label push:       1 labels
    P2P:
      Head: 100 interfaces,   0 active signalling attempts, 0 established
            87733091 activations,  87733091 deactivations
            144287155 failed activations
            0 SSO recovery attempts, 0 SSO recovered
      Midpoints: 0, Tails: 0

    P2MP:
      Head: 0 interfaces,   0 active signalling attempts, 0 established
            0 sub-LSP activations,  0 sub-LSP deactivations
            0 LSP successful activations,  0 LSP deactivations
            0 SSO recovery attempts, LSP recovered: 0 full, 0 partial, 0 fail
      Midpoints: 0, Tails: 0

Bidirectional Tunnel Summary:
    Tunnel Head: 100 total, 0 connected, 100 associated, 100 co-routed
    LSPs Head:   0 established, 0 proceeding, 0 associated, 0 standby
    LSPs Mid:    0 established, 0 proceeding, 0 associated, 0 standby
    LSPs Tail:   0 established, 0 proceeding, 0 associated, 0 standby
```

To verify the co-routed LSP, use the **Show mpls traffic-eng tunnel bidirectional co-routed** command.

```
Router#Show mpls traffic-eng tunnel bidirectional co-routed

Name: tunnel-te2  Destination: 192.168.0.3
  Status:
    Admin:    up Oper:   up   Path:  valid   Signalling: connected
    path option 1,  type dynamic  (Basis for Setup, path weight 3 (reverse 3))
    Bandwidth Requested: 80000 kbps  CT0
Config Parameters:
    Association Type: Single Sided Bidirectional LSPs, Co-routed: Yes
    Association ID: 100, Source: 9.9.9.9[, Global ID: 9]
    Reverse Bandwidth: 2 kbps CT0, Standby: 2 kbps CT0
    BFD Fast Detection: Enabled
    BFD Parameters: Min-interval 10000 ms, Multiplier 3 (default)
     BFD Bringup Timeout: Interval 60 seconds (default)
     BFD Initial Dampening: 16000 ms (default)
     BFD Maximum Dampening: 600000 ms (default)
     BFD Secondary Dampening: 20000 ms (default)
     Periodic LSP Ping: Interval 120 seconds (default)
     BFD Encap Mode: IP (default) | GAL
    Soft Preemption: Enabled, Current Status: Preemption not pending
```

How to Configure Non Co-routed Inter-area Flex LSP Tunnels

Note The working and protect LSPs for PE1 (head-end) is different from PE2 (tail-end).

At PE1 (head-end):

```
interface Tunnel1001
 ip unnumbered Loopback0
 mpls ip
 tunnel mode mpls traffic-eng
 tunnel destination 1.1.1.1
 tunnel mpls traffic-eng priority 7 7
 tunnel mpls traffic-eng bandwidth 200
 tunnel mpls traffic-eng path-option 1 explicit name ThruHunG verbatim
 tunnel mpls traffic-eng path-option protect 1 explicit name PROT1 verbatim
 tunnel mpls traffic-eng bidirectional association id 1001 source-address 1.1.1.1 global-id
 1001
!
interface Tunnel1002
 ip unnumbered Loopback0
 mpls ip
 tunnel mode mpls traffic-eng
 tunnel destination 1.1.1.1
 tunnel mpls traffic-eng priority 7 7
 tunnel mpls traffic-eng bandwidth 200
 tunnel mpls traffic-eng path-option 1 explicit name ThruHunG verbatim
 tunnel mpls traffic-eng path-option protect 1 explicit name PROT1 verbatim
 tunnel mpls traffic-eng bidirectional association id 1002 source-address 1.1.1.1 global-id
 1002

ip explicit-path name ThruTenG enable
 next-address loose 22.1.1.2
 next-address loose 10.1.1.1
 next-address loose 1.1.1.1
!
ip explicit-path name ThruHunG enable
 next-address loose 23.1.1.2
 next-address loose 10.1.1.1
 next-address loose 1.1.1.1

ip explicit-path name PROT1 enable
 next-address loose 30.1.1.2
 next-address loose 40.1.1.1
 next-address loose 1.1.1.1
```

At PE2 (tail-end):

```
interface Tunnel1001
 ip unnumbered Loopback0
 mpls ip
 tunnel mode mpls traffic-eng
 tunnel destination 4.4.4.4
 tunnel mpls traffic-eng priority 7 7
 tunnel mpls traffic-eng bandwidth 200
 tunnel mpls traffic-eng path-option 1 explicit name ThruTenG verbatim
 tunnel mpls traffic-eng path-option protect 1 explicit name PROT2 verbatim
 tunnel mpls traffic-eng bidirectional association id 1001 source-address 1.1.1.1 global-id
 1001
!
```

```
interface Tunnel1002
 ip unnumbered Loopback0
 mpls ip
 tunnel mode mpls traffic-eng
 tunnel destination 4.4.4.4
 tunnel mpls traffic-eng priority 7 7
 tunnel mpls traffic-eng bandwidth 200
 tunnel mpls traffic-eng path-option 1 explicit name ThruTenG verbatim
 tunnel mpls traffic-eng path-option protect 1 explicit name PROT2 verbatim
 tunnel mpls traffic-eng bidirectional association id 1002 source-address 1.1.1.1 global-id
 1002

ip explicit-path name ThruTenG enable
 next-address loose 10.1.1.2
 next-address loose 22.1.1.1
 next-address loose 4.4.4.4
!
ip explicit-path name ThruHunG enable
 next-address loose 10.1.1.2
 next-address loose 23.1.1.1
 next-address loose 4.4.4.4

ip explicit-path name PROT2 enable
 next-address loose 41.1.1.2
 next-address loose 31.1.1.1
 next-address loose 4.4.4.4
```

Configuring OSFP for Non Co-routed Flex LSP

Note Add the new area into OSPF based on where you want the Inter-area to run.

```
router ospf 1
 router-id 3.3.3.3
 nsr
 nsf cisco
 microloop avoidance
 passive-interface Loopback0
 network 3.3.3.3 0.0.0.0 area 0
 mpls traffic-eng router-id Loopback0
 mpls traffic-eng area 0
 mpls traffic-eng area 1
```

Verifying the Non Co-routed Inter-area Flex LSP Tunnels

At the PE1

```
Router# show mpls traffic-eng tunnels tunnel 1001

Name: PE1_t1001                           (Tunnel1001) Destination: 4.4.4.4
  Status:
    Admin: up         Oper: up     Path: valid       Signalling: connected
    path option 1, type explicit (verbatim) ThruTenG (Basis for Setup, path weight 0)
    Path Protection: Requested
    path protect option 1, type explicit (verbatim) PROT2 (Basis for Protect, path weight
0)

  Config Parameters:
    Bandwidth: 200      kbps (Global)  Priority: 7  7   Affinity: 0x0/0xFFFF
```

```
    Metric Type: TE (default)
    AutoRoute: disabled LockDown: disabled Loadshare: 200 [10000000] bw-based
    auto-bw: disabled
    Association Type: Double Sided Bidirectional LSPs, Co-routed: NO
    Association ID: 1001, Source: 1.1.1.1, Global ID: 1001
    Fault-OAM: disabled, Wrap-Protection: disabled, Wrap-Capable: No
  Active Path Option Parameters:
    State: explicit path option 1 is active
    BandwidthOverride: disabled  LockDown: disabled  Verbatim: enabled

  InLabel  :  -
  OutLabel : BDI100, 242
  Next Hop : 10.1.1.2
  Reverse Associated LSP Information:
    Signaled Name: 4.4.4.4 1001
    Tunnel: 1001, Source: 4.4.4.4, Dest: 1.1.1.1, LSP: 9 State: Up
  Lockout Info:
    Locked out: No
    Lockout Originated By: None
  Association:
    Association Type: Double Sided Bidirectional LSPs
    Association ID: 1001 Source: 1.1.1.1
  Extended Association:
    Global source: 1001
    Extended ID: None
  RSVP Signalling Info:
       Src 1.1.1.1, Dst 4.4.4.4, Tun_Id 1001, Tun_Instance 9
    RSVP Path Info:
      My Address: 10.1.1.1
      Explicit Route: 10.1.1.2 10.1.1.2* 22.1.1.1* 4.4.4.4*
      Record   Route:
      Tspec: ave rate=200 kbits, burst=1000 bytes, peak rate=200 kbits
    RSVP Resv Info:
      Record   Route:  22.1.1.2 22.1.1.1
      Fspec: ave rate=200 kbits, burst=1000 bytes, peak rate=200 kbits
  Shortest Unconstrained Path Info:
    Path Weight: 2 (TE)
    Explicit Route: 11.1.1.2 20.1.1.1 4.4.4.4
  Reason for the tunnel being down: Bidirectional: standby error from [1.1.1.1][UNK] LSP[8]


  History:
    Tunnel:
      Time since created: 7 minutes, 51 seconds
      Number of LSP IDs (Tun_Instances) used: 9
    Current LSP: [ID: 9]
      Uptime: 5 minutes, 59 seconds
```

At PE2

```
Router# show mpls traffic-eng tunnels tunnel 1001

Name: PE2_t1001                           (Tunnel1001) Destination: 1.1.1.1
  Status:
    Admin: up         Oper: up     Path: valid       Signalling: connected
    path option 1, type explicit (verbatim) ThruHunG (Basis for Setup, path weight 0)
    Path Protection: Requested
    path protect option 1, type explicit (verbatim) PROT1 (Basis for Protect, path weight
0)

  Config Parameters:
    Bandwidth: 200      kbps (Global)  Priority: 7  7   Affinity: 0x0/0xFFFF
    Metric Type: TE (default)
    AutoRoute: disabled LockDown: disabled Loadshare: 200 [10000000] bw-based
    auto-bw: disabled
```

```
    Association Type: Double Sided Bidirectional LSPs, Co-routed: NO
    Association ID: 1001, Source: 1.1.1.1, Global ID: 1001
    Fault-OAM: disabled, Wrap-Protection: disabled, Wrap-Capable: No
  Active Path Option Parameters:
    State: explicit path option 1 is active
    BandwidthOverride: disabled  LockDown: disabled  Verbatim: enabled

  InLabel  :  -
  OutLabel : BDI221, 980
  Next Hop : 23.1.1.2
  Reverse Associated LSP Information:
    Signaled Name: 1.1.1.1 1001
    Tunnel: 1001, Source: 1.1.1.1, Dest: 4.4.4.4, LSP: 9 State: Up
  Lockout Info:
    Locked out: No
    Lockout Originated By: None
  Association:
    Association Type: Double Sided Bidirectional LSPs
    Association ID: 1001 Source: 1.1.1.1
  Extended Association:
    Global source: 1001
    Extended ID: None
  RSVP Signalling Info:
       Src 4.4.4.4, Dst 1.1.1.1, Tun_Id 1001, Tun_Instance 9
    RSVP Path Info:
      My Address: 23.1.1.1
      Explicit Route: 23.1.1.2 23.1.1.2* 10.1.1.1* 1.1.1.1*
      Record   Route:
      Tspec: ave rate=200 kbits, burst=1000 bytes, peak rate=200 kbits
    RSVP Resv Info:
      Record   Route:  10.1.1.2 10.1.1.1
      Fspec: ave rate=200 kbits, burst=1000 bytes, peak rate=200 kbits
  Shortest Unconstrained Path Info:
    Path Weight: 2 (TE)
    Explicit Route: 20.1.1.2 11.1.1.1 1.1.1.1
  Reason for the tunnel being down: Bidirectional: standby error from [4.4.4.4][UNK] LSP[8]


  History:
    Tunnel:
      Time since created: 8 minutes, 9 seconds
      Time since path change: 6 minutes, 10 seconds
      Number of LSP IDs (Tun_Instances) used: 9
    Current LSP: [ID: 9]
      Uptime: 6 minutes, 10 seconds
```

Troubleshooting Flex LSP

Step 1: Verifying that the Flex LSP Tunnel is in UP State

```
Router# show mpls traffic-eng tunnels bidirectional-associated association id 1

P2P TUNNELS/LSPs:
Name: RP1_t3                              (Tunnel3) Destination: 10.5.0.1
  Status:
    Admin: up         Oper: up     Path: valid       Signalling: connected
    path option 2, type explicit expl_route_m2_tail (Basis for Setup, path weight 40)
    path option 3, type explicit expl_route_m3_tail
    Path Protection: 0 Common Link(s), 0 Common Node(s)
    path protect option 2, type explicit expl_route_m3_tail (Basis for Protect, path weight
 40)
    path protect option 3, type list name xtd
```

```
    Lockout Info:
      Locked Out: No
  Config Parameters:
    Bandwidth: 500      kbps (Global)  Priority: 7  7   Affinity: 0x0/0xFFFF
    Metric Type: TE (default)
    AutoRoute: disabled LockDown: disabled Loadshare: 500 [4000000] bw-based
    auto-bw: disabled
    Association Type: Single Sided Bidirectional LSPs, Co-routed: YES
    Association ID: 1, Source: 2.3.4.5, Global ID: 6
    Fault-OAM: disabled
  Active Path Option Parameters:
    State: explicit path option 2 is active
    BandwidthOverride: disabled  LockDown: disabled  Verbatim: disabled
  InLabel  :  -
  OutLabel : Ethernet0/0, 16
  Next Hop : 10.1.2.2
-------~Full Output not provided ~-------
```

Step 2: Verifying RSVP Signaling

```
Router# show ip rsvp sender detail
PATH:
  Tun Dest:   10.255.255.1  Tun ID: 15  Ext Tun ID: 10.255.255.8
  Tun Sender: 10.255.255.8  LSP ID: 40
  Path refreshes:
    arriving: from PHOP 10.5.2.1 on Et0/1 every 30000 msecs. Timeout in 136 sec
    sent:     to   NHOP 10.1.4.1 on Ethernet0/0
  Session Attr:
    Setup Prio: 7, Holding Prio: 7
    Flags: (0x4) SE Style
    Session Name: R3_t15
  ERO: (incoming)
    10.5.2.2 (Strict IPv4 Prefix, 8 bytes, /32)
    10.1.4.2 (Strict IPv4 Prefix, 8 bytes, /32)
    10.1.4.1 (Strict IPv4 Prefix, 8 bytes, /32)
    10.255.255.1 (Strict IPv4 Prefix, 8 bytes, /32)
  ERO: (outgoing)
    10.1.4.1 (Strict IPv4 Prefix, 8 bytes, /32)
    10.255.255.1 (Strict IPv4 Prefix, 8 bytes, /32)
  ASSOCIATION:
    Extended Association type: Single sided provisioned bidirectional LSPs IPv4
    Association ID: 1, Source: 1.1.1.1
    Global source: 0
    ExtID[0]: 0xAFFFF08
    ExtID[1]: 0x28
-------~Full Output not provided ~-------
```

Step 3: Verifying RSVP Reservation

```
Router# show ip rsvp reservation detail
Reservation:
  Tun Dest:   10.255.255.1  Tun ID: 15  Ext Tun ID: 10.255.255.8
  Tun Sender: 10.255.255.8  LSP ID: 327
  Resv refreshes:
    arriving: from NHOP 10.1.4.1 on Et0/0 every 30000 msecs. Timeout in 382 sec
  Next Hop: 10.1.4.1 on Ethernet0/0
  Label: 23 (outgoing)
  Reservation Style is Shared-Explicit, QoS Service is Controlled-Load
  Resv ID handle: 1200040C.
  Created: 11:08:07 EST Fri Aug 28 2015
  Average Bitrate is 0 bits/sec, Maximum Burst is 1K bytes
  Min Policed Unit: 0 bytes, Max Pkt Size: 1500 bytes
  Status:
  Policy: Accepted. Policy source(s): MPLS/TE
```

```
Reservation:
  Tun Dest:   10.255.255.8  Tun ID: 15  Ext Tun ID: 10.255.255.1
  Tun Sender: 10.255.255.1  LSP ID: 338
  Resv refreshes:
    arriving: from NHOP 10.5.2.1 on Et0/1 every 30000 msecs. Timeout in 382 sec
  Next Hop: 10.5.2.1 on Ethernet0/1
  Label: 17 (outgoing)
  Reservation Style is Shared-Explicit, QoS Service is Controlled-Load
  Resv ID handle: 05000410.
  Created: 11:08:07 EST Fri Aug 28 2015
  Average Bitrate is 0 bits/sec, Maximum Burst is 1K bytes
  Min Policed Unit: 0 bytes, Max Pkt Size: 1500 bytes
  RRO:
    10.3.2.2/32, Flags:0x0 (No Local Protection)
    10.3.2.1/32, Flags:0x0 (No Local Protection)
  Status:
  Policy: Accepted. Policy source(s): MPLS/TE
```

Step 4: Verifying Wrap Functionality

```
Router# show mpls traffic-eng tunnels
P2P TUNNELS/LSPs:
Name: R1_t15                              (Tunnel15) Destination: 10.255.255.8
  Status:
    Admin: up         Oper: up     Path: valid       Signalling: connected
    path option 1, type explicit Primary (Basis for Setup, path weight 60)
    path option 2, type dynamic
    Path Protection: 0 Common Link(s), 0 Common Node(s)
    path protect option 1, type explicit Secondary (Basis for Protect, path weight 40)
    Lockout Info:
      Locked Out: No
  Config Parameters:
    Bandwidth: 0        kbps (Global)  Priority: 7  7   Affinity: 0x0/0xFFFF
    Metric Type: TE (default)
    AutoRoute: enabled  LockDown: disabled Loadshare: 0 [0] bw-based
    auto-bw: disabled
    Association Type: Single Sided Bidirectional LSPs, Co-routed: YES
    Association ID: 1, Source: 1.1.1.1
   Fault-OAM: enabled, Path-Protection: ready, Wrap-Protection: enabled, Wrap-Capable: Yes

    FlexLSP Event History:
  Active Path Option Parameters:
    State: explicit path option 1 is active
    BandwidthOverride: disabled  LockDown: disabled  Verbatim: disabled

Router# show mpls traffic-eng tunnels protection
P2P TUNNELS:
R1_t15
  LSP Head, Tunnel15, Admin: up, Oper: up
  Src 10.255.255.1, Dest 10.255.255.8, Instance 34
  Fast Reroute Protection: None
  Lockout Info:
    Locked Out: No
  Path Protection: Backup lsp in use.
  Prior Working LSP details:
    LSP ID: 33 (Delayed Clean)
    Deactivates In: (2796) ms
    InLabel  :  -
    OutLabel : Ethernet0/1, 16
    Next Hop : 10.1.4.2
    Reverse Associated LSP Information:
        Signaled Name: 10.255.255.8 15
        Tunnel: 15, Source: 10.255.255.8, Dest: 10.255.255.1, LSP: 29 State: Up
    Lockout Info:
        Locked out: No
```

```
        Lockout Originated By: None
    Association:
        Association Type: Single Sided Bidirectional LSPs
        Association ID: 1 Source: 1.1.1.1
-------~Full Output not provided ~-------
```

Step 5: Verifying BFD and OAM Operations

```
Router# show mpls traffic-eng tunnels detail | sec Fault
    Fault-OAM: enabled, Path-Protection: no protection, Wrap-Protection: disabled,
Wrap-Capable: No
    Fault-OAM Events:
      LSP 4638 (deleted) bfd-delete,
          at 07:32:08 IST Fri Jun 3 2016 (1 days, 8 hours, 35 mins, 30 secs ago)
      LSP 4638 (deleted) fault-delete,
          at 07:32:08 IST Fri Jun 3 2016 (1 days, 8 hours, 35 mins, 30 secs ago)
      LSP 4638 (working) bfd-up,
          at 10:15:31 IST Thu Jun 2 2016 (2 days, 5 hours, 52 mins, 7 secs ago)
      LSP 4637 (working) bfd-delete,
          at 10:15:20 IST Thu Jun 2 2016 (2 days, 5 hours, 52 mins, 18 secs ago)
      LSP 4637 (working) fault-delete,
          at 10:15:20 IST Thu Jun 2 2016 (2 days, 5 hours, 52 mins, 18 secs ago)
      LSP 4636 (working) bfd-delete,
          at 10:15:17 IST Thu Jun 2 2016 (2 days, 5 hours, 52 mins, 21 secs ago)
      LSP 4636 (working) fault-delete,
          at 10:15:17 IST Thu Jun 2 2016 (2 days, 5 hours, 52 mins, 21 secs ago)
-------~Full Output not provided ~-------
```

```
Router# show mpls fault-oam session end-point detail
  MPLS Fault-OAM End-point Sessions
====================================
    Session handle : 0x6
    Client handle : 0x2B9FAE02B750
    Local label : 18
    Tunnel interface : Tunnel3 (0x15)
    Tunnel number : 3
    LSP number : 49
    Global ID : 0
    Node ID : 10.1.0.1
    Local event : Fault Clear
    Sender Information
        Fault source : End-point
        Refresh seconds : 20
        Initial count : 0
        Fault type : CLR
        Tx Fault-CLR count : 0
        Tx Fault-AIS count : 0
        Tx Fault-LDI count : 0
        Tx Fault-LKR count : 0
        Tx Lockout-CLR count : 0
        Tx Lockout count : 0
        Tx Error count : 0
    Receiver Information
        Source global ID : 0
        Source node ID : 0
        Source intf number : 0
        Fault type : CLR
        Rx Fault-CLR count : 0
        Rx Fault-AIS count : 0
        Rx Fault-LDI count : 0
        Rx Fault-LKR count : 0
        Rx Lockout-CLR count : 0
        Rx Lockout count : 0
        Rx Error count : 0
    -------~Full Output not provided ~-------
```

Step 6: Verifying that Pseudowire is in UP State

```
Router# show mpls l2transport vc vcid 1            (HEAD router)

Local intf     Local circuit              Dest address    VC ID      Status
-------------  -------------------------- --------------- ---------- ----------
Gi6            Eth VLAN 30                53.0.0.1        1          UP
 #show mpls l2transport vc vcid 1 detail
Local interface: Gi6 up, line protocol up, Eth VLAN 30 up
  Interworking type is Ethernet
  Destination address: 53.0.0.1, VC ID: 1, VC status: up
    Output interface: Tu10, imposed label stack {29 29780}
    Preferred path: Tunnel10,  active
    Required BW = 15000, Admitted BW = 15000
    Default path: ready
    Next hop: point2point
  Create time: 00:01:13, last status change time: 00:01:13
    Last label FSM state change time: 00:01:13
  Signaling protocol: LDP, peer 53.0.0.1:0 up
    Targeted Hello: 52.0.0.1(LDP Id) -> 53.0.0.1, LDP is UP
    Graceful restart: configured and enabled
    Non stop routing: configured and not enabled

    -------~Full Output not provided ~-------
```

Use the **show adjacency tunnel internal** command to view the software forwarding of the tunnel:

```
Router# show adjacency tunnel1 internal | i lsp-num

          GigabitEthernet0/5/2 55.0.0.1 label 21 lsp-num 20
          Path protected by GigabitEthernet0/5/3 label 22 lsp-num 21
          Reopt of working: Null0 0.0.0.0 label none lsp-num 0
          Reopt of protect: Null0 label none lsp-num 0
```