MPLS Embedded Management—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 Embedded Management—LSP Ping/Traceroute and Any Transport over MPLS Virtual Circuit Connection Verification (AToM VCCV) feature helps them do this.

MPLS Embedded Management—LSP Ping/Traceroute and AToM VCCV 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.
- AToM VCCV allows you to use MPLS LSP Ping to test the Pseudo-Wire (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 Embedded Management—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.

MPLS Embedded Management—LSP Ping/Traceroute and AToM VCCV use 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* (draft-ietf-mpls-lsp-ping-03.txt).

Feature History for MPLS Embedded Management—LSP Ping/Traceroute and AToM VCCV

Release	Modification
12.0(27)S	This feature was introduced.
12.2(18)SXE	This feature was integrated into Cisco IOS Release 12.2(18)SXE.

Corporate Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

Note Software images for Cisco 12000 series Internet routers have been deferred to Cisco IOS Release 12.0(27)S1.

Finding Support Information for Platforms and Cisco IOS Software Images

Use Cisco Feature Navigator to find information about platform support and Cisco IOS software image support. Access Cisco Feature Navigator at http://www.cisco.com/go/fn. You must have an account on Cisco.com. If you do not have an account or have forgotten your username or password, click **Cancel** at the login dialog box and follow the instructions that appear.

Contents

- Prerequisites for MPLS Embedded Management—LSP Ping/Traceroute and AToM VCCV, page 2
- Restrictions for MPLS Embedded Management—LSP Ping/Traceroute and AToM VCCV, page 3
- Information About MPLS Embedded Management—LSP Ping/Traceroute and AToM VCCV, page 3
- Additional References, page 34
- Command Reference, page 36
- Glossary, page 49

Prerequisites for MPLS Embedded Management—LSP Ping/Traceroute and AToM VCCV

Before you use the MPLS Embedded Management—LSP Ping/Traceroute and AToM VCCV feature, you should:

- Determine the baseline behavior of your 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 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, AToM, and LDP. You need to
  - Know how LDP is configured
  - Understand AToM concepts
  - Be able to troubleshoot a TE tunnel
- Understand label switching, forwarding, and load balancing.

Restrictions for MPLS Embedded Management—LSP Ping/Traceroute and AToM VCCV

The following restrictions apply to the MPLS Embedded Management—LSP Ping/Traceroute and AToM VCCV feature:

- You cannot use MPLS LSP Traceroute to trace the path taken by 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/Traceroute to validate/trace MPLS Virtual Private Networks (VPNs).
- You cannot use MPLS LSP Traceroute to troubleshoot LSPs that employ Time to Live (TTL) hiding.

Information About MPLS Embedded Management—LSP Ping/Traceroute and AToM VCCV

Before using the MPLS Embedded Management—LSP Ping/Traceroute and AToM VCCV feature, you need an understanding of the following concepts:

- MPLS LSP Ping Operation, page 3
- MPLS LSP Traceroute Operation, page 5
- Any Transport over MPLS Virtual Circuit Connection Verification, page 7
- MPLS LSP Ping/Traceroute Command Options, page 9
- MPLS Echo Request Packets Not Forwarded by IP, page 16
- Information Provided by the Router Processing LSP Ping or LSP Traceroute, page 17
- MTU Discovery in an LSP, page 18
- Managing an LSP Network, page 19
- Troubleshooting with LSP Ping/Traceroute, page 20
- Load Balancing for IPv4 LDP LSPs, page 32

MPLS LSP Ping Operation

MPLS LSP Ping uses MPLS echo request and reply packets to validate an 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 router 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 router generating the echo reply. The destination address is the source address of the router in the MPLS echo request packet.

Figure 1 shows MPLS LSP Ping echo request and echo reply paths.

Figure 1 MPLS LSP Ping Echo Request and Echo Reply Paths

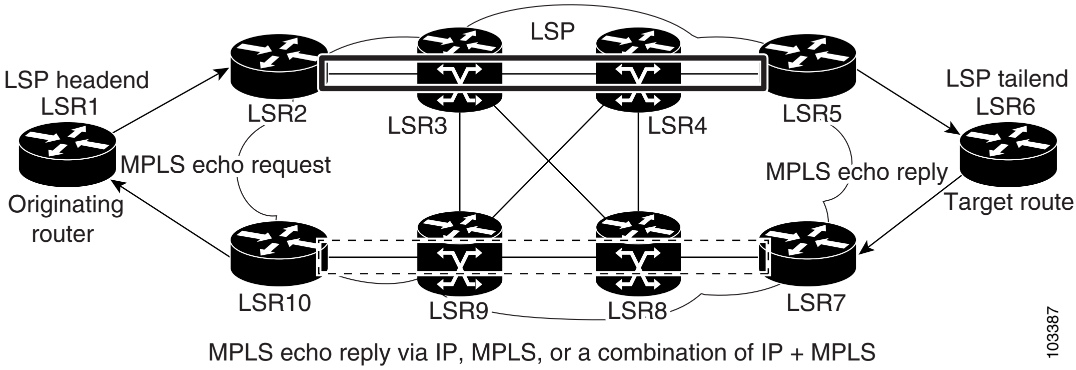

If you initiate an MPLS LSP Ping request at LSR1 to an FEC at LSR6, you get the results shown in Table 1.

Table 1 MPLS LSP Ping Example from Figure 1

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

You can use MPLS LSP Ping to validate IPv4 LDP, AToM, and IPv4 Resource Reservation Protocol (RSVP) FECs by using appropriate keywords and arguments with the **ping mpls** command:

```
ping mpls {ipv4 destination-address destination-mask | pseudowire ipv4-address vc-id vc-id
| traffic-eng tunnel-interface tunnel-number}
```

MPLS LSP Traceroute Operation

MPLS LSP Traceroute also uses MPLS echo request and reply packets to validate an LSP. The echo request and echo reply are 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 router processing the MPLS echo request when it receives a labeled packet with a TTL = 1. On Cisco routers, when the TTL expires, the packet is sent to the Route Processor (RP) for processing. The transit router returns an MPLS echo reply containing information about the transit hop in response to the TTL-expired MPLS packet.

Figure 2 shows an MPLS LSP Traceroute example with an LSP from LSR1 to LSR4.

Figure 2 MPLS LSP Traceroute Example

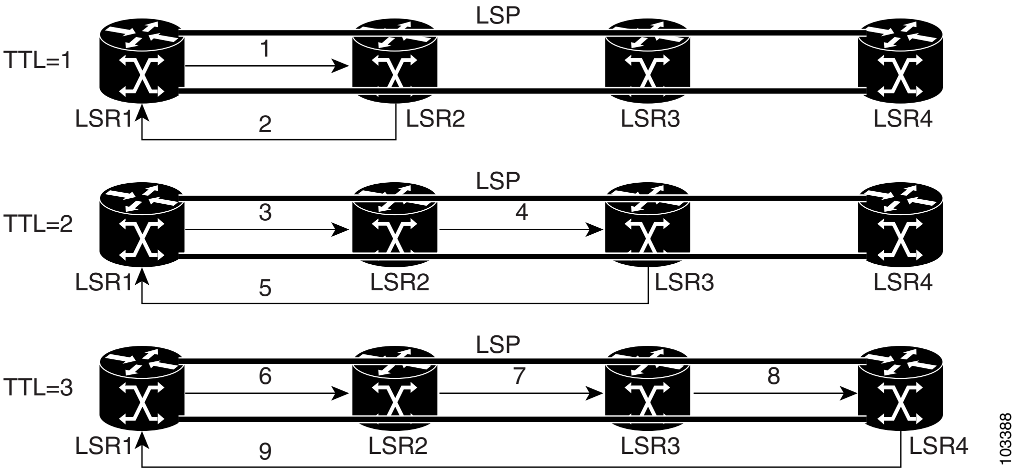

If you enter an LSP traceroute to a FEC at LSR4 from LSR1, you get the results shown in Table 2.

Table 2 MPLS LSP Traceroute Example Based on Figure 2

Step	Router	MPLS Packet Type and Description	Router 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
2.	LSR2	MPLS echo reply	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 and sends a reply
3.	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

Table 2 MPLS LSP Traceroute Example Based on Figure 2 (continued)

Step	Router	MPLS Packet Type and Description	Router Action
4.	LSR2	MPLS echo request	Receives packet with TTL = 2 • Decrements the TTL • Forwards the echo request to LSR3
5.	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
6.	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
7.	LSR2	MPLS echo request	Receives packet with TTL = 3 • Decrements the TTL • Forwards the echo request to LSR3
8.	LSR3	MPLS echo request	Receives packet with TTL = 2 • Decrements the TTL • Forwards the echo request to LSR4
9.	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 router is the egress router for the target FEC • Replies to LSR1

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

```
trace mpls {ipv4 destination-address destination-mask | traffic-eng tunnel-interface
tunnel-number}
```

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 router address of the last point that the trace reaches is repeated until the ouput is 30 lines. You can ignore the duplicate entries. The following example shows that the trace encountered an LSP problem at the router that has an IP address of 10.6.1.6:

```
Router# 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.

```
Router# 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 PW from the originating provider edge (PE) router. The transmission is intercepted at the destination PE router, instead of being forwarded to the customer edge (CE) router. 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 AToM VC setup is the signaling of VC labels and AToM VCCV capabilities between AToM VC endpoints. The router 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 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.

Table 3 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 routers always use Type 1 switching, if available, when they send MPLS LSP Ping packets over an AToM VC control channel. Type 2 switching accommodates those VC types and implementations that do not support or interpret the AToM control word.

Table 4 shows the AToM 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 router (PE1) to the destination router (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 router. For example:

```
PE1# show mpls l2transport binding

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

MPLS LSP Ping/Traceroute Command Options

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

The **ping mpls** command provides the following options:

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*] [**timeout** *seconds*][{**size** *packet-size*} | {**sweep** *minimum maximum size-increment*}] [**pad** *pattern*] [**reply mode** *reply-mode*] [**interval** *msec*] [**exp** *exp-bits*] [**verbose**]

The **trace mpls** command provides the following options:

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*] [**reply mode** *reply-mode*] [**ttl** *maximum-time-to-live*] [**exp** *exp-bits*]

The following sections describe some command options of the MPLS LSP Ping/Traceroute features:

- Selection of FECs for Validation, page 10
- Reply Mode Options for MPLS LSP Ping/Traceroute, page 10
- Other MPLS LSP Ping/Traceroute Command Options, page 12
- MPLS LSP Ping/Traceroute Option Interactions and Loops, page 14

Selection of FECs for Validation

An LSP is formed by labels. Routers learn labels through LDP, TE, AToM, or other MPLS applications. You can use MPLS LSP Ping/Traceroute to validate an LSP used for forwarding traffic for a given FEC. Table 5 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*	—[1]

[1] MPLS LSP Traceroute does not support the AToM tunnel LSP type for this release.

Reply Mode Options for MPLS LSP Ping/Traceroute

The reply mode is used to control how the responding router replies to an MPLS echo request sent by an MPLS LSP Ping or MPLS LSP Traceroute command. Table 6 describes the reply mode options.

Table 6 Reply Mode Options for a Responding Router

Option	Description
ipv4	Reply with an IPv4 UDP packet (default). This is the most common reply mode selected for use with an MPLS LSP Ping/Traceroute command when you want to periodically poll the integrity of an 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 router fails to receive a reply, select the router-alert option, "Reply with an IPv4 UDP packet with a router alert." The responding router sets the IP precedence of the reply packet to 6. You implement this option using the **reply mode ipv4** keywords.
router-alert	Reply with an IPv4 UDP packet with a router alert. This reply mode adds the router alert option to the IP header. This forces the packet to be special handled by the Cisco router 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

On Cisco routers, the reply with an IPv4 UDP packet implies that the router 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.

On Cisco routers, the reply with an IPv4 UDP packet that contains a router 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) routers fail to receive a reply to the MPLS echo request.

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

```
ping mpls {ipv4 destination-address destination-mask | pseudowire ipv4-address vc-id vc-id
| traffic-eng tunnel-interface tunnel-number} reply mode router-alert
```

or

```
trace mpls {ipv4 destination-address destination-mask | traffic-eng tunnel-interface
tunnel-number} reply mode router-alert
```

However, the reply with a router alert adds overhead to the process of getting a reply back to the originating router. 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) router 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 an MPLS packet with a router alert label as its outermost label arrives at a router, the router punts (redirects) the packet to the RP process level for handling. This allows these packets to bypass the forwarding failures in hardware routing tables. Table 7 describes how IP and MPLS packets with an IP router alert option are handled by the router switching path processes.

Table 7 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	Router alert option in 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
	Router alert option in IP 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
MPLS packet—Outermost label contains a router alert	If the router alert label is the outermost label, it causes the packet to be 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, it causes the packet to be 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/Traceroute Command Options

Table 8 describes other MPLS LSP Ping/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 for you to use only on the **ping mpls** command are indicated as such.

Table 8 Other MPLS LSP Ping/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 (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 ICMP ping sweep parameter. The lower boundary on the sweep range varies depending on the 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 **mpls ping** 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 MPLS echo response. Use the **source** *source-address* keyword and argument. For use with the MPLS LSP Ping and Traceroute features.
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 router. The label stack that is used to forward the echo request routes the MPLS packet to the destination router. The 127/8 address guarantees that the packets are routed to the localhost (the default loopback address of the router 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 AToM FECs with the MPLS LSP Ping feature and with IPv4 FECs with the MPLS LSP Traceroute feature.

Table 8 Other MPLS LSP Ping/Traceroute and AToM VCCV Options (continued)

Option	Description
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 TTL field in a packet is decremented by 1 each time it travels through a router. 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 router. 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 router. MPLS LSP Traceroute incrementally increases the TTL value in its MPLS echo requests (TTL = 1, 2, 3,4, ...) to accomplish this. Use the **ttl** *maximum-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* argument and keyword. For use with the MPLS LSP Ping and Traceroute features.
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 Table 8 can be implemented by the use of 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] [interval msec] [exp exp-bits] [verbose]
```

MPLS LSP Traceroute options described in Table 8 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]
```

MPLS LSP Ping/Traceroute Option Interactions and Loops

Usage examples for the MPLS Embedded Management—LSP Ping/Traceroute and AToM VCCV feature in this and subsequent sections are based on the sample topology shown in Figure 3.

Figure 3 Sample Topology for Configuration Examples

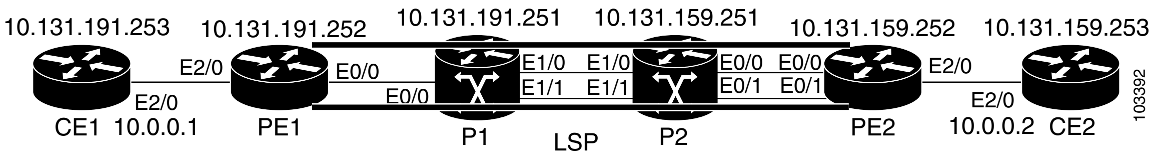

The interaction of some MPLS Embedded Management—LSP Ping/Traceroute and AToM VCCV options can cause loops. See the following topic for a description of the loops you might encounter with the **ping mpls** and **trace mpls** commands:

- Possible Loops with MPLS LSP Ping, page 14
- Possible Loop with MPLS LSP Traceroute, page 15

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 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:

```
Router# ping mpls ipv4 10.131.159.251/32 destination 127.0.0.1 127.0.0.2 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 127.0.0.2

!

!
```

```
Destination address 127.0.0.1
!
!

Destination address 127.0.0.2
!
!
```

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, 127.0.0.5, 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 UDP destination address range option and the time-to-live option.

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

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

```
Router# trace mpls ipv4 10.131.159.251/32 destination 127.0.0.1 127.0.0.3 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 127.0.0.3. 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

MPLS echo request packets sent during an 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. Routers 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 UDP packet is significant in that the MPLS echo request packet fails to make it to the target router if a transit router does not label switch the LSP. This allows for the detection of LSP breakages.

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

Figure 4 shows the path of the MPLS echo request and reply when a transit router fails to label switch a packet in an LSP.

Figure 4 Path When Transit Router Fails to Label Switch a Packet

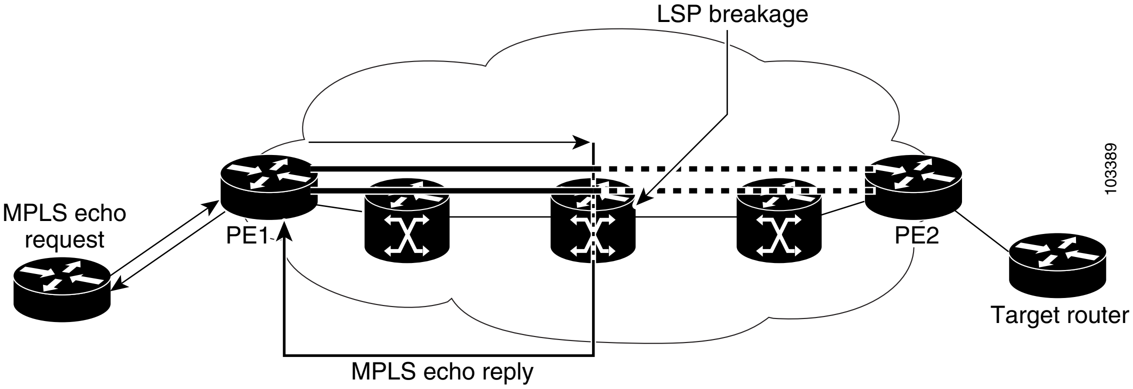

Note An AToM payload does not contain usable forwarding information at a transit router because the payload might not be an IP packet. An MPLS VPN packet, although an IP packet, does not contain usable forwarding information at a transit router because the destination IP address is only significant to the VRFs at the endpoints of the MPLS network.

Information Provided by the Router Processing LSP Ping or LSP Traceroute

Table 9 describes the characters that the router 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 **verbose** keyword on the **ping mpls** command.

Table 9 LSP Ping and Traceroute Reply Characters

Character	Meaning
Period “.”	A timeout occurs before the target router can reply.
U	The target router is unreachable.
R	The router processing the MPLS echo request is a downstream router but is not the destination.
Exclamation mark “!”	Replying router 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 LSP exists that matches the FEC information.
C	Replying router rejected the echo request because it was malformed.

MTU Discovery in an LSP

During an MLPS LSP Ping, 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 LSP without fragmentation.

Figure 5 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

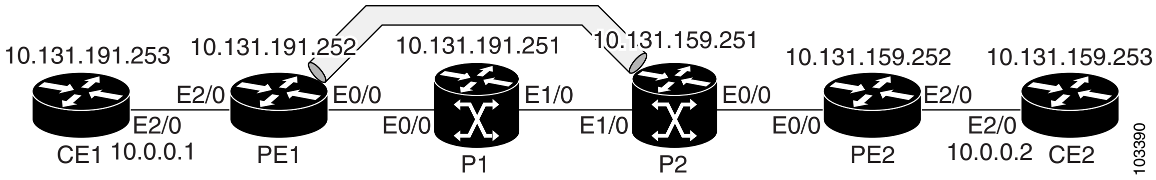

You can determine the maximum receive unit (MRU) at each hop by tracing the LSP using the MPLS Traceroute feature. The MRU is the maximum size of a labeled packet that can be forwarded through an LSP. The following example shows the results of a **trace mpls** command when the LSP is formed with labels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

```
PE1# 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:

```
PE1# 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** *interface-name* command.

```
PE1# show interface e0/0

Ethernet0/0 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 1/255
  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** *interface-name* 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 Figure 5, is 1500 – (2 x 4) = 1492.

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

```
PE1# 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.

Managing an LSP Network

To manage an 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 LDP prefixes, TE tunnels, and 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 router 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.

Troubleshooting with LSP Ping/Traceroute

ICMP **ping** and **trace** commands are often used to help diagnose the root cause of a failure. When an LSP is broken, the packet might make its way to the target router by way of IP forwarding, thus making ICMP ping and traceroute unreliable for detecting MPLS forwarding problems. The MPLS Embedded Management—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.

Figure 6 shows a sample topology with an LDP LSP and TE tunnel LSP.

Figure 6 Sample Topology with LDP and TE Tunnel LSPs

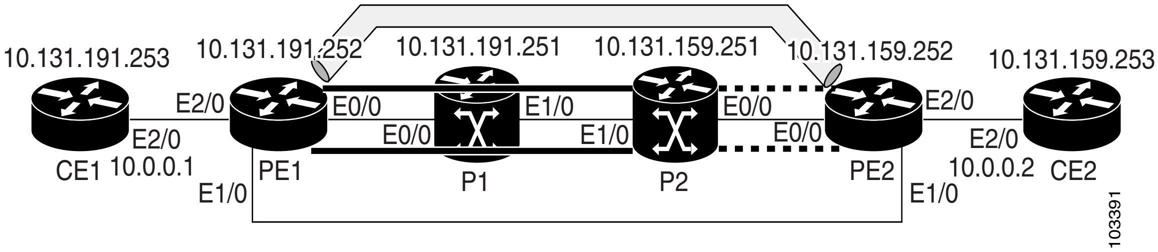

This section contains the following topics:

- MPLS LSP Ping/Traceroute Discovers LSP Breakage, page 20
- MPLS LSP Traceroute Tracks Untagged Cases, page 29
- MPLS LSP Ping/Traceroute Returns a Q, page 31

MPLS LSP Ping/Traceroute Discovers LSP Breakage

This section contains the following topics:

- Configuration for Sample Topology, page 21
- Verifying That the LSP Is Set Up Correctly, page 26
- Discovering LSP Breakage, page 27

Configuration for Sample Topology

These are sample topology configurations for the troubleshooting examples in the following sections (see Figure 6). There are the six sample router configurations.

Router CE1 Configuration

Following is the configuration for the CE1 router:

```
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 Ethernet2/0
 ip address 10.0.0.1 255.255.255.0
 no ip directed-broadcast
 no keepalive
 no cdp enable
!
end
```

Router PE1 Configuration

Following is the configuration for the PE1 router:

```
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 10.131.159.251
 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 10.131.159.252
 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 Ethernet0/0
 ip address 10.131.191.230 255.255.255.252
 no ip directed-broadcast
 mpls traffic-eng tunnels
 mpls ip
 ip rsvp bandwidth 1500 1500
 ip rsvp signalling dscp 0
!
interface Ethernet1/0
 ip address 10.131.159.246 255.255.255.252
 no ip directed-broadcast
 no shutdown
 mpls ip
 ip rsvp bandwidth 1500 1500
 ip rsvp signalling dscp 0

!
interface Ethernet2/0
 no ip address
 no ip directed-broadcast
 no cdp enable
 xconnect 10.131.159.252 333 encapsulation mpls
!
interface Ethernet3/0
 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
```

Router P1 Configuration

Following is the configuration for the P1 router:

```
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 Ethernet0/0
 ip address 10.131.191.229 255.255.255.252
 no ip directed-broadcast
 mpls traffic-eng tunnels
 mpls ip
 ip rsvp bandwidth 1500 1500
 ip rsvp signalling dscp 0
!
interface Ethernet1/0
 ip address 10.131.159.226 255.255.255.252
 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
```

Router P2 Configuration

Following is the configuration for the P2 router:

```
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 Ethernet0/0
 ip address 10.131.159.229 255.255.255.252
 no ip directed-broadcast
 mpls traffic-eng tunnels
 mpls ip
 ip rsvp bandwidth 1500 1500
 ip rsvp signalling dscp 0
!
interface Ethernet1/0
```

```
 ip address 10.131.159.225 255.255.255.252
 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
```

Router PE2 Configuration

Following is the configuration for the PE2 router:

```
version 12.0
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname as2_pe
!
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 Ethernet0/0
 ip address 10.131.159.230 255.255.255.252
 no ip directed-broadcast
 mpls traffic-eng tunnels
 tag-switching ip
 ip rsvp bandwidth 1500 1500
 ip rsvp signalling dscp 0
!
```

```
interface Ethernet1/0
 ip address 10.131.159.245 255.255.255.252
 no ip directed-broadcast
 mpls traffic-eng tunnels
 tag-switching ip
 ip rsvp bandwidth 1500 1500
 ip rsvp signalling dscp 0
!
interface Ethernet2/0
 no ip address
 no ip directed-broadcast
 no cdp enable
 xconnect 10.131.191.252 333 encapsulation mpls
!
interface Ethernet3/0
 no ip address
 no ip directed-broadcast
!
interface Serial4/0
 no ip address
 no ip directed-broadcast
 shutdown
!
interface Serial5/0
 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
```

Router CE2 Configuration

Following is the configuration for the CE2 router:

```
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 Ethernet2/0
 ip address 10.0.0.2 255.255.255.0
 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 MPLS forwarding table.

```
PE1# 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.

```
PE1# 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 : Ethernet0/0, 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.

```
PE1# 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

An LDP target-session is established between routers PE1 and P2, as shown in the output of the following **show mpls ldp discovery** command:

```
PE1# show mpls ldp discovery

 Local LDP Identifier:
    10.131.191.252:0
    Discovery Sources:
    Interfaces:
        Ethernet0/0 (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 router in global configuration mode:

```
P2(config)# no mpls ldp discovery targeted-hello accept
```

The LDP configuration change causes the targeted LDP session between the headend and tailend of the 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:

```
PE1# show mpls ldp discovery

 Local LDP Identifier:
    10.131.191.252:0
    Discovery Sources:
    Interfaces:
        Ethernet0/0 (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 router. The display shows that the outgoing packets are untagged as a result of the LDP configuration changes.

```
PE1# 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 router displays the following:

```
PE1# 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 MPLS echo reply had a routing entry but no MPLS FEC. Entering the **verbose** keyword to the **ping mpls** command displays the MPLS LSP echo reply sender address and the return code. You should be able to solve the problem by Telneting to the replying router and inspecting its forwarding and label tables. You might need to look at the neighboring upstream router as well, because the breakage might be on the upstream router.

```
PE1# 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 router caused the breakage. In the following example, for subsequent values of TTL greater than 2, the same router keeps responding (10.131.159.225). This suggests that the MPLS echo request keeps getting processed by the router 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.

```
PE1# 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.

- Troubleshooting Implicit Null Cases, page 30
- Troubleshooting Untagged Cases, page 30

Untagged output interfaces at a penultimate hop do not impact the forwarding of IP packets through an LSP because the forwarding decision is made at the penultimate hop through use of the incoming label. The untagged case causes AToM and MPLS VPN traffic to be dropped at the penultimate hop.

Troubleshooting Implicit Null Cases

In the following example, Tunnel 1 is shut down, and only an LSP formed with LDP labels is established. An implicit null is advertised between the P2 and PE2 routers. Entering an MPLS LSP Traceroute at the PE1 router results in the following display:

```
PE1# 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 Ethernet 0/0 out interface for the PE2 router.

Troubleshooting Untagged Cases

Untagged cases are valid configurations for IGP LSPs that could cause problems for MPLS VPNs.

A **show mpls forwarding-table** command and a **show mpls ldp discovery** command issued at the P2 router show that LDP is properly set up:

```
P2# show mpls for 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

P2# show mpls ldp discovery

 Local LDP Identifier:
    10.131.159.251:0
    Discovery Sources:
    Interfaces:
        Ethernet0/0 (ldp): xmit/recv
            LDP Id: 10.131.159.252:0
        Ethernet1/0 (ldp): xmit/recv
            LDP Id: 10.131.191.251:0
```

The **show mpls ldp discovery** command output shows that Ethernet0/0, which connects PE2 to P2, is sending and receiving packets.

If a **no mpls ip** command is entered on Ethernet 0/0, this could prevent an LDP session between the P2 and PE2 routers from being established. A **show mpls ldp discovery** command entered on the PE router shows that the MPLS LDP session with the PE2 router is down:

```
P2# show mpls ldp discovery

 Local LDP Identifier:
    10.131.159.251:0
    Discovery Sources:
```

```
    Interfaces:
        Ethernet0/0 (ldp): xmit
        Ethernet1/0 (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:

```
P2# 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:

```
PE1# 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/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 an LSP could not be found that matches the 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.

Table 10 lists commands you can use for troubleshooting when the MPLS echo reply returns a Q.

Table 10 Troubleshooting a Q

Database	Command to View Contents
Routing Information Base	**show ip route**
Label information base/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:

```
PE1# 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 **show ip route** command demonstrate that the address is not in either routing table:

```
PE1# 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

PE1# 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 ICMP ping or trace follows one path from the originating router to the target router. Round robin load balancing of IP packets from a source router is used to discover the various output paths to the target IP address.

For MPLS LSP Ping/Traceroute, Cisco routers use the source and destination addresses in the IP header for load balancing when multiple paths exist through the network to a target router. 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 ipv4** *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 router and the target router.

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

```
PE1# configure terminal

Enter configuration commands, one per line.  End with CNTL/Z.

PE1(config)# interface ethernet 1/0

PE1(config-if)# no shut

PE1(config-if)# end
```

```
*Dec 31 19:14:10.034: %LINK-3-UPDOWN: Interface Ethernet1/0, changed state to up
*Dec 31 19:14:11.054: %LINEPROTO-5-UPDOWN: Line protocol on Interface Ethernet1/0,
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 Ethernet1/0
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:

```
PE1# 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          Et0/0      10.131.191.229
       20          10.131.159.251/32 0          Et1/0      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:

```
Router# 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 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 127.0.0.3 shows that the path selected has a path index of 1:

```
PE1# ping mpls ipv4 10.131.159.251/32 dest 127.0.0.3 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 20:43:09.518: 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 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.

Additional References

The following sections provide references related to the MPLS Embedded Management—LSP Ping/Traceroute and AToM VCCV feature.

- Related Documents, page 34
- Standards, page 35
- MIBs, page 35
- RFCs, page 35
- Technical Assistance, page 35

Related Documents

Related Topic	Document Title
Usage examples for the IP ping and IP traceroute commands	*Cisco—Understanding the Ping and Traceroute Commands*
Usage examples for the extended ping and extended traceroute commands	*Cisco—Using the Extended Ping and Traceroute Commands*
Configuration and verification tasks for MPLS LDP	*MPLS Label Distribution Protocol (LDP)*

Related Topic	Document Title
Configuration task for MPLS traffic engineering	*Multiprotocol Label Switching (MPLS) Traffic Engineering*
Configuration and verification tasks for AToM	*Any Transport over MPLS (AToM)*
Troubleshooting procedures for MPLS	*Cisco—MPLS Troubleshooting*
Switching services commands	*Cisco IOS Switching Services Command Reference, Release 12.3*
Configuration and verification tasks for MPLS applications	Part 3: Multiprotocol Label Switching, *Cisco IOS Switching Services Configuration Guide, Release 12.3*

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
No new or modified MIBs are supported by this feature, and support for existing MIBs has not been modified by this feature.	To locate and download MIBs for selected platforms, Cisco IOS releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

RFCs

RFCs	Title
draft-ietf-mpls-lsp-ping-03.txt	*Detecting MPLS Data Plane Failures*
RFC 2113	*IP Router Alert Option*
draft-ietf-pwe3-vccv-01.txt	*Pseudo-Wire (PW) Virtual Circuit Connection Verification (VCCV)*

Technical Assistance

Description	Link
Technical Assistance Center (TAC) home page, containing 30,000 pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/public/support/tac/home.shtml

Command Reference

This section documents only new and modified commands.

- **debug mpls lspv**
- **ping mpls**
- **trace mpls**

debug mpls lspv

To display information related to the Multiprotocol Label Switching (MPLS) Label Switched Path (LSP) Ping/Traceroute feature, use the **debug mpls lspv** command in privileged EXEC mode. To disable debugging output, use the **no** form of this command.

debug mpls lspv [**tlv**] [**error**] [**event**] [**packet** [**data** | **error**]]

no debug mpls lspv

Syntax Description

tlv	(Optional) Displays MPLS echo packet type-length-value (TLV) information as it is being coded and decoded.
error	(Optional) Displays error conditions encountered during MPLS echo request and echo reply encoding and decoding. See Table 11.
event	(Optional) Displays MPLS echo request and reply send and receive event information.
packet data	(Optional) Displays detailed debug information for the MPLS echo packets sent and received. This output is seen only on the originating router and the router generating the reply.
packet error	(Optional) Displays packet errors for MPLS echo request and reply. No output is expected for this command at this time.

Defaults

MPLS LSP debugging is disabled.

Command Modes

Privileged EXEC

Command History

Release	Modification
12.0(27)S	This command was introduced.

Usage Guidelines

Use this command to monitor activity associated with the **ping mpls** and the **trace mpls** commands.

Table 11 lists the messages displayed by the **debug mpls lspv error** command and the reason for the error message.

Table 11 Messages Displayed by the debug mpls lspv error Command

Message	Reason
Echo reply discarded because not routable	If an echo reply message is sent because the IP header indicates that the packet has the Router Alert set and the packet is not routable, this message is displayed.
UDP checksum error, packet discarded	If a packet is received on the port being used by Label Switched Path Verification (LSPV) and there is a checksum error on the packet, this message is displayed.

Table 11 Messages Displayed by the debug mpls lspv error Command (continued)

Message	Reason
Invalid echo message type	If an MPLS echo packet with an invalid echo message type (neither a request nor a reply) is received, this message is displayed.
Illegal Action	If the state machine that drives the LSPV software detects an invalid condition, this message is displayed.

Examples

The following example shows the syntax for the **debug mpls lspv** command:

```
PE1_Router# debug mpls lspv ?

  error   error debugging
  event   event debugging
  packet  packet debug options
  tlv     TLV debugging
```

The following example shows sample output for the **ping mpls** command when LSPV event debugging is enabled:

```
PE1_Router# debug mpls lspv event

LSPV event debugging is on
PE1_Router#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.
!
Success rate is 100 percent (1/1), round-trip min/avg/max = 48/48/48 ms
PE1_Router#
*Dec 31 19:31:15.366: LSPV:
waiting for 2 seconds
*Dec 31 19:31:15.366: LSPV: sender_handle: 2000002D, Event Echo Requests Start,
[Idle->Waiting for Echo Reply]
*Dec 31 19:31:15.414: LSPV: sender_handle: 2000002D, Event Echo Reply Received,
[Waiting for Echo Reply->Waiting for Interval]
*Dec 31 19:31:15.466: LSPV: sender_handle: 2000002D, Event Echo Requests Cancel,
[Waiting for Interval->Idle]

PE1_Router# undebug all
All possible debugging has been turned off
```

The following example shows sample output for the **ping mpls** command when LSPV TLV debugging is enabled:

```
PE1_Router# debug mpls lspv tlv

LSPV tlv debugging is on

PE1_Router# 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.
!
Success rate is 100 percent (1/1), round-trip min/avg/max = 40/40/40 ms
PE1_Router#
*Dec 31 19:32:32.566: LSPV: Echo Hdr encode: version 1, msg type 1, reply mode 2
, return_code 0, return_subcode 0, sender handle 9400002E, sequence number 1,
timestamp sent 14:32:32 EST Wed Dec 31 2003, timestamp rcvd 19:00:00 EST Thu Dec 31 1899
*Dec 31 19:32:32.566: LSPV: IPV4 FEC encode: destaddr 10.131.159.252/32
*Dec 31 19:32:32.566: LSPV: Pad TLV encode: type 1, size 18, pattern 0xABCD
*Dec 31 19:32:32.606: LSPV: Echo Hdr decode: version 1, msg type 2, reply mode 2,
return_code 3, return_subcode 0, sender handle 9400002E, sequence number 1,
timestamp sent 14:32:32 EST Wed Dec 31 2003, timestamp rcvd 14:32:32 EST Wed Dec 31 2003

PE1_Router# undebug all
All possible debugging has been turned off
```

Related Commands

Command	Description
ping mpls	Checks MPLS LSP connectivity.
trace mpls	Discovers MPLS LSP routes that packets will actually take when traveling to their destinations.

ping mpls

To check Multiprotocol Label Switching (MPLS) label switched path (LSP) connectivity, use the **ping mpls** command in privileged EXEC mode.

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*] [**timeout** *seconds*] [{**size** *packet-size*} | {**sweep** *minimum maximum size-increment*}] [**pad** *pattern*] [**reply mode** *reply-mode*] [**interval** *msec*] [**exp** *exp-bits*] [**verbose**]

Syntax Description

ipv4	Specifies the destination type as a Label Distribution Protocol (LDP) IPv4 address.
destination-address	Address prefix of the target to be tested.
destination-mask	Number of bits in the network mask of the target address.
destination	(Optional) Specifies a network 127 address.
address-start	(Optional) Beginning network 127 address.
address-end	(Optional) Ending network 127 address.
increment	(Optional) Number by which to increment the network 127 address.
ttl *time-to-live*	(Optional) Specifies a time-to-live (TTL) value.
pseudowire	Specifies the destination type as an Any Transport over MPLS (AToM) virtual circuit (VC).
ipv4-address	IPv4 address of the AToM VC to be tested.
vc-id *vc-id*	Specifies the VC identifier of the AToM VC to be tested.
traffic-eng	Specifies the destination type as an MPLS traffic engineering (TE) tunnel.
tunnel-interface	Tunnel interface to be tested.
tunnel-number	Tunnel interface number.
source	(Optional) Specifies the source address or name.
source-address	(Optional) Source address or name.
repeat *count*	(Optional) Specifies the number of times to resend the same packet. The range is from 1 to 2147483647. The default is 5.
timeout *seconds*	(Optional) Specifies the timeout interval in seconds. The range is from 0 to 3600. The default is 2 seconds.
size *packet-size*	(Optional) Specifies the packet size. Packet size is the number of bytes in each ping. The range is from 40 to 18024.
sweep	(Optional) Specifies a range of sizes for the echo packets sent.
minimum	(Optional) Minimum or start size for an MPLS echo packet. The lower boundary of the **sweep** range varies depending on the LSP type.
maximum	(Optional) Maximum or end size for an echo packet.
size-increment	(Optional) Number by which to increment the echo packet size.
pad *pattern*	(Optional) Pad pattern for MPLS echo request so that a payload of a given size can be tested.

reply mode *reply-mode*	(Optional) Specifies the reply mode for the echo request packet. The *reply-mode* is one of the following: **ipv4** = Reply with an IPv4 UDP packet (default) **router-alert** = Reply with an IPv4 UDP packet with router alert
interval *msec*	(Optional) Specifies a send interval between requests in milliseconds. Default is 0.
exp *exp-bits*	(Optional) Specifies the MPLS experimental field value in the MPLS header for an MPLS echo reply. Valid values are from 0 to 7. Default is 0.
verbose	(Optional) Displays the MPLS echo reply sender address of the packet and displays return codes.

Defaults

repeat count = 1
timeout = 2 seconds
packet size = 100 bytes
sweep minimum = 100 bytes
sweep maximum = 17986 bytes
sweep size increment = 100 bytes
pad TLV pattern = 0xABCD
reply mode = ipv4 via UDP (2)
time-to-live = 255 seconds
send interval = 0 msec
exp bits in MPLS header = 0
verbose = no

Command Modes

Privileged EXEC

Command History

Release	Modification
12.0(27)S	This command was introduced.

Usage Guidelines

Use the **mpls ping** command to validate, test, or troubleshoot IPv4 LDP LSPs, IPv4 RSVP TE tunnels, and AToM VCs.

Examples

The following example shows how to use the **ping mpls** command to test connectivity of an IPv4 LDP LSP:

```
PE# ping mpls ipv4 10.131.191.252/32 exp 5 repeat 5 verbose

Sending 5, 100-byte MPLS Echos to 10.131.191.252, timeout is 2 seconds:

Codes: '!' - success, 'Q' - request not transmitted,
       '.' - timeout, 'U' - unreachable,
       'R' - downstream router but not destination

Type escape sequence to abort.
!     10.131.191.230, return code 3
!     10.131.191.230, return code 3
!     10.131.191.230, return code 3
!     10.131.191.230, return code 3
!     10.131.191.230, return code 3

Success rate is 100 percent (5/5), round-trip min/avg/max = 100/102/112 ms
```

The following example shows how to invoke the **ping mpls** command in the interactive mode:

```
PE1# ping

Protocol [ip]: mpls
Target IPv4, pseudowire or traffic-eng [ipv4]: ipv4
Target IPv4 address: 10.131.159.252
Target mask: 255.255.255.255
Repeat count [5]: 1
Datagram size [100]:
Timeout in seconds [2]:
Send interval in msec [0]:
Extended commands? [no]: yes
Destination address or destination start address: 127.0.0.1
Destination end address: 127.0.0.3
Destination address increment: 0.0.0.1
Source address:
EXP bits in mpls header [0]:
Pad TLV pattern [ABCD]:
Time To Live [255]:
Reply mode ( 2-ipv4 via udp, 3-ipv4 via udp with router alert) [2]:
Verbose mode? [no]: yes
Sweep range of sizes? [no]:
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.
Destination address 127.0.0.1
!   10.131.159.245, return code 3


Destination address 127.0.0.2
!   10.131.159.245, return code 3


Destination address 127.0.0.3
!   10.131.159.245, return code 3


Success rate is 100 percent (3/3), round-trip min/avg/max = 40/48/52 ms
```

Note The “Destination end address” and “Destination address increment” prompts display only if you enter an address at the “Destination address or destination start address” prompt. Also, the “Sweep min size,” “Sweep max size,” and “Sweep interval” prompts display only if you enter **yes** at the “Sweep range of size? [no]” prompt.

The following example shows how to use the **ping mpls** command to test connectivity of an AToM VC:

```
PE# show mpls l2transport vc

Local intf     Local circuit          Dest address    VC ID      Status
-------------  ---------------------- --------------- ---------- ----------
Et2/0          Ethernet               10.131.191.252  333        UP
as2_pe#sh mpls l2transport vc det
Local interface: Et2/0 up, line protocol up, Ethernet up
  Destination address: 10.131.191.252, VC ID: 333, VC status: up
    Preferred path: not configured
    Default path: active
    Tunnel label: imp-null, next hop 10.131.159.246
    Output interface: Et1/0, imposed label stack {16}
  Create time: 06:46:08, last status change time: 06:45:51
  Signaling protocol: LDP, peer 10.131.191.252:0 up
    MPLS VC labels: local 16, remote 16
    Group ID: local 0, remote 0
    MTU: local 1500, remote 1500
    Remote interface description:
  Sequencing: receive disabled, send disabled
  VC statistics:
    packet totals: receive 0, send 0
    byte totals:   receive 0, send 0
    packet drops:  receive 0, send 0
```

MPLS LSP Ping can be used for Pseudo-Wires (PWs) as follows:

```
PE# ping mpls pseudowire 10.131.191.252 333

Sending 1, 100-byte MPLS Echos to 10.131.191.252, timeout is 2 seconds:

Codes: '!' - success, 'Q' - request not transmitted,
       '.' - timeout, 'U' - unreachable,
       'R' - downstream router but not destination

Type escape sequence to abort.
!
Success rate is 100 percent (1/1), round-trip min/avg/max = 92/92/92 ms
```

This ping is particularly useful because the VC might be up and the LDP session between the PE and its downstream neighbor might be also up, but LDP might be broken somewhere in between. In such cases an LSP ping helps verify that the LSP is actually up.

A related point concerns the situation when a PW has been configured to use a specific TE tunnel. For example:

```
PE(config-if)# do show running-config interface ethernet 2/0

Building configuration...

Current configuration : 129 bytes
!
interface Ethernet2/0
 no ip address
 no ip directed-broadcast
 no cdp enable
 xconnect 10.131.191.252 333 pw-class test1
end
```

```
PE# show running-config | begin pseudowire

pseudowire-class test1
 encapsulation mpls
 preferred-path interface Tunnel0
!
```

In such cases, you can use MPLS LSP Ping to simply verify the connectivity of the LSP that a certain PW is taking, be it LDP based or a TE tunnel:

```
PE# ping mpls pseudowire 10.131.191.252 333 repeat 200 size 1400

Sending 200, 1400-byte MPLS Echos to 10.131.191.252, timeout is 2 seconds:

Codes: '!' - success, 'Q' - request not transmitted,
       '.' - timeout, 'U' - unreachable,
       'R' - downstream router but not destination

Type escape sequence to abort.
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Success rate is 100 percent (200/200), round-trip min/avg/max = 72/85/112 ms
```

You can also use MPLS LSP Ping to verify the maximum packet size that can successfully be transmitted. The following command uses a packet size of 1500 bytes:

```
PE# ping mpls pseudowire 10.131.191.252 333 repeat 5 size 1500

Sending 5, 1500-byte MPLS Echos to 10.131.191.252, timeout is 2 seconds:

Codes: '!' - success, 'Q' - request not transmitted,
       '.' - timeout, 'U' - unreachable,
       'R' - downstream router but not destination

Type escape sequence to abort.
QQQQQ
Success rate is 0 percent (0/5)
```

The Qs indicate that the packets are not transmitted.

The following command uses a packet size of 1476 bytes:

```
PE# ping mpls pseudowire 10.131.191.252 333 repeat 5 size 1476

Sending 5, 1476-byte MPLS Echos to 10.131.191.252, timeout is 2 seconds:

Codes: '!' - success, 'Q' - request not transmitted,
       '.' - timeout, 'U' - unreachable,
       'R' - downstream router but not destination

Type escape sequence to abort.
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 80/83/92 ms
```

The following example shows how to test the connectivity of an MPLS TE tunnel:

```
pe# ping mpls traffic-eng tunnel 3 repeat 5 verbose

Sending 5, 100-byte MPLS Echos to Tunnel3,
     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.
!   10.131.159.198, return code 3
!   10.131.159.198, return code 3
!   10.131.159.198, return code 3
!   10.131.159.198, return code 3
!   10.131.159.198, return code 3

Success rate is 100 percent (5/5), round-trip min/avg/max = 32/37/40 ms
```

The MPLS LSP Ping feature again would be very useful if you want to verify TE tunnels before actually mapping traffic onto them.

Related Commands

Command	Description
trace mpls	Discovers MPLS LSP routes that packets will actually take when traveling to their destinations.

trace mpls

To discover Multiprotocol Label Switching (MPLS) label switched path (LSP) routes that packets actually take when traveling to their destinations, use the **trace mpls** command in privileged EXEC mode.

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*] [**reply mode** *reply-mode*] [**ttl** *maximum-time-to-live*] [**exp** *exp-bits*]

Syntax Description

ipv4	Specifies the destination type as a Label Distribution Protocol (LDP) IPv4 address.
destination-address	Address prefix of the target to be tested.
destination-mask	Number of bits in the network mask of the target address.
destination	(Optional) Specifies a network 127 address.
address-start	(Optional) The beginning network 127 address.
address-end	(Optional) The ending network 127 address.
address-increment	(Optional) Number by which to increment the network 127 address.
traffic-eng	Specifies the destination type as an MPLS traffic engineering (TE) tunnel.
tunnel-interface	Tunnel interface to be tested.
tunnel-number	Tunnel interface number.
source	(Optional) Specifies the source address or name.
source-address	(Optional) Source address or name.
timeout *seconds*	(Optional) Specifies the timeout interval in seconds. The range is from 0 to 3600. The default is 2 seconds.
reply mode *reply-mode*	(Optional) Specifies the reply mode for the echo request packet. The *reply-mode* is one of the following: **ipv4** = Reply with an IPv4 UDP packet (default) **router-aler**t = Reply with an IPv4 UDP packet with router alert
ttl *maximum-time-to-live*	(Optional) Specifies a maximum hop count.
exp *exp-bits*	(Optional) Specifies the MPLS experimental field value in the MPLS header for an MPLS echo reply. Valid values are from 0 to 7. Default is 0.

Defaults

timeout = 2 seconds
maximum time-to-live = 30
exp bits in MPLS header = 0
reply mode = ipv4 via UDP (2)

Command Modes

Privileged EXEC

Command History

Release	Modification
12.0(27)S	This command was introduced.

Usage Guidelines

Use the **trace mpls** command to validate, test, or troubleshoot IPv4 LDP LSPs and IPv4 RSVP TE tunnels.

Examples

The following example shows how to trace packets through an MPLS LDP LSP:

```
pe# trace mpls ipv4 10.131.191.252/32
```

Or, alternatively, using the interactive mode:

```
Protocol [ip]: mpls
Target IPv4, pseudowire or traffic-eng [ipv4]: <ipv4 |pseudowire |tunnel> ipv4
Target IPv4 address: 10.131.191.252
Target mask: /32
Repeat [1]:
Packet size [100]:
Timeout in seconds [2]:
Extended commands? [no]: yes
Destination start address:
Destination end address:
Source address:
EXP bits in mpls header [0]:
TimeToLive [255]:
Reply mode (2-ipv4 via udp, 3-ipv4 via udp with router alert) [2]:
Reply ip header DSCP bits [0]:

Tracing MPLS Label Switched Path to 10.131.191.252/32, timeout is 2 seconds

Codes: '!' - success, 'Q' - request not transmitted,
       '.' - timeout, 'U' - unreachable,
       'R' - downstream router but not destination

Type escape sequence to abort.

  0 10.131.159.245 mtu 1500 []
! 1 10.131.191.252 100 ms
```

The following example shows how to trace packets through an MPLS TE tunnel:

```
PE# trace mpls tunnel ?

  Tunnel  Tunnel interface
```

Or, alternatively, using the interactive mode:

```
pe# traceroute
Protocol [ip]: mpls
Target IPv4 or tunnel [ipv4]: traffic-eng
Tunnel number [0]:
Repeat [1]:
Timeout in seconds [2]:
Extended commands? [no]:
Tracing MPLS TE Label Switched Path on Tunnel0, timeout is 2 seconds

Codes: '!' - success, 'Q' - request not transmitted,
       '.' - timeout, 'U' - unreachable,
       'R' - downstream router but not destination
```

```
Type escape sequence to abort.

  0 10.131.159.230 mtu 1500 [Labels: 22 Exp: 0]
R 1 10.131.159.225 mtu 1500 [Labels: 22 Exp: 6] 72 ms
R 2 10.131.191.229 mtu 1504 [implicit-null] 72 ms
! 3 10.131.191.252 92 ms
```

Tunnel 0 is configured as follows:

```
PE# show running-config interface tunnel 0

Building configuration...

Current configuration : 210 bytes
!
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
end

PE# show mpls traffic-eng tunnels tunnel 0 brief

Signalling Summary:
    LSP Tunnels Process:            running
    RSVP Process:                   running
    Forwarding:                     enabled
    Periodic reoptimization:        every 3600 seconds, next in 1369 seconds
    Periodic FRR Promotion:         Not Running
    Periodic auto-bw collection:    disabled
TUNNEL NAME                      DESTINATION      UP IF     DOWN IF   STATE/PROT
PE_t0                        10.131.191.252   -         Et0/0     up/up

PE# show ip cef 10.131.191.252

10.131.191.252/32, version 37, epoch 0, cached adjacency 10.131.159.246
0 packets, 0 bytes
  tag information set, all rewrites owned
    local tag: 21
  via 10.131.159.246, Ethernet1/0, 0 dependencies
    next hop 10.131.159.246, Ethernet1/0
    valid cached adjacency
    tag rewrite with Et1/0, 10.131.159.246, tags imposed {}
```

The tunnel destination is the same IP address as the one in the earlier trace IPv4 example, except that the trace takes a different path, even though tunnel 0 is not configured to forward traffic by means of auto-route or static routing. This makes the **trace mpls traffic-eng** command very powerful, because you can now test the tunnels to make sure that they actually work before you map traffic onto them.

Related Commands

Command	Description
ping mpls	Checks MPLS LSP connectivity.

Glossary

FEC—Forwarding Equivalence Class. A set of packets that can be handled equivalently for forwarding purposes and are thus suitable for binding to a single label. Examples include the set of packets destined for one address prefix and any flow.

flow—Generally, a set of packets traveling between a pair of hosts, or a pair of transport protocol ports on a pair of hosts. For example, packets with the same source address, source port, destination address, and destination port might be considered a flow.

A flow is also a stream of data traveling between two endpoints across a network (for example, from one LAN station to another). Multiple flows can be transmitted on a single circuit.

fragmentation—Process of breaking a packet into smaller units when transmitting over a network medium that cannot support the original size of the packet.

ICMP— Internet Control Message Protocol. A network layer Internet protocol that reports errors and provides other information relevant to IP packet processing. It is documented in RFC 792.

LFIB—label forwarding information base. A data structure and way of managing forwarding in which destinations and incoming labels are associated with outgoing interfaces and labels.

localhost—A name that represents the host name of a router (device). The localhost uses the reserved loopback IP address 127.0.0.1.

LSP—label switched path. A connection between two routers that uses MPLS to carry the packets.

LSPV—Label Switched Path Verification. An LSP Ping subprocess that encodes and decodes MPLS echo requests and replies; interfaces with IP, MPLS, and AToM switching for sending and receiving MPLS echo requests and replies; and, at the MPLS echo request originator router, maintains a database of outstanding echo requests for which echo responses have not been received.

MPLS router alert label—An MPLS label of 1. An MPLS packet with a router alert label is redirected by the router to the Route Processor (PR) processing level for handling. This allows these packets to bypass any forwarding failures in hardware routing tables.

MRU—maximum receive unit. Maximum size, in bytes, of a labeled packet that can be forwarded through an LSP.

MTU—maximum transmission unit. Maximum packet size, in bytes, that a particular interface can handle.

punt—Redirect packets with a router alert from the line card or interface to Route Processor (RP) level processing for handling.

PW—Pseudo-Wire. A mechanism that carries the essential elements of an emulated circuit from one provider edge (PE) router to another PE router over a packet-switched network.

RP—Route Processor. Processor module in the Cisco 7000 series routers that contains the CPU, system software, and most of the memory components that are used in the router. It is sometimes called a supervisory processor.

RSVP—Resource Reservation Protocol. A protocol that supports the reservation of resources across an IP network. Applications running on IP end systems can use RSVP to indicate to other nodes the nature (bandwidth, jitter, maximum burst, and so on) of the packet streams they want to receive. RSVP depends on IPv6. Is is also known as Resource Reservation Setup Protocol.

UDP—User Datagram Protocol. Connectionless transport layer protocol in the TCP/IP protocol stack. UDP is a simple protocol that exchanges datagrams without acknowledgments or guaranteed delivery, requiring that error processing and retransmission be handled by other protocols. UDP is defined in RFC 768.

Note Refer to the Cisco Systems *Dictionary of Internetworking Terms and Acronyms* for terms not included in this glossary.