



RSVP-TE Transport Solution with Flex LSP

Table 1: Feature History Table

Feature Name	Release Information	Feature Description
RSVP-TE Transport Solution with Flex LSP	Release 24.4.1	<p>Introduced in this release on: Fixed Systems(8200, 8700[ASIC:K100]); Modular Systems (8800 [LC ASIC: P100]) (select variants only*).</p> <p>*This feature is now supported on:</p> <ul style="list-style-type: none">• 8212-32FH-M• 8711-32FH-M• 88-LC1-12TH24FH-E• 8712-MOD-M
RSVP-TE Transport Solution with Flex LSP	Release 7.11.1	<p>You can configure bidirectional label-switched paths dynamically through Resource Reservation Protocol-Traffic Engineering (RSVP-TE) using flex LSP. With this configuration, you can dynamically alter settings for LSPs such as QoS, bandwidth, fast reroute, and path selection to tailor the associated bidirectional LSP to network requirements and traffic patterns.</p>

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

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

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

- [Signaling Methods and Object Association for Flex LSPs, on page 2](#)
- [Associated Bidirectional Co-routed LSPs, on page 2](#)
- [Restrictions for Flex LSP, on page 3](#)
- [Key Features supported in Flex LSP, on page 3](#)
- [Flex LSP Scale Details , on page 5](#)
- [How to Configure Co-routed Flex LSPs, on page 5](#)
- [RSVP-TE Transport Solution with Flex LSP Headend, on page 11](#)
- [RSVP-TE transport solution using PCE, on page 22](#)

Signaling Methods and Object Association for Flex LSPs

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


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

Associated Bidirectional Co-routed LSPs

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

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

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



In the above topology:

- Paths at the top of the figure (in green) indicate working co-routed LSP pairs.
- Paths at the bottom of the figure (in red) indicate protecting co-routed LSP pairs.

- Router 1 sets up working LSP to Router 3 (in green) after performing bidirectional CSPF and sends reverse explicit route object (ERO) to Router 3. Node Router 3 uses the received reverse ERO to set up reverse green working LSP to Router 1.
- Router 3 sets up protecting LSP to Router 1 (in red) after performing bidirectional CSPF and sends reverse ERO to Router 1. Node Router 1 uses the received reverse ERO to set up reverse red protecting LSP to Router 3.

Restrictions for Flex LSP

- Exp-null over Flex-LSP is not supported.
- 50 msec convergence is not guaranteed without WRAP protection. WRAP protection is mandatory to achieve 50 msec convergence for remote failures.
- TE NSR and IGP NSR are mandatory for RSP switchover.
- VPLS over Flex-LSP is not supported.
- Non-co routed Flex LSP is not supported.
- Sub interface shut will not guarantee 50 msec convergence.
- MPLS forwarding table stats is not supported.
- 1000 tunnels are supported with wrap protection and path protection.

Key Features supported in Flex LSP

Following list outlines key features supported:

• Protection -

Following Protection features are supported:

- **Lockout:** Using Lockout feature, user can perform lockout protection on a selected LSP and can switch traffic to protecting LSP, if the selected LSP carries the traffic. LSP on locked out interface remains up but no traffic flows on the locked out path. To configure lockout under an MPLS-TE enabled interface:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# mpls traffic-eng
RP/0/RP0/CPU0:router(config-mpls-te)# interface tenGigE0/1/0/1
RP/0/RP0/CPU0:router(config-mpls-te-if)# fault-oam lockout
```

- **Wrap Protection:** Using Wrap Protection, each LSP signals unique wrap label for head-end to identify lookback traffic and sends it over protect LSP. To configure Wrap Protection:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# interface tunnel-te1
RP/0/RP0/CPU0:router(config-if)# ipv4 unnumbered Loopback0
RP/0/RP0/CPU0:router(config-if)# destination 49.49.49.2
RP/0/RP0/CPU0:router(config-if)# path-option 10 explicit name PATH1-2-3
RP/0/RP0/CPU0:router(config-if)# bidirectional association id 100 source-address
```

```
49.49.49.2
RP/0/RP0/CPU0:router(config-if)# bidirectional association association type co-routed
RP/0/RP0/CPU0:router(config-if)# wrap-protection
RP/0/RP0/CPU0:router(config-if)# fault-oam
```

- **MPLS-OAM:** MPLS-OAM supports single segment pseudowire going over the associated bidirectional TE tunnels. This support includes pseudowires signaled dynamically, statically, or using a mix of both modes. To configure MPLS-OAM, use the following command:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# mpls oam
RP/0/RP0/CPU0:router(config-oam)# echo reply-mode control-channel allow-reverse-lsp
```

Following features are supported:

- **LSP Ping:** Using LSP ping, use can enable on demand ping. It supports IP encapsulation for both request and reply messages. It also performs reverse path verification.

Following is the example of LSP Ping configuration:

```
RP/0/RP0/CPU0:router# ping mpls traffic-eng tunnel-te 1 reply mode control-channel

Tue May 21 11:04:12.211 EDT
Sending 5, 100-byte MPLS Echos to tunnel-tel,
    timeout is 2 seconds, send interval is 0 msec:

Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
       'L' - labeled output interface, 'B' - unlabeled output interface,
       'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
       'M' - malformed request, 'm' - unsupported tlvs, 'N' - no rx label,
       'P' - no rx intf label prot, 'p' - premature termination of LSP,
       'R' - transit router, 'I' - unknown upstream index,
       'X' - unknown return code, 'x' - return code 0

Type escape sequence to abort.

!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 3/4/5 ms
```

- **Traceroute:** Traceroute supports IP-encapsulation but not for echo reply. Following is the example of Traceroute configuration:

```
RP/0/RP0/CPU0:router# traceroute mpls traffic-eng tunnel-te 1

Tue May 21 11:06:16.056 EDT
Tracing MPLS TE Label Switched Path on tunnel-tel, timeout is 2 seconds

Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
       'L' - labeled output interface, 'B' - unlabeled output interface,
       'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
       'M' - malformed request, 'm' - unsupported tlvs, 'N' - no rx label,
       'P' - no rx intf label prot, 'p' - premature termination of LSP,
       'R' - transit router, 'I' - unknown upstream index,
       'X' - unknown return code, 'x' - return code 0

Type escape sequence to abort.

 0 10.10.10.1 MRU 1500 [Labels: 16005 Exp: 0]
L 1 10.10.10.2 MRU 1500 [Labels: implicit-null Exp: 0] 11 ms
! 2 13.13.13.4 4 ms
```

Flex LSP Scale Details

Scale details for Flex LSP:

Table 2: Supported LSPs for Flex LSP

Flex LSP with wrap protection	Head/Tail Node: 18750 LSPs Mid Node: 12500 LSPs
Flex LSP without wrap protection	Head/Tail Node: 18750 LSPs Mid Node: 75000 LSPs

How to Configure Co-routed Flex LSPs

A co-routed bidirectional packet LSP is a combination of two LSPs (one in the forward direction and the other in reverse direction) sharing the same path between a pair of ingress and egress nodes. It is established using the extensions to RSVP-TE. This type of LSP can be used to carry any of the standard types of MPLS-based traffic, including Layer 2 VPNs.

The configuration includes the following steps:

1. Enable basic MPLS Traffic Engineering on hostname PE1 and RSVP Configuration.
2. Configure Flex LSP.
3. Enable Wrap Protection.
4. Enable Fault OAM.
5. Map pseudowire to a specific Flex LSP tunnel.

Configuring Co-routed Flex LSPs

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



Note Up to 1000 Flex LSP tunnels are supported.

Before you begin

Procedure

1. **Enable basic MPLS Traffic Engineering on hostname PE1 and RSVP Configuration:**
Configure MPLS-TE;

```
Router# configure
Router(config)# mpls traffic-eng
Router(config-mpls-te)# interface TenGigE0/9/0/0/12.1
```

Configure RSVP:

```
Router# configure
Router(config)# rsvp
Router(config-rsvp)# signalling hello graceful-restart refresh interval 3000
Router(config-rsvp)# interface TenGigE0/8/0/0/102.1 bandwidth 1000000
Router(config-rsvp)# interface TenGigE0/5/0/4.1 bandwidth 1000000
```

2. Configure Flex LSP:

```
Router# configure
Router(config)# interface tunnel-tel
Router(config-if)# ipv4 unnumbered Loopback0
Router(config-if)# destination 49.49.49.2
Router(config-if)# path-option 10 explicit name PATH1-2-3
Router(config-if)# bidirectional
Router(config-if-bidir)# association id 100 source-address 49.49.49.2
Router(config-if-bidir-co-routed)# association type co-routed
```

3. Wrap Protection:

For Wrap Protection:

```
Router# configure
Router(config)# interface tunnel-tel
Router(config-if)# ipv4 unnumbered Loopback0
Router(config-if)# destination 49.49.49.2
Router(config-if)# path-option 10 explicit name PATH1-2-3
Router(config-if)# bidirectional
Router(config-if-bidir)# association id 100 source-address 49.49.49.2
Router(config-if-bidir)# association type co-routed
Router(config-if-bidir-co-routed)# wrap-protection
```

4. Enable Fault OAM

```
Router# configure
Router(config)# interface tunnel-tel
Router(config-if)# ipv4 unnumbered Loopback0
Router(config-if)# destination 49.49.49.2
Router(config-if)# path-option 10 explicit name PATH1-2-3
Router(config-if)# bidirectional
Router(config-if-bidir)# association id 100 source-address 49.49.49.2
Router(config-if-bidir)# association type co-routed
Router(config-if-bidir-co-routed)# fault-oam
```

5. Map pseudowire to a specific Flex LSP tunnel:

```
Router# configure
Router(config)# l2vpn
Router(config-l2vpn)# pw-class foo
Router(config-l2vpn-pwc)# encapsulation mpls
Router(config-l2vpn-pwc-mpls)# preferred-path interface tunnel-te 1
Router(config-l2vpn-pwc)# exit
Router(config-l2vpn)# exit
Router(config-l2vpn)# xconnect group gold
Router(config-l2vpn-xc)# p2p cust_one
Router(config-l2vpn-xc-p2p)# neighbor ipv4 49.49.49.2 pw-id 1
```

```
Router(config-l2vpn-xc-p2p-pw)# pw-class foo
```

Verifying the Co-routed Flex LSP Configuration

To verify the co-routed LSP, use the **show mpls traffic-eng tunnels detail** command.

```
Router# show mpls traffic-eng tunnels 7001 detail

Name: tunnel-te7001 Destination: 104.0.0.1 Ifhandle:0x8000aa4
  Signalled-Name: Cisco8K-R11_t7001
  Status:
    Admin:      up Oper:      up (Uptime 136y10w)

    path option 1, type explicit path01 (Basis for Setup, path weight 30 (reverse 30))
      Protected-by PO index: 2
    path option 2, type explicit path02 (Basis for Standby, path weight 100010 (reverse
100010))
      Protected-by PO index: 1
      G-PID: 0x0800 (derived from egress interface properties)
      Bandwidth Requested: 10 kbps CT0
      Creation Time: Wed Jan 11 03:08:36 2017 (136y10w ago)
  Config Parameters:
    Bandwidth:      10 kbps (CT0) Priority:  7  7 Affinity: 0x0/0xffff
    Metric Type: TE (interface)
    Path Selection:
      Tiebreaker: Min-fill (default)
    Hop-limit: disabled
    Cost-limit: disabled
    Path-invalidation timeout: 10000 msec (default), Action: Tear (default)
    AutoRoute: disabled LockDown: disabled Policy class: not set
    Forward class: 0 (default)
    Forwarding-Adjacency: disabled
    Autoroute Destinations: 0
    Loadshare:      0 equal loadshares
    Auto-bw: disabled
    Fast Reroute: Disabled, Protection Desired: None
    Path Protection: Enabled
    Association Type: Single Sided Bidirectional LSPs, Co-routed: YES
    Association ID: 86, Source: 192.0.0.0
    Reverse Bandwidth: 10 kbps (CT0), Standby: 10 kbps (CT0)
    LSP Wrap Protection: Enabled

    Reoptimization after affinity failure: Enabled
    Soft Preemption: Disabled
  Fault-OAM Info:
    Last Fault Msg: Clear
  SNMP Index: 25
  Binding SID: None
  Path Protection Info:
    Standby Path: User defined [explicit path option: 2],
    Last Switchover:
      136y10w ago, From LSP 14 To LSP 16
      No subcause recorded
      Reopt time remaining: 0 seconds
    Number of Switchovers 1, Standby Ready 3 times, Standby Reopt 0 times
  Lockout Info:
    Locked Out: NO
    Locked out LSP ID: 0
    Lockout Originated By: None
  LSP Wrap Protection: Enabled
    LSP Wrap Label: 24182
```

```

History:
  Reopt. LSP:
    Last Failure:
      LSP not signalled, identical to the [CURRENT] LSP
      Date/Time: Tue Jan 10 21:42:41 UTC 2017 [00:03:42 ago]
  Standby Reopt LSP:
    Last Failure:
      LSP not signalled, identical to the [STANDBY] LSP
      Date/Time: Tue Jan 10 21:42:41 UTC 2017 [00:03:42 ago]
      First Destination Failed: 104.0.0.1
  Prior LSP:
    ID: 14 Path Option: 1
    Removal Trigger: path protection switchover
  Current LSP Info:
    Instance: 18, Signaling Area: IS-IS 100 level-2
    Uptime: 136y10w (since Wed Jan 11 03:09:56 UTC 2017)
    Outgoing Interface: TenGigE0/4/0/2.1, Outgoing Label: 24157
    Router-IDs: local      102.0.0.1
                  downstream 107.0.0.1
  Soft Preemption: None
  SRLGs: not collected
  Path Info:
    Outgoing:
      Explicit Route:
        Strict, 1.27.1.2
        Strict, 3.67.1.2
        Strict, 3.67.1.1
        Strict, 1.46.1.2
        Strict, 1.46.1.1
        Strict, 104.0.0.1

  Record Route: Disabled
  Tspec: avg rate=10 kbits, burst=1000 bytes, peak rate=10 kbits
  Session Attributes: Local Prot: Not Set, Node Prot: Not Set, BW Prot: Not Set
                    Soft Preemption Desired: Not Set
  Reverse Associated LSP Information:
    Signaled Name: Cisco8K-R10 t7001
    Tunnel: 7001, Source: 104.0.0.1, Dest: 102.0.0.1, LSP: 9, State: Up
  Association:
    Association Type: Single Sided Bidirectional LSPs
    Association ID: 86, Source: 192.0.0.0
  Extended Association:
    Global source: 0
    Extended ID:
      0x66000001 (102.0.0.1)
      0x12 (0.0.0.18)
  Protection:
    Secondary (S): 0, Protecting (P): 0, Notification (N): 0, Oper (O): 0
    Link Flags: Any, LSP Flags: 1:N Protection with Extra-Traffic
    Reverse Tspec: avg rate=10 kbits, burst=1000 bytes, peak rate=10 kbits
  Reverse ERO:
    Explicit Route:
      Strict, 1.46.1.1
      Strict, 1.46.1.2
      Strict, 3.67.1.1
      Strict, 3.67.1.2
      Strict, 1.27.1.2
      Strict, 1.27.1.1
      Strict, 102.0.0.1

  Resv Info: None
  Record Route: Disabled
  Fspec: avg rate=10 kbits, burst=1000 bytes, peak rate=10 kbits
  Standby LSP Info:

```



```

Instance: 19, Signaling Area: IS-IS 100 level-2
Uptime: 136y10w (since Wed Jan 11 03:10:04 UTC 2017), Oper State: Up
Outgoing Interface: TenGigE0/4/0/11.1, Outgoing Label: 24176
Router-IDs: local      102.0.0.1
             downstream 109.0.0.1
Soft Preemption: None
SRLGs: not collected
Path Info:
  Outgoing:
    Explicit Route:
      Strict, 1.29.1.2
      Strict, 1.49.1.2
      Strict, 1.49.1.1
      Strict, 104.0.0.1

Record Route: Disabled
Tspec: avg rate=10 kbits, burst=1000 bytes, peak rate=10 kbits
Session Attributes: Local Prot: Not Set, Node Prot: Not Set, BW Prot: Not Set
                   Soft Preemption Desired: Not Set
Reverse Associated LSP Information:
  Signaled Name: Cisco8K-R10_t7001
  Tunnel: 7001, Source: 104.0.0.1, Dest: 102.0.0.1, LSP: 10, State: Up
Association:
  Association Type: Single Sided Bidirectional LSPs
  Association ID: 86, Source: 192.0.0.0
Extended Association:
  Global source: 0
  Extended ID:
    0x68000001 (104.0.0.1)
    0xa (0.0.0.10)
Protection:
  Secondary (S): 0, Protecting (P): 1, Notification (N): 0, Oper (O): 0
  Link Flags: Any, LSP Flags: 1:N Protection with Extra-Traffic
Resv Info: None
Record Route: Disabled
Fspec: avg rate=10 kbits, burst=1000 bytes, peak rate=10 kbits
Persistent Forwarding Statistics:
  Out Bytes: 20272384
  Out Packets: 79189

LSP Tunnel 104.0.0.1 7001 [9] is signalled, Signaling State: up
Tunnel Name: Cisco8K-R10_t7001 Tunnel Role: Tail
InLabel: TenGigE0/4/0/2.1, 24164
Signalling Info:
  Src 104.0.0.1 Dst 102.0.0.1, Tun ID 7001, Tun Inst 9, Ext ID 104.0.0.1
Router-IDs: upstream 107.0.0.1
             local    102.0.0.1
Bandwidth: 10 kbps (CT0) Priority: 7 7 DSTE-class: 0
Soft Preemption: None
SRLGs: not collected
Path Info:
  Incoming Address: 1.27.1.1
  Incoming:
    Explicit Route:
      Strict, 1.27.1.1
      Strict, 102.0.0.1

Record Route: Disabled
Tspec: avg rate=10 kbits, burst=1000 bytes, peak rate=10 kbits
Session Attributes: Local Prot: Not Set, Node Prot: Not Set, BW Prot: Not Set
                   Soft Preemption Desired: Not Set
Reverse Associated LSP Information:
  Signaled Name: Cisco8K-R11_t7001
  Tunnel: 7001, Source: 102.0.0.1, Dest: 104.0.0.1, LSP: 18, State: Up

```

```

Association:
  Association Type: Single Sided Bidirectional LSPs (Tie breaking slave)
  Association ID: 86, Source: 192.0.0.0
Extended Association:
  Global source: 0
  Extended ID:
    0x66000001 (102.0.0.1)
    0x12 (0.0.0.18)
Protection:
  Secondary (S): 0, Protecting (P): 0, Notification (N): 0, Oper (O): 0
  Link Flags: Any, LSP Flags: 1:N Protection with Extra-Traffic
Resv Info: None
  Record Route: Disabled
  Fspec: avg rate=10 kbits, burst=1000 bytes, peak rate=10 kbits

LSP Tunnel 104.0.0.1 7001 [10] is signalled, Signaling State: up
Tunnel Name: Cisco8K-R10_t7001 Tunnel Role: Tail
InLabel: TenGigE0/4/0/11.1, 24463
Signalling Info:
  Src 104.0.0.1 Dst 102.0.0.1, Tun ID 7001, Tun Inst 10, Ext ID 104.0.0.1
  Router-IDs: upstream 109.0.0.1
               local 102.0.0.1
  Bandwidth: 10 kbps (CT0) Priority: 7 7 DSTE-class: 0
  Soft Preemption: None
  SRLGs: not collected
  Path Info:
    Incoming Address: 1.29.1.1
    Incoming:
    Explicit Route:
      Strict, 1.29.1.1
      Strict, 102.0.0.1

  Record Route: Disabled
  Tspec: avg rate=10 kbits, burst=1000 bytes, peak rate=10 kbits
  Session Attributes: Local Prot: Not Set, Node Prot: Not Set, BW Prot: Not Set
                     Soft Preemption Desired: Not Set
  Reverse Associated LSP Information:
    Signaled Name: Cisco8K-R11_t7001
    Tunnel: 7001, Source: 102.0.0.1, Dest: 104.0.0.1, LSP: 19, State: Up
  Association:
    Association Type: Single Sided Bidirectional LSPs (Tie breaking slave)
    Association ID: 86, Source: 192.0.0.0
    Extended Association: Fspec: avg rate=8K, burst=1K, peak rate=8K

```

To verify the forwarding interface, use the **show mpls forwarding tunnels detail** command.

```
Router# show mpls forwarding tunnels 7001 detail
```

Tunnel Name	Outgoing Label	Outgoing Interface	Next Hop	Bytes Switched
tt7001	24157	Te0/4/0/2.1	1.27.1.2	0

```

Updated: Jan 10 21:40:04.966
Version: 17852, Priority: 2
Label Stack (Top -> Bottom): { 24157 }
Local Label: 24354
NHID: 0x0, Encap-ID: INVALID, Path idx: 0, Backup path idx: 0, Weight: 0
MAC/Encaps: 18/22, MTU: 1500
Packets Switched: 0

Interface Name: tunnel-te7001, Interface Handle: 0x08000aa4, Local Label: 24354
Forwarding Class: 0, Weight: 0
Packets/Bytes Switched: 79189/20272384

```

RSVP-TE Transport Solution with Flex LSP Headend

Table 3: Feature History Table

Feature Name	Release Information	Feature Description
RSVP-TE Transport Solution with Flex LSP Headend	Release 25.1.1	<p>Introduced in this release on: Fixed Systems (8200 [ASIC: Q200, P100], 8700 [ASIC: P100, K100])(select variants only*); Centralized Systems (8600 [ASIC:Q200]); Modular Systems (8800 [LC ASIC: Q100, Q200, P100])</p> <p>You can now configure up to four backup paths for both primary and secondary paths. This is an improvement over previous configurations that supported only two backup paths.</p> <p>This enhance the network's ability to handle failures gracefully and optimize performance</p> <p>You can now configure headend router or a centralized Path Computation Element (PCE) server to perform flex LSP path computation. For more information about PCE, see RSVP-TE transport solution using PCE.</p> <p>*This feature is now supported on Cisco 8712-MOD-M router.</p>

Feature Name	Release Information	Feature Description
RSVP-TE Transport Solution with Flex LSP Headend	Release 24.4.1	<p>Introduced in this release on: Fixed Systems (8200 [ASIC: Q200, P100], 8700 [ASIC: P100]); Centralized Systems (8600 [ASIC: Q200]); Modular Systems (8800 [LC ASIC: Q100, Q200, P100])</p> <p>You can enable the router to compute</p> <ul style="list-style-type: none"> • a primary LSP path, • a secondary LSP path, • a dynamic restore path serves as a backup to the primary LSP path, and • another dynamic restore path acts as a backup to the secondary LSP path. <p>During the primary LSP path failure, the headend router computes a new dynamic restore primary path by reusing the links from the failed primary LSP path, and during the secondary LSP path or backup path or protecting path failure, the headend router computes a new dynamic restore path by reusing the links from the failed secondary LSP path.</p> <p>CLI:</p> <p>The protected-by keyword is enhanced to include two backup paths in the path-option command.</p>

RSVP-TE (Resource Reservation Protocol-Traffic Engineering) is a protocol used to reserve resources across an MPLS (Multiprotocol Label Switching) network to ensure Quality of Service (QoS) for specific traffic flows.

Flex LSP (Flexible Label Switched Path) is an enhancement of RSVP-TE that allows for more flexible and efficient selection of co-routed bidirectional LSP path selection.

Co-routed bidirectional LSP paths refer to LSP paths where traffic flows from headend to tailend and vice versa over the same physical path.

You can enable the router to configure these co-routed bidirectional LSP paths:

- **Primary flex LSP path:** This is an explicit path computed by the router between a headend node and a tailend node. This path serves as the main route for traffic.
- **Secondary flex LSP path:** In addition to the primary path, the router computes a secondary flex LSP path. This path acts as a backup to the primary path.
- **Dynamic restore paths:**
 - **Primary backup:** The headend router computes a dynamic LSP path using Constrained Shortest Path First (CSPF). This path serves as a backup to the primary flex LSP path.
 - **Secondary backup:** Another dynamic restore path is configured as a backup for the secondary flex LSP path. It provides additional redundancy and ensures continuous traffic flow even if the secondary path becomes unavailable.

For more information, see [RSVP-TE Transport Solution with Flex LSP](#).

Once the failure is resolved, the router restores the previously used primary or secondary LSP paths.

Before Release 24.4.1, you can enable the router to configure

- a secondary LSP path as a backup to the primary LSP path, and
- dynamic restore path as a backup to the secondary LSP path.

From Release 24.4.1 onwards, you can enable the router to configure

- a secondary LSP path and a dynamic restore primary path as backups to the primary LSP path, and
- a dynamic restore path as a backup to the secondary LSP path.

From Release 25.1.1 onwards, you can configure a centralized PCE server to perform flex LSP path computation for primary, secondary, and restore paths. For more information, see [RSVP-TE transport solution using PCE, on page 22](#).

How flex LSP headend traffic flow works

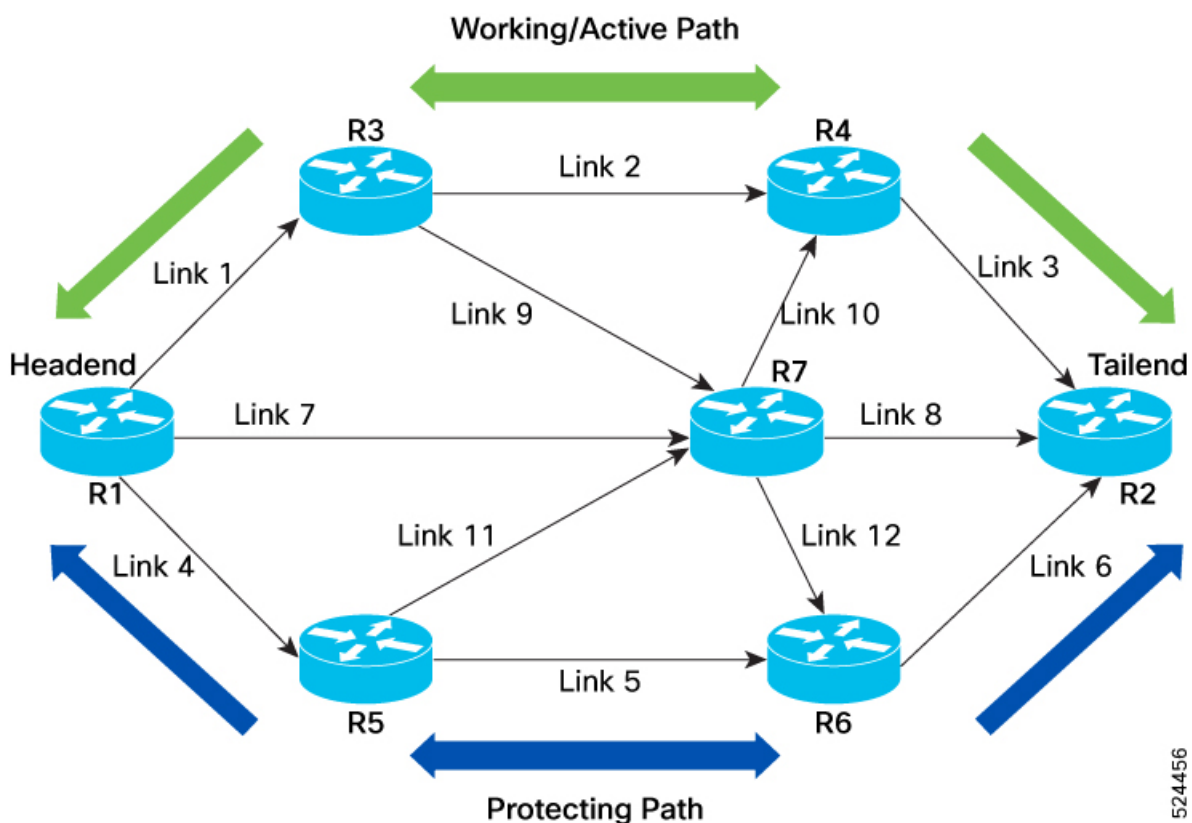
Summary

The key components that are involved in flex LSP headend traffic flow are:

- R1: Headend router
- R2: Tailend router
- R3, R4, R5, R6, and R7: Transit routers
- Path 10 (Link 1-2-3): Primary LSP path or Working or Active path
- Path 20 (Link 4-5-6): Secondary LSP path or Protecting path
- Path 30 (Link 1-9-10-3): Dynamic restore primary path
- Path 40 (Link 4-11-12-6): Dynamic restore path
- Path 50 (Link 7-8): Dynamic LSP path
- All links are 100G
- Weight:
 - Link 1 to Link 6: 10000 weights
 - Link 7 to Link 12: 100 weights

Workflow

Figure 1: Flex LSP Headend Traffic Flow



524456

These stages describe how flex LSP headend traffic flow works.

Before Release 24.4.1,

When	Then...
Link 2 of the primary LSP path is broken	the headend router activates the secondary LSP Path 20.
both primary and secondary LSP paths are broken	the headend router computes a new dynamic path using CSPF, without prioritizing the primary or secondary paths.
for example, Link 2 and Link 5 are broken	the headend router computes a new dynamic LSP Path 50, selecting the shortest path compared to Path 30 or Path 40.

From Release 24.4.1 onwards,

When	Then...
Link 2 of the primary LSP path is broken	the headend router activates the secondary LSP Path 20.

When	Then...
both primary and secondary LSP paths are broken	the headend router computes a new dynamic restore primary path, prioritizing the primary LSP path.
for example, Link 2 and Link 5 are broken	the headend router prioritizes Link 1 and assigns a new dynamic primary Path 30.
primary, secondary, and dynamic primary LSP paths are broken	the headend router computes a new dynamic restore path, prioritizing the secondary LSP path.
for example, Link 2, Link 5, and Link 10 are broken	the headend router prioritizes Link 4 and assigns a new dynamic restore Path 40.

Flex LSP headend scale details

By default, 1000 RSVP-TE tunnels are enabled on the edge routers. If you want to scale up the edge routers to support more than 1000 RSVP-TE tunnels

1. Scale up the edge routers to support maximum of 8000 RSVP-TE tunnels with **hw-module profile cef te-tunnel highscale-no-ldp-over-te** command.
2. Reload all the LCs in the chassis with the **reload location all** command to enable the maximum of 8000 RSVP-TE tunnels.
3. Enable the RSVP-TE transport solution with flex LSP headend feature.

Below are the scale details for flex LSP headend.

Table 4: Supported LSPs for flex LSP headend

Supported LSPs for flex LSP headend	Scale Coverage
Flex LSP with path protection	Head/Tail Node: 4000 LSPs Mid Node: 8000 LSPs
Flex LSP without path protection	Head/Tail Node: 4000 LSPs Mid Node: 8000 LSPs

Limitations and restrictions for flex LSP headend

The flex LSP headend has these limitations and restrictions:

- This feature does not support the exp-null over Flex-LSP feature.
- This feature does not guarantee 50 msec convergence without WRAP protection. WRAP protection is mandatory to achieve 50 msec convergence for remote failures.
- TE NSR and IGP NSR are mandatory for RSP switchover.
- This feature does not support VPLS over Flex-LSP.

- This feature does not support non-co routed Flex LSP.
- Sub interface shut does not guarantee 50 msec convergence.
- This feature does not support MPLS forwarding table stats on mid nodes.
- This feature supports 4000 tunnels with and without wrap protection and path protection.
- This feature does not support the following features:
 - Bidirectional Forwarding Detection (BFD) over Flex-LSP
 - Multicast over Flex-LSP
 - Policy-Based Traffic Steering (PBTS)
 - Policy-Based Routing (PBR)
 - Access Control Lists (ACL)
 - L3VPN over Flex-LSP

Configure the co-routed flex LSPs headend

Perform these steps to configure co-routed flex LSPs headend:

Before you begin

- You must have symmetric source and destination TE router IDs for bidirectional LSPs to be associated.
- Tunnels attributes must be configured identically on both sides of co-routed bidirectional LSP.
- Before enabling the Flex LSP feature, scale up the RSVP-TE tunnels on the edge routers.

Procedure

- Step 1** Scale up the edge routers with the **hw-module profile cef te-tunnel highsacle-no-ldp-over-te** command to support a maximum of 8000 RSVP-TE tunnels.

Example:

```
Router# configure
Router(config)# hw-module profile cef te-tunnel highsacle-no-ldp-over-te
Router(config)# commit
```

- Step 2** You must reload all the LCs in the chassis with the **reload location all** command to enable the 8000 RSVP-TE tunnels.

Example:

```
Router# reload location all
```

- Step 3** Configure MPLS-TE on Provider Edge (PE) router with **mpls traffic-eng** command to enable the MPLS traffic engineering on PE.

```
Router# configure
Router(config)# mpls traffic-eng
Router(config-mpls-te)# interface HundredGigE0/9/0/0
```


Step 4 Configure RSVP on PE with **rsvp** command to enable the RSVP on PE.

```
Router# configure
Router(config)# rsvp
Router(config-rsvp)# interface HundredGigE0/0/0/37 bandwidth percentage 90
Router(config-rsvp)# interface HundredGigE0/1/0/40 bandwidth percentage 90
```

Step 5 Configure flex LSP with **association type co-routed** command to enable bidirectional co-routed LSPs.

```
Router# configure
Router(config)# interface tunnel-te7001
Router(config-if)# ipv4 unnumbered Loopback0
Router(config-if)# destination 192.0.2.1
Router(config-if)# bidirectional
Router(config-if-bidir)# association id 1 source-address 192.0.2.0
Router(config-if-bidir-co-routed)# association type co-routed
```

Step 6 Enable wrap protection with **wrap-protection** command to enable LSP signals unique wrap label for headend.

```
Router(config-if-bidir-co-routed)# wrap-protection
```

Step 7 Configure path protection for a mpls-te tunnel with the **path-protection** command to create explicit and dynamic path.

```
Router(config-if-bidir-co-routed)# path-protection
```

Step 8 Configure working co-routed LSP path, protecting co-routed LSP path, working dynamic path, and protecting dynamic path using **path-option** command to enable the primary, secondary, and dynamic backup LSP paths.

From Release 25.1.1, You can configure up to four protected-by paths for the working and protecting paths, which serve as backup paths for the primary and secondary paths. In this configuration, only two protected-by paths are considered.

```
Router# configure
Router(config)# interface tunnel-te7001
Router(config-if)# path-option 10 explicit name r1-r2-working protected-by 20 40
Router(config-if)# path-option 20 explicit name r1-r2-protecting protected-by 10 30
Router(config-if)# path-option 30 dynamic attribute-set r1-r2-restoreworking protected-by 40
Router(config-if)# path-option 40 dynamic attribute-set r1-r2-restoreprotecting protected-by 30
```

Step 9 Map Pseudowire to a specific Flex LSP tunnel with **pw-class** command to ensure the Pseudowire traffic is transported over the MPLS TE tunnel.

```
Router# configure
Router(config)# l2vpn
Router(config-l2vpn)# pw-class pw1
Router(config-l2vpn-pwc)# encapsulation mpls
Router(config-l2vpn-pwc-mpls)# preferred-path interface tunnel-te1
Router(config-l2vpn-pwc)# exit
Router(config-l2vpn)# exit
Router(config-l2vpn)# xconnect group gold
Router(config-l2vpn-xc)# p2p cust_one
Router(config-l2vpn-xc-p2p)# neighbor ipv4 192.2.2.1 pw-id 1
Router(config-l2vpn-xc-p2p-pw)# pw-class pw1
```

Step 10 Configure explicit working dynamic path using **attribute-set** command to specify criteria for the working dynamic path computation.

```
Router# configure
Router(config)# mpls traffic
Router(config-if-mpls-te)# biased-to explicit-path r1-r2-working
Router(config-te-attribute-set)# commit
```

- Step 11** Configure explicit protecting dynamic path using **attribute-set** command to specify criteria for the protecting dynamic path computation.

```
Router# configure
Router(config)# mpls traffic
Router(config-if-mpls-te)# attribute-set path-option r1-r2-restoreprotecting
Router(config-te-attribute-set)# commit
```

- Step 12** Verify the configured co-routed flex LSP with **show mpls traffic-eng tunnels detail** command.

Example:

The following example displays the co-routed flex LSP configuration with interface tunnel 7001 and destination address 104.0.0.1:

```
Router# show mpls traffic-eng tunnels 7001 detail

Name: tunnel-te7001 Destination: 104.0.0.1 Ifhandle:0x8000aa4
Signalled-Name: Cisco8K-R11_t7001
Status:
  Admin:      up Oper:      up (Uptime 136y10w)

  path option 1, type explicit path01 (Basis for Setup, path weight 30 (reverse 30))
    Protected-by PO index: 2
  path option 2, type explicit path02 (Basis for Standby, path weight 100010 (reverse 100010))
    Protected-by PO index: 1
  G-PID: 0x0800 (derived from egress interface properties)
  Bandwidth Requested: 10 kbps CT0
  Creation Time: Wed Jan 11 03:08:36 2017 (136y10w ago)
Config Parameters:
  Bandwidth:      10 kbps (CT0) Priority:  7  7 Affinity: 0x0/0xffff
  Metric Type: TE (interface)
  Path Selection:
    Tiebreaker: Min-fill (default)
  Hop-limit: disabled
  Cost-limit: disabled
  Path-invalidation timeout: 10000 msec (default), Action: Tear (default)
  AutoRoute: disabled LockDown: disabled Policy class: not set
  Forward class: 0 (default)
  Forwarding-Adjacency: disabled
  Autoroute Destinations: 0
  Loadshare:      0 equal loadshares
  Auto-bw: disabled
  Fast Reroute: Disabled, Protection Desired: None
  Path Protection: Enabled
  Association Type: Single Sided Bidirectional LSPs, Co-routed: YES
  Association ID: 86, Source: 192.0.0.0
  Reverse Bandwidth: 10 kbps (CT0), Standby: 10 kbps (CT0)
  LSP Wrap Protection: Enabled

  Reoptimization after affinity failure: Enabled
  Soft Preemption: Disabled
Fault-OAM Info:
  Last Fault Msg: Clear
SNMP Index: 25
Binding SID: None
Path Protection Info:
  Standby Path: User defined [explicit path option: 2],
  Last Switchover:
    136y10w ago, From LSP 14 To LSP 16
    No subcause recorded
    Reopt time remaining: 0 seconds
  Number of Switchovers 1, Standby Ready 3 times, Standby Reopt 0 times
```

```

Lockout Info:
  Locked Out: NO
  Locked out LSP ID: 0
  Lockout Originated By: None
LSP Wrap Protection: Enabled
  LSP Wrap Label: 24182
History:
  Reopt. LSP:
    Last Failure:
      LSP not signalled, identical to the [CURRENT] LSP
      Date/Time: Tue Jan 10 21:42:41 UTC 2017 [00:03:42 ago]
  Standby Reopt LSP:
    Last Failure:
      LSP not signalled, identical to the [STANDBY] LSP
      Date/Time: Tue Jan 10 21:42:41 UTC 2017 [00:03:42 ago]
      First Destination Failed: 104.0.0.1
  Prior LSP:
    ID: 14 Path Option: 1
    Removal Trigger: path protection switchover
Current LSP Info:
  Instance: 18, Signaling Area: IS-IS 100 level-2
  Uptime: 136y10w (since Wed Jan 11 03:09:56 UTC 2017)
  Outgoing Interface: TenGigE0/4/0/2.1, Outgoing Label: 24157
  Router-IDs: local      102.0.0.1
               downstream 107.0.0.1
  Soft Preemption: None
  SRLGs: not collected
  Path Info:
    Outgoing:
      Explicit Route:
        Strict, 1.27.1.2
        Strict, 3.67.1.2
        Strict, 3.67.1.1
        Strict, 1.46.1.2
        Strict, 1.46.1.1
        Strict, 104.0.0.1

Record Route: Disabled
Tspec: avg rate=10 kbits, burst=1000 bytes, peak rate=10 kbits
Session Attributes: Local Prot: Not Set, Node Prot: Not Set, BW Prot: Not Set
                   Soft Preemption Desired: Not Set
Reverse Associated LSP Information:
  Signaled Name: Cisco8K-R10_t7001
  Tunnel: 7001, Source: 104.0.0.1, Dest: 102.0.0.1, LSP: 9, State: Up
Association:
  Association Type: Single Sided Bidirectional LSPs
  Association ID: 86, Source: 192.0.0.0
Extended Association:
  Global source: 0
  Extended ID:
    0x66000001 (102.0.0.1)
    0x12 (0.0.0.18)
Protection:
  Secondary (S): 0, Protecting (P): 0, Notification (N): 0, Oper (O): 0
  Link Flags: Any, LSP Flags: 1:N Protection with Extra-Traffic
Reverse Tspec: avg rate=10 kbits, burst=1000 bytes, peak rate=10 kbits
Reverse ERO:
  Explicit Route:
    Strict, 1.46.1.1
    Strict, 1.46.1.2
    Strict, 3.67.1.1
    Strict, 3.67.1.2
    Strict, 1.27.1.2
    Strict, 1.27.1.1

```

```

    Strict, 102.0.0.1

Resv Info: None
Record Route: Disabled
Fspec: avg rate=10 kbits, burst=1000 bytes, peak rate=10 kbits
Standby LSP Info:
Instance: 19, Signaling Area: IS-IS 100 level-2
Uptime: 136y10w (since Wed Jan 11 03:10:04 UTC 2017), Oper State: Up
Outgoing Interface: TenGigE0/4/0/11.1, Outgoing Label: 24176
Router-IDs: local      102.0.0.1
             downstream 109.0.0.1
Soft Preemption: None
SRLGs: not collected
Path Info:
Outgoing:
Explicit Route:
Strict, 1.29.1.2
Strict, 1.49.1.2
Strict, 1.49.1.1
Strict, 104.0.0.1

Record Route: Disabled
Tspec: avg rate=10 kbits, burst=1000 bytes, peak rate=10 kbits
Session Attributes: Local Prot: Not Set, Node Prot: Not Set, BW Prot: Not Set
                   Soft Preemption Desired: Not Set
Reverse Associated LSP Information:
Signaled Name: Cisco8K-R10_t7001
Tunnel: 7001, Source: 104.0.0.1, Dest: 102.0.0.1, LSP: 10, State: Up
Association:
Association Type: Single Sided Bidirectional LSPs
Association ID: 86, Source: 192.0.0.0
Extended Association:
Global source: 0
Extended ID:
0x68000001 (104.0.0.1)
0xa (0.0.0.10)
Protection:
Secondary (S): 0, Protecting (P): 1, Notification (N): 0, Oper (O): 0
Link Flags: Any, LSP Flags: 1:N Protection with Extra-Traffic
Resv Info: None
Record Route: Disabled
Fspec: avg rate=10 kbits, burst=1000 bytes, peak rate=10 kbits
Persistent Forwarding Statistics:
Out Bytes: 20272384
Out Packets: 79189

LSP Tunnel 104.0.0.1 7001 [9] is signalled, Signaling State: up
Tunnel Name: Cisco8K-R10_t7001 Tunnel Role: Tail
InLabel: TenGigE0/4/0/2.1, 24164
Signalling Info:
Src 104.0.0.1 Dst 102.0.0.1, Tun ID 7001, Tun Inst 9, Ext ID 104.0.0.1
Router-IDs: upstream 107.0.0.1
             local    102.0.0.1
Bandwidth: 10 kbps (CT0) Priority: 7 7 DSTE-class: 0
Soft Preemption: None
SRLGs: not collected
Path Info:
Incoming Address: 1.27.1.1
Incoming:
Explicit Route:
Strict, 1.27.1.1
Strict, 102.0.0.1

Record Route: Disabled

```

```

Tspec: avg rate=10 kbits, burst=1000 bytes, peak rate=10 kbits
Session Attributes: Local Prot: Not Set, Node Prot: Not Set, BW Prot: Not Set
                    Soft Preemption Desired: Not Set
Reverse Associated LSP Information:
  Signaled Name: Cisco8K-R11_t7001
  Tunnel: 7001, Source: 102.0.0.1, Dest: 104.0.0.1, LSP: 18, State: Up
Association:
  Association Type: Single Sided Bidirectional LSPs (Tie breaking slave)
  Association ID: 86, Source: 192.0.0.0
Extended Association:
  Global source: 0
  Extended ID:
    0x66000001 (102.0.0.1)
    0x12 (0.0.0.18)
Protection:
  Secondary (S): 0, Protecting (P): 0, Notification (N): 0, Oper (O): 0
  Link Flags: Any, LSP Flags: 1:N Protection with Extra-Traffic
Resv Info: None
Record Route: Disabled
Fspec: avg rate=10 kbits, burst=1000 bytes, peak rate=10 kbits

LSP Tunnel 104.0.0.1 7001 [10] is signalled, Signaling State: up
Tunnel Name: Cisco8K-R10_t7001 Tunnel Role: Tail
InLabel: TenGigE0/4/0/11.1, 24463
Signalling Info:
  Src 104.0.0.1 Dst 102.0.0.1, Tun ID 7001, Tun Inst 10, Ext ID 104.0.0.1
  Router-IDs: upstream 109.0.0.1
               local 102.0.0.1
Bandwidth: 10 kbps (CT0) Priority: 7 7 DSTE-class: 0
Soft Preemption: None
SRLGs: not collected
Path Info:
  Incoming Address: 1.29.1.1
  Incoming:
  Explicit Route:
    Strict, 1.29.1.1
    Strict, 102.0.0.1

Record Route: Disabled
Tspec: avg rate=10 kbits, burst=1000 bytes, peak rate=10 kbits
Session Attributes: Local Prot: Not Set, Node Prot: Not Set, BW Prot: Not Set
                    Soft Preemption Desired: Not Set
Reverse Associated LSP Information:
  Signaled Name: Cisco8K-R11_t7001
  Tunnel: 7001, Source: 102.0.0.1, Dest: 104.0.0.1, LSP: 19, State: Up
Association:
  Association Type: Single Sided Bidirectional LSPs (Tie breaking slave)
  Association ID: 86, Source: 192.0.0.0
Extended Association: Fspec: avg rate=8K, burst=1K, peak rate=8K

```

Step 13 Verify the forwarding interface with the **show mpls forwarding tunnels detail** command.

Example:

```
Router# show mpls forwarding tunnels 7001 detail
```

Tunnel Name	Outgoing Label	Outgoing Interface	Next Hop	Bytes Switched
tt7001	24157	Te0/4/0/2.1	1.27.1.2	0

```

Updated: Jan 10 21:40:04.966
Version: 17852, Priority: 2
Label Stack (Top -> Bottom): { 24157 }
Local Label: 24354

```

NHID: 0x0, Encap-ID: INVALID, Path idx: 0, Backup path idx: 0, Weight: 0
 MAC/Encaps: 18/22, MTU: 1500
 Packets Switched: 0

Interface Name: tunnel-te7001, Interface Handle: 0x08000aa4, Local Label: 24354
 Forwarding Class: 0, Weight: 0
 Packets/Bytes Switched: 79189/20272384

RSVP-TE transport solution using PCE

Table 5: Feature History Table

Feature Name	Release Information	Feature Description
RSVP-TE transport solution using PCE	Release 25.1.1	<p>Introduced in this release on: Fixed Systems (8200 [ASIC: Q200, P100], 8700 [ASIC: P100, K100])(select variants only*); Centralized Systems (8600 [ASIC:Q200]); Modular Systems (8800 [LC ASIC: Q100, Q200, P100])(select variants only*)</p> <p>You can now enable a centralized Path Computation Element (PCE) server to perform LSP path computations for primary, secondary, and restore paths.</p> <p>The PCE efficiently tracks and manages network resources, handling multiple LSP requests and complex network environments.</p> <p>It optimizes LSP routing by dynamically adapting to network changes, reducing crank-back, and enhancing traffic engineering.</p> <p>*This feature is now supported on:</p> <ul style="list-style-type: none"> • 8212-48FH-M • 8711-32FH-M • 8712-MOD-M • 88-LC1-12TH24FH-E • 88-LC1-52Y8H-EM

MPLS Traffic Engineering Path Computation Element (TE-PCE) extends the existing IOS XR Path Computation Element Protocol (PCEP) functionality to provide stateless PCE capabilities, supporting operations on both the MPLS data plane and the IPv4 control plane. For more information, see [RFC 5440](#).

PCE characteristics:

- **Stateless operation:** The PCE operates in a stateless manner, meaning it does not maintain ongoing state information about the LSPs or paths it computes. It provides paths on demand based on real-time network conditions.
- **Request and reply Model:** The headend router sends a Path Computation Request (PCReq) to the PCE and receives a Path Computation Reply (PCRep) in return, containing the computed path details.
- **Topology learning:** The PCE learns about the network topology from the Interior Gateway Protocol (IGP).
- **Bandwidth management:** The PCE maintains information on in-progress bandwidth to reduce the need for computing the path again in the network.
- **Path disjointness:** The PCE can compute disjoint paths, ensuring that two or more LSPs do not share the network resources. This capability enhances network resilience by providing diverse path options that can originate from the same or different headend routers.

The Cisco IOS XRv 9000 is the recommended platform to act as the TE-PCE. For more information, see [Cisco IOS XRv 9000 Router Installation and Configuration Guide](#).

Path definition

- **Explicit path:** A manually configured path that specifies a predetermined sequence of nodes or links for traffic passage from the headend to the tailend node.
- **Dynamic path:** A path automatically computed by either the PCE server or headend router using routing protocols and algorithms to adapt to real-time network conditions.

Before Release 25.1.1, the headend router were used to perform flex LSP path computation for these paths:

- **Primary path:** Depending on network requirements, you could configure an explicit path or enable the headend router to compute a dynamic path.
- **Secondary path:** Depending on network requirements, you could configure an explicit path or enable the headend router to compute a dynamic path.
- **Restore path for primary:** Headend router computes a new dynamic path by reusing the links from the failed primary path.
- **Restore path for secondary:** Headend router computes a new dynamic path by reusing the links from the failed secondary path.

For more information, see [RSVP-TE Transport Solution with Flex LSP Headend, on page 11](#).

From Release 25.1.1, a centralized TE-PCE is introduced to perform the flex LSP paths computation for these paths:

- **Primary path:** Depending on network requirements, you can configure an explicit path or enable the PCE or headend router to compute a dynamic path.
- **Secondary path:** Depending on network requirements, you can configure an explicit path or enable the PCE or headend router to compute a dynamic path.
- **Restore path for primary:** The PCE server computes a new dynamic path by reusing links from the failed primary path.
- **Restore path for secondary:** The PCE server computes a new dynamic path by reusing links from the failed secondary path.

If the PCE connection is broken, the headend router performs flex LSP path computation.

Table 6: Difference between flex LSP path computation using headend and flex LSP path computation using PCE

Aspects	Flex LSP path computation using headend	Flex LSP path computation using PCE
Path computation	Performed by headend router.	Centralized computation by the PCE.
State maintenance	Stateful: Maintains state at mid-point nodes.	Stateless: Does not maintain an ongoing database of LSPs.
Bandwidth management	Prevents overbooking by tracking “in-progress” bandwidth based on reservations at mid-point nodes.	Prevents overbooking by computing bandwidth on real-time.
Topology updates	Synchronizes path and bandwidth information upon receiving new topology updates.	Computes based on the latest topology updates, without storing past states.
Network view	Limited to the information available at each node.	Utilizes a global view of the network, enabling optimized path computation.
Flexibility	Requires manual configuration for complex path requirements.	Offers flexibility through centralized computation, accommodating various constraints and requirements.
Redundancy handling	Need additional mechanisms for redundancy and failover management.	Incorporates redundancy and failover within its centralized computation model.

How PCEP session works

The headend or tailend routers which tries to connect to the PCE are referred to as Path Computation Client (PCC).

Connect the PCC to PCE server using the **mpls traffic pce server ipv4** command.

1. The PCC maintains the PCEP session until connection is active.
2. If the connection between PCC and PCE is broken, the PCC performs these steps to retry connecting to PCE:
 - a. The PCC sends the PC request and waits to receive the PC reply from PCE.
 - b. If the PCE does not respond to the PCC within this timeout, the PCC skips path-options that require the PCE and moves on to the next available path-option.
 By default, the PCC timeout for a PC request is 10 seconds. You can modify this timeout to any value between 5 to 100 seconds using the request-timeout command.
 - c. If there is no more available path-options, then the PCC will retry the PCE path-option.

How PCE handles multiple PC request

The PCE performs these steps to compute paths concurrently when it receives multiple PC request:

When...	Then...
PCE receives the first LSP request with a specified disjoint-group ID	the PCE computes the first LSP, encoding the shortest path from the first source to the first destination.
PCE receives the second LSP request with the same disjoint-group ID	the PCE utilizes the information from both requests to compute two disjoint paths concurrently.

Advantages of PCE

These are the advantages of RSVP-TE transport solution using PCE:

- You can connect multiple PCE servers to the headend and tailend router to enhance network redundancy.
- The PCC continues to verify PCE-computed paths even if the connection to the PCE is lost, ensuring that tunnels remain unaffected.
- The PCE computes co-routed bidirectional paths automatically when the reverse-bandwidth constraint is present, removing the need for special indication for co-routed and bidirectional computation.

Limitations and restrictions for RSVP-TE transport solution using PCE

The PCE has these limitations and restrictions:

- For the flex LSP headend limitations, see [RSVP-TE Transport Solution with Flex LSP Headend, on page 11](#).
- In addition to the traditional path request parameters such as source and destination in the endpoints object as per RFC 5440 section 7.6, the path computation service supports the following:
 - Co-routed bidirectional dynamic computation
 - Forward and reverse asymmetrical bandwidth constraints
 - Affinity constraints
 - Priority for the bandwidth constraints
 - Diversity with an existing path
 - SRLG Diversity: on or off
 - Bias path to influence the computation to use certain links in the topology

How RSVP-TE transport solution with PCE works

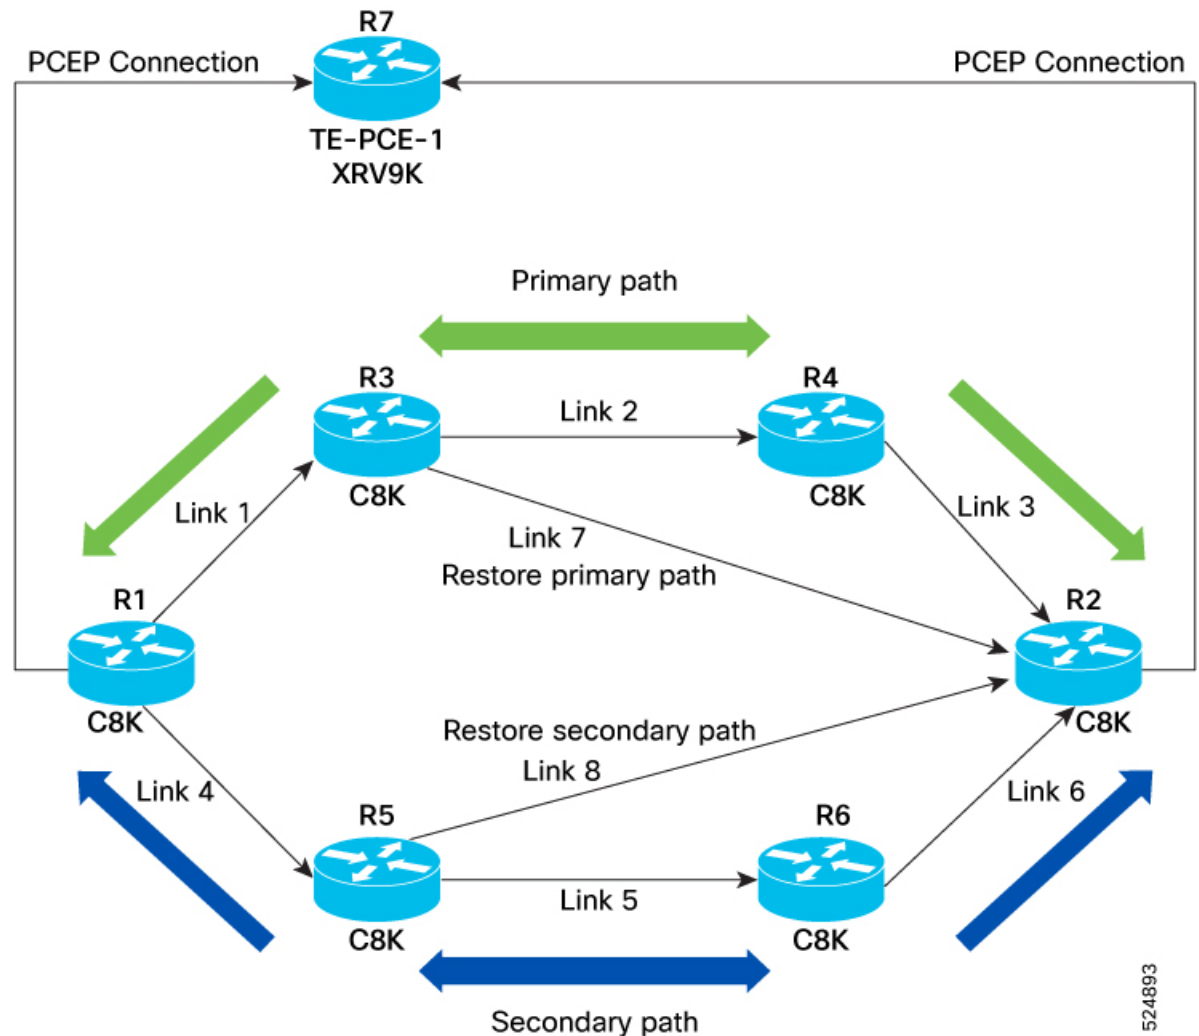
Summary

You can configure primary and secondary paths as explicit or dynamic path based on the network requirements. For explaining this feature, we are considering these key components:

- **Headend router:** R1
 - **Tailend router:** R2
 - **Mid-point routers:** R3, R4, R5, and R6
 - **PCE server:** R7 TE-PCE1, the Cisco IOS XRv 9000 platform serving as the centralized path computation element.
 - This feature provides the flexibility to connect multiple PCEs to the headend and tailend router for network redundancy.
 - To explain the workflow, we are prioritizing the PCE to perform flex LSP path computation, If the primary and secondary paths are broken, and If the PCE connection is broken then the headend router perform flex LSP path computation.
 - **Paths:**
 - Path 1 (Link 1-2-3): Working or active or primary path, explicitly configured path.
 - Path 2 (Link 4-5-6): Protecting or standby or secondary path, explicitly configured path.
 - Path 3 (Link 1-7): Restore primary path, dynamically calculated through the PCE server or headend router.
 - Path 4 (Link 4-8): Restore secondary path, dynamically calculated through the PCE server or headend router.
- The headend router and PCE use the shortest path first (SPF) algorithm to compute flex LSP paths, allowing them to potentially compute identical paths.
- **Link Capacities:** All links are 100G
 - **Weights:**
 - Link 1 to 6 have weights of 10,000
 - Link 7 and 8 have weights of 800

Workflow

Figure 2: Flex LSP path computation using PCE



These scenarios explain how PCE performs the flex LSP path computation:

When...	Then...
a primary path breaks	<p>Active path: The secondary path becomes active, which is Path 2.</p> <p>Standby path: The PCE computes a restore primary path, which is Path 3, by reusing the links from the failed primary path.</p>

When...	Then...
both primary and secondary path breaks	<p>Active path: The restore primary path becomes active, which is Path 3.</p> <p>Standby path: The PCE computes a restore secondary path, which is Path 4, by reusing the links from the failed secondary path.</p>
the broken link comes up again	<p>the primary or secondary path becomes active</p> <ul style="list-style-type: none"> • after reoptimization or • you can perform manual reoptimization using the mpls traffic-eng reoptimize command. <p>Reoptimization: The broken primary or secondary path becomes active after the configured reoptimization time limit. By default, the reoptimization time limit is 3600 seconds, but you can configure it between 0 to 604800 seconds using the reoptimize command under MPLS traffic engineering.</p>

In a scenario where the headend and tailend routers lose connection with the PCE, the headend router performs the flex LSP path computation. For more information, see [How flex LSP headend traffic flow works, on page 13](#).

Configure the RSVP-TE transport solution with flex LSP headend and PCE

Perform these steps on headend and tailend routers to configure the RSVP-TE transport solution with flex LSP headend and PCE:

Before you begin

- **Symmetric Configuration:** Ensure symmetric source and destination TE router IDs for bidirectional LSPs to be associated.
- **Tunnel Attributes:** Configure tunnel attributes identically on both sides of co-routed bidirectional LSPs.
- **Path Configuration:** You can configure primary and secondary paths as explicit or dynamic path based on the network requirements. For explaining this configuration, we are considering these key components:
 - **Path 10:** Explicit primary path protected by Paths 20, 50, and 60.
 - **Path 20:** Explicit secondary path protected by Paths 10, 30, and 40.
 - **Path 30:** Restore primary path dynamically computed by the PCE server.
 - **Path 40:** Restore primary path dynamically computed by the headend router.
 - **Path 50:** Restore secondary path dynamically computed by the PCE server.
 - **Path 60:** Restore secondary path dynamically computed by the headend router.
 - **Path Restoration Priority:**

- If both primary and secondary paths are broken, the PCE is prioritized to compute restored primary and secondary paths.
- If the connection of the headend and tailend routers with the PCE is broken, the headend router computes restored primary and secondary paths.
- The headend and tailend routers, along with the PCE, use the Shortest Path First (SPF) algorithm to compute flex LSP paths, potentially resulting in identical paths.
- Configure these setups:
 - **Bundle ethernet interfaces:** Configure bundle ethernet interfaces for efficient link aggregation.
 - **Loopback interface:** Set up a loopback interface for reliable routing.
 - **L2VPN configuration:** Define Pseudowire (PW) classes and point-to-point (P2P) cross-connects within an xconnect group.
 - **RSVP configuration:** Configure RSVP on PE devices to reserve bandwidth.
 - **Administrative weight:** Assign administrative weights to interfaces for path preference.
 - **MPLS LDP configuration:** Configure MPLS Label Distribution Protocol (LDP) on network devices for label switching.

Procedure

Step 1

If your network requires more than 1,000 tunnels, perform these steps:

- Scale up the routers with the **hw-module profile cef te-tunnel highscale-no-ldp-over-te** command to support a maximum of 8,000 RSVP-TE tunnels.

Example:

```
Router# configure
Router(config)# hw-module profile cef te-tunnel highscale-no-ldp-over-te
Router(config)# commit
```

- Reload all the LCs in the chassis with the **reload location all** command to enable the 8000 RSVP-TE tunnels.

Example:

```
Router# reload location all
```

Step 2

Configure primary path as explicit or dynamic path based on the network requirements. In this configuration, the primary LSP path is explicit.

Example:

```
Router# configure
Router(config)# explicit-path name R1-R2-Working-1
Router(config-exp-path)# index 1 next-address strict ipv4 unicast 101.1.3.17
Router(config-exp-path)# index 2 next-address strict ipv4 unicast 101.1.3.18
Router(config-exp-path)# index 3 next-address strict ipv4 unicast 101.3.4.17
Router(config-exp-path)# index 4 next-address strict ipv4 unicast 101.3.4.18
Router(config-exp-path)# index 5 next-address strict ipv4 unicast 101.2.4.6
Router(config-exp-path)# index 6 next-address strict ipv4 unicast 101.2.4.5
Router(config-exp-path)# exit
```

Step 3 Configure secondary path as explicit or dynamic path based on the network requirements. In this configuration, the secondary LSP path is explicit.

Example:

```
Router# configure
Router(config)# explicit-path name R1-R2-Protecting-1
Router(config-exp-path)# index 1 next-address strict ipv4 unicast 101.1.5.97
Router(config-exp-path)# index 2 next-address strict ipv4 unicast 101.1.5.98
Router(config-exp-path)# index 3 next-address strict ipv4 unicast 101.5.6.1
Router(config-exp-path)# index 4 next-address strict ipv4 unicast 101.5.6.2
Router(config-exp-path)# index 5 next-address strict ipv4 unicast 101.2.6.6
Router(config-exp-path)# index 6 next-address strict ipv4 unicast 101.2.6.5
Router(config-exp-path)# exit
```

Step 4 Configure a MPLS TE tunnel with bidirectional co-routed LSP, and enable wrap protection.

Example:

```
Router# configure
Router(config)# interface tunnel-te51
Router(config-if)# ipv4 unnumbered Loopback0
Router(config-if)# ipv6 enable
Router(config-if)# autoroute announce include-ipv6
Router(config-if)# destination 100.2.1.1
Router(config-if)# bidirectional
Router(config-if-bidir)# association id 51 source-address 100.1.1.1
Router(config-if-bidir)# association type co-routed
Router(config-if-bidir-co-routed)# wrap-protection
Router(config-if-bidir-co-routed)# exit
Router(config-if-bidir)# exit
```

Step 5 Configure path protection, path option for primary path, secondary path, and dynamic restore paths for primary and secondary.

You can configure up to four protected-by paths for the working and protecting paths, which serve as backup paths for the primary and secondary paths. In this configuration, only three protected-by paths are considered.

Example:

```
Router(config-if)# path-protection
Router(config-if)# path-option 10 explicit name R1-R2-Working-1 protected-by 20 50 60
Router(config-if)# path-option 20 explicit name R1-R2-Protecting-1 protected-by 10 30 40
Router(config-if)# path-option 30 dynamic attribute-set R1-R2-RestoreWorking-pce-1 protected-by 50
Router(config-if)# path-option 40 dynamic attribute-set R1-R2-RestoreWorking-distributed-1
protected-by 60
Router(config-if)# path-option 50 dynamic attribute-set R1-R2-RestoreProtecting-pce-1 protected-by
30
Router(config-if)# path-option 60 dynamic attribute-set R1-R2-RestoreProtecting-distributed-1
protected-by 40
```

Step 6 Perform the these steps to configure PCEP connection:

a) Configure PCE server to setup a centralized TE-PCE.

The Cisco IOS XRv 9000 is the recommended platform to act as the TE-PCE. For more information, see [Cisco IOS XRv 9000 Router Installation and Configuration Guide](#).

Example:

```
Router# configure
Router(config)# mpls traffic-eng
Router(config-mpls-te)# pce
Router(config-mpls-te-pce)# server ipv4 100.7.1.1
Router(config-mpls-te-pce)# commit
```

- b) Configure PCE peer on headend and tailend routers to enable the PCEP connection.

By default, the PCC timeout for a PC request is 10 seconds. You can modify this timeout to any value between 5 to 100 seconds using the **request-timeout** command.

Example:

```
Router# configure
Router(config)# mpls traffic-eng
Router(config-mpls-te)# pce
Router(config-mpls-te-pce)# peer ipv4 100.7.1.1
Router(config-mpls-te-pce)# request-timeout 100
Router(config-mpls-te-pce)# exit
Router(config-mpls-te)# commit
```

- Step 7** Configure attribute set for explicit primary and secondary LSP path.

Example:

```
Router# configure
Router(config)# mpls traffic-eng
Router(config-mpls-te)# attribute-set path-option R1-R2-RestoreWorking-distributed-1
Router(config-te-attribute-set)# biased-to explicit-path R1-R2-Working-1
Router(config-te-attribute-set)# commit
Router(config-mpls-te)# attribute-set path-option R1-R2-RestoreProtecting-distributed-1
Router(config-te-attribute-set)# biased-to explicit-path R1-R2-Protecting-1
Router(config-te-attribute-set)# commit
```

- Step 8** Configure attribute set for primary and secondary LSP path with PCE.

Example:

```
Router# configure
Router(config)# mpls traffic-eng
Router(config-mpls-te)# attribute-set path-option R1-R2-RestoreWorking-pce-1
Router(config-te-attribute-set)# pce
Router(config-te-attribute-set)# biased-to explicit-path R1-R2-Working-1
Router(config-te-attribute-set)# commit
Router(config-mpls-te)# attribute-set path-option R1-R2-RestoreProtecting-pce-1
Router(config-te-attribute-set)# pce
Router(config-te-attribute-set)# biased-to explicit-path R1-R2-Protecting-1
Router(config-te-attribute-set)# commit
```

- Step 9** Run the **sh mpls traffic-eng pce peer** command on the headend router to verify the PCEP connection with headend.

Example:

```
Router# sh mpls traffic-eng pce peer

Address Precedence State Learned From
-----
100.7.1.1 255 Up Static config
```

- Step 10** Run the **show mpls traffic-eng tunnels <tunnel id> detail** command on the headend router, to verify MPLS TE tunnel with explicit paths and dynamic restore paths for primary and secondary routes computed by the PCE server or headend routers.

These examples explain scenarios of how the PCE server compute dynamic restore paths for primary and secondary routes during a failure:

- a) In this example, Path 10 is the active path, and Path 20 is the standby path.

Example:

```

Router# show mpls traffic-eng tunnels 51 detail
Name: tunnel-te51 Destination: 100.2.1.1 Ifhandle:0x7800a684
  Signalled-Name: R1_t51
  Status:
    Admin:    up Oper:    up (Uptime 00:05:21)
    path option 10, type explicit R1-R2-Working-1 (Basis for Setup, path weight 30000 (reverse
30000))
      Accumulative metrics: TE 30000 (reverse 30000) IGP 30 (reverse 30) Delay 900000 (reverse
900000)
      Accumulative biased metrics: TE 30000 (reverse 30000) IGP 30 (reverse 30) Delay 900000
(reverse 900000)
      Protected-by PO index: 20,50,60
    path option 20, type explicit R1-R2-Protecting-1 (Basis for Standby, path weight 30000
(reverse 30000))
      Accumulative metrics: TE 30000 (reverse 30000) IGP 30 (reverse 30) Delay 900000 (reverse
900000)
      Accumulative biased metrics: TE 30000 (reverse 30000) IGP 30 (reverse 30) Delay 900000
(reverse 900000)
      Protected-by PO index: 10,30,40
    path option 30, type dynamic pce
      Protected-by PO index: 50
      Path-option attribute: R1-R2-RestoreWorking-pce-1
      PCE enabled
      Biased-to: explicit-path R1-R2-Working-1
    path option 40, type dynamic
      Protected-by PO index: 60
      Path-option attribute: R1-R2-RestoreWorking-distributed-1
      Biased-to: explicit-path R1-R2-Working-1
    path option 50, type dynamic pce
      Protected-by PO index: 30
      Path-option attribute: R1-R2-RestoreProtecting-pce-1
      PCE enabled
      Biased-to: explicit-path R1-R2-Protecting-1
    path option 60, type dynamic
      Protected-by PO index: 40
      Path-option attribute: R1-R2-RestoreProtecting-distributed-1
      Biased-to: explicit-path R1-R2-Protecting-1
*
*
Current LSP Info:
  Role: Working
  Instance: 11, Signaling Area: PCE computed path: server 100.7.1.1
  In-use path-option: 10
  Uptime: 00:05:21 (since Wed Mar 19 12:03:16 PDT 2025)
  Outgoing Interface: HundredGigE0/8/0/10, Outgoing Label: 24050
  Router-IDs: local      100.1.1.1
               downstream 100.3.1.1
  Soft Preemption: None
  SRLGs: not collected
  Path Info:
    Outgoing:
      Explicit Route:
        Strict, 101.1.3.18
        Strict, 101.3.4.18
        Strict, 101.2.4.5
        Strict, 100.2.1.1
*
*
Standby LSP Info:

```



```

Role: Protecting
Instance: 13, Signaling Area: PCE computed path: server 100.7.1.1
In-use path-option: 20
Uptime: 00:04:52 (since Wed Mar 19 12:03:45 PDT 2025), Oper State: Up
Outgoing Interface: HundredGigE0/8/0/0, Outgoing Label: 24051
Router-IDs: local      100.1.1.1
            downstream 100.5.1.1
Soft Preemption: None
SRLGs: not collected
Path Info:
  Outgoing:
    Explicit Route:
      Strict, 101.1.5.98
      Strict, 101.5.6.2
      Strict, 101.2.6.5
      Strict, 100.2.1.1

```

The above sample displays only a part of the actual output; the actual output displays more details.

b) In this example:

- The explicit primary path, Path 10, is broken.
- Path 20 becomes the active path, protected by Paths 10, 30, and 40.
- The PCE computes a restored primary path, which is Path 30, as a standby path by reusing the links from the original primary path.

Example:

```

Router# show mpls traffic-eng tunnels 51 detail
Name: tunnel-te51 Destination: 100.2.1.1 Ifhandle:0x7800a684
  Signalled-Name: R1_t51
  Status:
    Admin:      up Oper:      up (Uptime 00:09:56)

    path option 20, type explicit R1-R2-Protecting-1 (Basis for Setup, path weight 30000 (reverse
30000))
      Accumulative metrics: TE 30000 (reverse 30000) IGP 30 (reverse 30) Delay 900000 (reverse
900000)
      Accumulative biased metrics: TE 30000 (reverse 30000) IGP 30 (reverse 30) Delay 900000
(reverse 900000)
      Protected-by PO index: 10,30,40
    path option 30, type dynamic pce (Basis for Standby, path weight 100 (reverse 100))
      Accumulative metrics: TE 10100 (reverse 10100) IGP 20 (reverse 20) Delay 600000 (reverse
600000)
      Accumulative biased metrics: TE 100 (reverse 100) IGP 10 (reverse 10) Delay 300000 (reverse
300000)
      Protected-by PO index: 50
      Path-option attribute: R1-R2-RestoreWorking-pce-1
      PCE enabled
      Biased-to: explicit-path R1-R2-Working-1
*
*
Current LSP Info:
  Role: Working
  Instance: 16, Signaling Area: PCE computed path: server 100.7.1.1
  In-use path-option: 20
  Uptime: 00:00:41 (since Wed Mar 19 12:12:31 PDT 2025)
  Outgoing Interface: HundredGigE0/8/0/0, Outgoing Label: 24065
  Router-IDs: local      100.1.1.1
              downstream 100.5.1.1

```

```

Soft Preemption: None
SRLGs: not collected
Path Info:
  Outgoing:
    Explicit Route:
      Strict, 101.1.5.98
      Strict, 101.5.6.2
      Strict, 101.2.6.5
      Strict, 100.2.1.1
*
*
Standby LSP Info:
  Role: Protecting
  Instance: 18, Signaling Area: PCE computed path: server 100.7.1.1
  In-use path-option: 30
  Uptime: 00:00:27 (since Wed Mar 19 12:12:45 PDT 2025), Oper State: Up
  Outgoing Interface: HundredGigE0/8/0/10, Outgoing Label: 24061
  Router-IDs: local      100.1.1.1
              downstream 100.3.1.1
  Soft Preemption: None
  SRLGs: not collected
  Path Info:
    Outgoing:
      Explicit Route:
        Strict, 101.1.3.18
        Strict, 101.2.3.17
        Strict, 100.2.1.1

```

The above sample displays only a part of the actual output; the actual output displays more details.

c) In this example:

- The explicit secondary path, Path 20, is broken.
- Path 30 becomes the active path, protected by Path 50.
- The PCE computes a restored secondary path, which is Path 50, as a standby path by reusing the links from the original secondary path.

Example:

```

Router# show mpls traffic-eng tunnels 51 detail
Name: tunnel-te51 Destination: 100.2.1.1 Ifhandle:0x7800a684
  Signalled-Name: R1_t51
  Status:
    Admin:      up Oper:      up (Uptime 00:10:47)

    path option 30, type dynamic pce (Basis for Setup, path weight 100 (reverse 100))
      Accumulative metrics: TE 10100 (reverse 10100) IGP 20 (reverse 20) Delay 600000 (reverse 600000)
      Accumulative biased metrics: TE 100 (reverse 100) IGP 10 (reverse 10) Delay 300000 (reverse 300000)
      Protected-by PO index: 50
      Path-option attribute: R1-R2-RestoreWorking-pce-1
      PCE enabled
      Biased-to: explicit-path R1-R2-Working-1
    path option 50, type dynamic pce (Basis for Standby, path weight 100 (reverse 100))
      Accumulative metrics: TE 10100 (reverse 10100) IGP 20 (reverse 20) Delay 600000 (reverse 600000)
      Accumulative biased metrics: TE 100 (reverse 100) IGP 10 (reverse 10) Delay 300000 (reverse 300000)
      Protected-by PO index: 30

```

```
Path-option attribute: R1-R2-RestoreProtecting-pce-1
PCE enabled
Biased-to: explicit-path R1-R2-Protecting-1
*
*
Current LSP Info:
  Role: Working
  Instance: 13, Signaling Area: PCE computed path: server 100.7.1.1
  In-use path-option: 30
  Uptime: 00:00:42 (since Wed Mar 19 12:13:24 PDT 2025)
  Outgoing Interface: Bundle-Ether2301, Outgoing Label: 24044
  Router-IDs: local      100.2.1.1
               downstream 100.3.1.1
  Soft Preemption: None
  SRLGs: not collected
  Path Info:
    Outgoing:
      Explicit Route:
        Strict, 101.2.3.18
        Strict, 101.1.3.17
        Strict, 100.1.1.1
*
*

Standby LSP Info:
  Role: Protecting
  Instance: 14, Signaling Area: PCE computed path: server 100.7.1.1
  In-use path-option: 50
  Uptime: 00:00:21 (since Wed Mar 19 12:13:45 PDT 2025), Oper State: Up
  Outgoing Interface: Bundle-Ether2501, Outgoing Label: 24045
  Router-IDs: local      100.2.1.1
               downstream 100.5.1.1
  Soft Preemption: None
  SRLGs: not collected
  Path Info:
    Outgoing:
      Explicit Route:
        Strict, 101.2.5.18
        Strict, 101.1.5.97
        Strict, 100.1.1.1
```

The above sample displays only a part of the actual output; the actual output displays more details.

