

MPLS Traffic Engineering

This chapter provides conceptual and configuration information for the following MPLS-TE features:

- MPLS-TE Automatic Bandwidth
- MPLS-TE Fast Reroute (FRR)

- Overview of MPLS Traffic Engineering , on page 1
- MPLS-TE Scale Details , on page 2
- MPLS-TE Automatic Bandwidth, on page 2
- Configure Automatic Bandwidth, on page 5
- Fast Reroute, on page 9
- FRR Node Protection , on page 9
- Protecting MPLS Tunnels with Fast Reroute, on page 10
- Path Computation Client Initiated RSVP-TE, on page 13

Overview of MPLS Traffic Engineering

MPLS-TE software enables an MPLS backbone to replicate and expand upon the traffic engineering capabilities of Layer 2 ATM and Frame Relay networks. MPLS is an integration of Layer 2 and Layer 3 technologies. By making traditional Layer 2 features available to Layer 3, MPLS enables traffic engineering. Thus, you can offer in a one-tier network what now can be achieved only by overlaying a Layer 3 network on a Layer 2 network.

MPLS-TE is essential for service provider and Internet service provider (ISP) backbones. Such backbones must support a high use of transmission capacity, and the networks must be very resilient so that they can withstand link or node failures. MPLS-TE provides an integrated approach to traffic engineering. With MPLS, traffic engineering capabilities are integrated into Layer 3, which optimizes the routing of IP traffic, given the constraints imposed by backbone capacity and topology.

Benefits of MPLS-TE

MPLS-TE enables ISPs to route network traffic to offer the best service to their users in terms of throughput and delay. By making the service provider more efficient, traffic engineering reduces the cost of the network.

Currently, some ISPs base their services on an overlay model. In the overlay model, transmission facilities are managed by Layer 2 switching. The routers see only a fully meshed virtual topology, making most destinations appear one hop away. If you use the explicit Layer 2 transit layer, you can precisely control how

traffic uses available bandwidth. However, the overlay model has numerous disadvantages. MPLS-TE achieves the TE benefits of the overlay model without running a separate network and without a non-scalable, full mesh of router interconnects.

How MPLS-TE works

MPLS-TE automatically establishes and maintains label switched paths (LSPs) across the backbone by using RSVP. The path that an LSP uses is determined by the LSP resource requirements and network resources, such as bandwidth. Available resources are flooded by means of extensions to a link-state-based Interior Gateway Protocol (IGP).

MPLS-TE tunnels are calculated at the LSP headend router, based on a fit between the required and available resources (constraint-based routing). The IGP automatically routes the traffic to these LSPs.

MPLS-TE Scale Details

Scale details for MPLS-TE:

Table 1: Supported LSPs for MPLS-TE

MPLS TE with FRR	Head/Tail Node: 75000 LSPs Mid Node: 37500 LSPs
MPLS TE without FRR	Head/Tail Node: 75000 LSPs Mid Node: 75000 LSPs

MPLS-TE Automatic Bandwidth

The MPLS-TE automatic bandwidth feature measures the traffic in a tunnel and periodically adjusts the signaled bandwidth for the tunnel.

MPLS-TE Automatic Bandwidth Overview

MPLS-TE automatic bandwidth is configured on individual Label Switched Paths (LSPs) at every head-end. MPLS-TE monitors the traffic rate on a tunnel interface. Periodically, MPLS-TE resizes the bandwidth on the tunnel interface to align it closely with the traffic in the tunnel. MPLS-TE automatic bandwidth can perform these functions:

- Monitors periodic polling of the tunnel output rate
- Resizes the tunnel bandwidth by adjusting the highest rate observed during a given period

For every traffic-engineered tunnel that is configured for an automatic bandwidth, the average output rate is sampled, based on various configurable parameters. Then, the tunnel bandwidth is readjusted automatically based upon either the largest average output rate that was noticed during a certain interval, or a configured maximum bandwidth value.

This table lists the automatic bandwidth functions.

Table 2: Automatic Bandwidth Variables

Function	Command	Description	Default Value
Application frequency	**application** command	Configures how often the tunnel bandwidths changed for each tunnel. The application period is the period of A minutes between the bandwidth applications during which the output rate collection is done.	24 hours
Requested bandwidth	**bw-limit** command	Limits the range of bandwidth within the automatic-bandwidth feature that can request a bandwidth.	0 Kbps
Collection frequency	**auto-bw collect** command	Configures how often the tunnel output rate is polled globally for all tunnels.	5 min
Highest collected bandwidth	—	You cannot configure this value.	—
Delta	—	You cannot configure this value.	—

The output rate on a tunnel is collected at regular intervals that are configured by using the **application** command in MPLS-TE auto bandwidth interface configuration mode. When the application period timer expires, and when the difference between the measured and the current bandwidth exceeds the adjustment threshold, the tunnel is reoptimized. Then, the bandwidth samples are cleared to record the new largest output rate at the next interval.

When reoptimizing the LSP with the new bandwidth, a new path request is generated. If the new bandwidth is not available, the last good LSP continues to be used. This way, the network experiences no traffic interruptions.

If minimum or maximum bandwidth values are configured for a tunnel, the bandwidth, which the automatic bandwidth signals, stays within these values.



Note When more than 100 tunnels are **auto-bw** enabled, the algorithm will jitter the first application of every tunnel by a maximum of 20% (max 1hour). The algorithm does this to avoid too many tunnels running auto bandwidth applications at the same time.

If a tunnel is shut down, and is later brought again, the adjusted bandwidth is lost and the tunnel is brought back with the initial configured bandwidth. In addition, the application period is reset when the tunnel is brought back.

Adjustment Threshold

Adjustment Threshold is defined as a percentage of the current tunnel bandwidth and an absolute (minimum) bandwidth. Both thresholds must be fulfilled for the automatic bandwidth to resignal the tunnel. The tunnel

bandwidth is resized only if the difference between the largest sample output rate and the current tunnel bandwidth is larger than the adjustment thresholds.

For example, assume that the automatic bandwidth is enabled on a tunnel in which the highest observed bandwidth B is 30 Mbps. Also, assume that the tunnel was initially configured for 45 Mbps. Therefore, the difference is 15 mbit/s. Now, assuming the default adjustment thresholds of 10% and 10kbps, the tunnel is signalled with 30 Mbps when the application timer expires. This is because 10% of 45Mbit/s is 4.5 Mbit/s, which is smaller than 15 Mbit/s. The absolute threshold, which by default is 10kbps, is also crossed.

Overflow Detection

Overflow detection is used if a bandwidth must be resized as soon as an overflow condition is detected, without having to wait for the expiry of an automatic bandwidth application frequency interval.

For overflow detection one configures a limit N, a percentage threshold Y% and optionally, a minimum bandwidth threshold Z. The percentage threshold is defined as the percentage of the actual signalled tunnel bandwidth. When the difference between the measured bandwidth and the actual bandwidth are both larger than Y% and Z threshold, for N consecutive times, then the system triggers an overflow detection.

The bandwidth adjustment by the overflow detection is triggered only by an increase of traffic volume through the tunnel, and not by a decrease in the traffic volume. When you trigger an overflow detection, the automatic bandwidth application interval is reset.

By default, the overflow detection is disabled and needs to be manually configured.

Underflow Detection

Underflow detection is used when the bandwidth on a tunnel drops significantly, which is similar to overflow but in reverse.

Underflow detection applies the highest bandwidth value from the samples which triggered the underflow. For example, if you have an underflow limit of three, and the following samples trigger the underflow for 10 kbps, 20 kbps, and 15 kbps, then, 20 kbps is applied.

Unlike overflow, the underflow count is not reset across an application period. For example, with an underflow limit of three, you can have the first two samples taken at the end of an application period and then the underflow gets triggered by the first sample of the next application period.

Restrictions for MPLS-TE Automatic Bandwidth

When the automatic bandwidth cannot update the tunnel bandwidth, the following restrictions are listed:

- Tunnel is in a fast reroute (FRR) backup, active, or path protect active state. This occurs because of the assumption that protection is a temporary state, and there is no need to reserve the bandwidth on a backup tunnel. You should prevent taking away the bandwidth from other primary or backup tunnels.
- Reoptimization fails to occur during a lockdown. In this case, the automatic bandwidth does not update the bandwidth unless the bandwidth application is manually triggered by using the **mpls traffic-eng auto-bw apply** command in EXEC mode.

Configure Automatic Bandwidth

Configuring automatic bandwidth involves the following tasks:

- Configuring Collection Frequency
- Forcing the current application period to expire immediately
- Configuring the automatic bandwidth functions

Configure Collection Frequency

Perform this task to configure the collection frequency. You can configure only one global collection frequency.

Procedure

Step 1 **configure**

Step 2 **mpls traffic-eng**

Example:

```
RP/0/(config)# mpls traffic-eng
RP/0/(config-mpls-te)#
```

Enters MPLS-TE configuration mode.

Step 3 **auto-bw collect frequency** *minutes*

Example:

```
RP/0/(config-mpls-te)# auto-bw collect frequency 1
```

Configures the automatic bandwidth collection frequency, and controls the manner in which the bandwidth for a tunnel collects output rate information; but does not adjust the tunnel bandwidth.

minutes

Configures the interval between automatic bandwidth adjustments in minutes. Range is from 1 to 10080.

Step 4 **commit**

Forcing the Current Application Period to Expire Immediately

Perform this task to force the current application period to expire immediately on the specified tunnel. The highest bandwidth is applied on the tunnel before waiting for the application period to end on its own.

Procedure

Step 1 **mpls traffic-eng auto-bw apply** {**all** | **tunnel-te** *tunnel-number*}

Example:

```
RP/0/# mpls traffic-eng auto-bw apply tunnel-te 1
```

Configures the highest bandwidth available on a tunnel without waiting for the current application period to end.

all

Configures the highest bandwidth available instantly on all the tunnels.

tunnel-te

Configures the highest bandwidth instantly to the specified tunnel. Range is from 0 to 65535.

Step 2 **commit**

Step 3 **show mpls traffic-eng tunnels** [**auto-bw**]

Example:

```
RP/0/# show mpls traffic-eng tunnels auto-bw
```

Displays information about MPLS-TE tunnels for the automatic bandwidth.

Configure Automatic Bandwidth Functions

Perform this task to configure the following automatic bandwidth functions:

Application frequency

Configures the application frequency in which a tunnel bandwidth is updated by the automatic bandwidth.

Bandwidth collection

Configures only the bandwidth collection.

Bandwidth parameters

Configures the minimum and maximum automatic bandwidth to set on a tunnel.

Adjustment threshold

Configures the adjustment threshold for each tunnel.

Overflow detection

Configures the overflow detection for each tunnel.

Procedure

Step 1 **configure**

Step 2 **interface tunnel-te** *tunnel-id*

Example:

```
RP/0/(config)# interface tunnel-te 6
RP/0/(config-if)#
```

Configures an MPLS-TE tunnel interface and enables traffic engineering on a particular interface on the originating node.

Step 3 **auto-bw**

Example:

```
RP/0/(config-if)# auto-bw
RP/0/(config-if-tunte-autobw)#
```

Configures automatic bandwidth on a tunnel interface and enters MPLS-TE automatic bandwidth interface configuration mode.

Step 4 **application** *minutes*

Example:

```
RP/0/(config-if-tunte-autobw)# application 1000
```

Configures the application frequency in minutes for the applicable tunnel.

minutes

Frequency in minutes for the automatic bandwidth application. Range is from 5 to 10080 (7 days). The default value is 1440 (24 hours).

Step 5 **bw-limit** {**min** *bandwidth* } {**max** *bandwidth*}

Example:

```
RP/0/(config-if-tunte-autobw)# bw-limit min 30 max 80
```

Configures the minimum and maximum automatic bandwidth set on a tunnel.

min

Applies the minimum automatic bandwidth in kbps on a tunnel. Range is from 0 to 4294967295.

max

Applies the maximum automatic bandwidth in kbps on a tunnel. Range is from 0 to 4294967295.

Step 6 **adjustment-threshold** *percentage* [**min** *minimum-bandwidth*]

Example:

```
RP/0/(config-if-tunte-autobw)# adjustment-threshold 50 min 800
```

Configures the tunnel bandwidth change threshold to trigger an adjustment.

percentage

Bandwidth change percent threshold to trigger an adjustment if the largest sample percentage is higher or lower than the current tunnel bandwidth. Range is from 1 to 100 percent. The default value is 5 percent.

min

Configures the bandwidth change value to trigger an adjustment. The tunnel bandwidth is changed only if the largest sample is higher or lower than the current tunnel bandwidth. Range is from 10 to 4294967295 kilobits per second (kbps). The default value is 10 kbps.

Step 7 **overflow threshold** *percentage* [**min** *bandwidth*] **limit** *limit*

Example:

```
RP/0/(config-if-tunte-autobw)# overflow threshold 100 limit 1
```

Configures the tunnel overflow detection.

percentage

Bandwidth change percent to trigger an overflow. Range is from 1 to 100 percent.

limit

Configures the number of consecutive collection intervals that exceeds the threshold. The bandwidth overflow triggers an early tunnel bandwidth update. Range is from 1 to 10 collection periods. The default value is none.

min

Configures the bandwidth change value in kbps to trigger an overflow. Range is from 10 to 4294967295. The default value is 10.

Step 8 **commit**

Fast Reroute

Table 3: Feature History

Feature Name	Release Information	Feature Description
Traffic Engineering (TE) over LAG	Cisco IOS XR Release 6.5.31	This feature allows the MPLS-TE tunnels to be protected with Fast Reroute (FRR) for interfaces on the LAG. If the LSPs in the MPLS-TE tunnel encounter a failed link, FRR reroutes the traffic carried by the LSPs. Commands added: • show mpls traffic-eng fast-reroute database • show mpls traffic-eng fast-reroute log

Fast Reroute (FRR) provides link protection to LSPs enabling the traffic carried by LSPs that encounter a failed link to be rerouted around the failure. The reroute decision is controlled locally by the router connected to the failed link. The headend router on the tunnel is notified of the link failure through IGP or through RSVP. When it is notified of a link failure, the headend router attempts to establish a new LSP that bypasses the failure. This provides a path to reestablish links that fail, providing protection to data transfer.

You should be aware of these requirements for the backup tunnel path

- Backup tunnel must not pass through the element it protects.
- Primary tunnel and a backup tunnel should intersect at least at two points (nodes) on the path: point of local repair (PLR) and merge point (MP). PLR is the headend of the backup tunnel, and MP is the tailend of the backup tunnel.

Note When you configure TE tunnel with multiple protection on its path and merge point is the same node for more than one protection, you must configure record-route for that tunnel.

FRR Node Protection

If a link failure occurs within an area, the upstream router directly connected to the failed link generates an RSVP path error message to the headend. As a response to the message, the headend sends an RSVP path tear message and the corresponding path option is marked as invalid for a specified period and the next path-option (if any) is evaluated.

To retry the ABR immediately, a second path option (identical to the first one) should be configured. Alternatively, the retry period (path-option hold-down, 2 minutes by default) can be tuned to achieve a faster retry.

Protecting MPLS Tunnels with Fast Reroute

From R6.5.3.1, the MPLS tunnels can be protected with Fast Reroute for LAG interfaces.

Before you begin

The following prerequisites are required to protect MPLS-TE tunnels:

- You must have a router ID for the neighboring router.
- Stable router ID is required at either end of the link to ensure that the link is successful. If you do not assign a router ID to the routers, the system defaults to the global router ID. Default router IDs are subject to change, which can result in an unstable link.
- You must first configure a primary tunnel.

Procedure

Step 1 **configure**

Step 2 **interface tunnel-te** *tunnel-id*

Example:

```
RP/0/# interface tunnel-te 1
```

Configures an MPLS-TE tunnel interface.

Step 3 **fast-reroute**

Example:

```
RP/0/(config-if)# fast-reroute
```

Enables fast reroute.

Step 4 **exit**

Example:

```
RP/0/(config-if)# exit
```

Exits the current configuration mode.

Step 5 **mpls traffic-eng**

Example:

```
RP/0/(config)# mpls traffic-eng
```

```
RP/0/(config-mpls-te)#
```

Enters MPLS-TE configuration mode.

Step 6 **reoptimize timers delay {cleanup** *delay-time* **| installation** *delay-time***}**

Example:

```
RP/0/(config-mpls-te)# reoptimize timers delay cleanup 180
RP/0/(config-mpls-te)# reoptimize timers delay installation 180
```

Delays removal of the old LSPs and installation of a new label after tunnel reoptimization. The minimum installation and cleanup time is 180 seconds.

Step 7 **interface** *type interface-path-id*

Example:

```
RP/0/(config-mpls-te)# interface TenGigE0/1/0/3
RP/0/(config-mpls-te-if)#
```

Example:

```
RP/0/(config-mpls-te)# interface bundle-ether 150
RP/0/(config-mpls-te-if)#
```

Enables traffic engineering on a particular interface on the originating node. From R6.5.3.1, you can also enable traffic engineering on LAG interface.

Step 8 **backup-path tunnel-te** *tunnel-number*

Example:

```
RP/0/(config-mpls-te-if)# backup-path tunnel-te 2
```

Sets the backup path to the backup tunnel.

Step 9 **exit**

Example:

```
RP/0/(config-mpls-te-if)# exit
RP/0/(config-mpls-te)#
```

Exits the current configuration mode.

Step 10 **exit**

Example:

```
RP/0/(config-mpls-te)# exit
RP/0/(config)#
```

Exits the current configuration mode.

Step 11 **interface tunnel-te** *tunnel-id*

Example:

```
RP/0/(config)# interface tunnel-te 2
```

Configures an MPLS-TE tunnel interface.

Step 12 **ipv4 unnumbered** *type interface-path-id*

Example:

```
RP/0/(config-if)# ipv4 unnumbered Loopback0
```

Assigns a source address to set up forwarding on the new tunnel.

Step 13 **path-option** *preference-priority* {**explicit name** *explicit-path-name*}

Example:

```
RP/0/(config-if)# path-option 1 explicit name backup-path
```

Sets the path option to explicit with a given name (previously configured) and assigns the path ID.

Step 14 **destination** *ip-address*

Example:

```
RP/0/(config-if)# destination 192.168.92.125
```

Assigns a destination address on the new tunnel.

- Destination address is the remote node's MPLS-TE router ID.
- Destination address is the merge point between backup and protected tunnels.

Note
When you configure TE tunnel with multiple protection on its path and merge point is the same node for more than one protection, you must configure record-route for that tunnel.

Step 15 **commit**

Step 16 (Optional) **show mpls traffic-eng tunnels backup**

Example:

```
RP/0/# show mpls traffic-eng tunnels backup
```

Displays the backup tunnel information.

Step 17 (Optional) **show mpls traffic-eng tunnels protection frr**

Example:

```
RP/0/# show mpls traffic-eng tunnels protection frr
```

Displays the tunnel protection information for Fast-Reroute (FRR).

Step 18 (Optional) **show mpls traffic-eng fast-reroute database**

Example:

```
RP/0/# show mpls traffic-eng fast-reroute database
```

Displays the protected tunnel state (for example, the tunnel's current ready or active state).

Step 19 (Optional) **show mpls traffic-eng fast-reroute log**

Example:

```
RP/0/# show mpls traffic-eng fast-reroute log
```

Displays the log of FRR events.

Path Computation Client Initiated RSVP-TE

Table 4: Feature History

Feature Name	Release Information	Feature Description
Path Computation Client Initiated RSVP-TE	Cisco IOS XR Release 6.5.31	This feature establishes Path Computation Element Communication Protocol (PCEP) between PCE (NCS 5500) and Path Computation Client (PCC) (NCS 4000) and creates RSVP-TE tunnels between the head end node (PCC) and a tail end node (another NCS 4000 device). It supports client such as Cisco Optimization Engine (COE) to preview the RSVP-TE path initiated by PCC, before deployment, thereby supporting Bandwidth on Demand (BWoD). Commands added: • show pce ipv4 • show pce lps • show mpls traffic-eng pce peer • show mpls traffic-eng pce lsp-database

Cisco IOS-XR Path Computation Element (PCE) collects network topology through IGP and/or BGP-LS, and provides path computation services for RSVP-TE tunnels. Also, it supports any external client (for example, Cisco Crosswork Optimization Engine (COE)) to deploy RSVP-TE tunnels based on the client's needs. COE obtains services such as topology collection, path computation, and RSVP-TE deployment services from PCE, to support Bandwidth on Demand (BWoD) and Bandwidth Optimization (BWOpt) applications.

PCE describes a set of procedures by which a Path Computation Client (PCC) can report and delegate control of head end tunnels that are sourced from the PCC to a PCE peer. The PCC and PCE establish a PCE Communication Protocol (PCEP) connection that PCE uses to push updates to the network.

About RSVP-TE tunnels, PCE support is restricted to disjoint path computation (node, link, SRLG). Also, RSVP-TE tunnels reports to PCE for discovery purpose.

From Release 7.3.1, the following features in PCE and PCC support clients such as COE, WAE, or any third-party tools:

- Clients can discover RSVP-TE tunnels that are delegated or simply reported to PCE.
- Neither PCE nor client can modify the path of a nondelegated tunnel.
- PCE always dynamically computes the path of a delegated tunnel that is initiated by PCC, and clients cannot modify the paths.
- Clients can preview an RSVP-TE path before deployment. Clients may also choose not to deploy that tunnel.

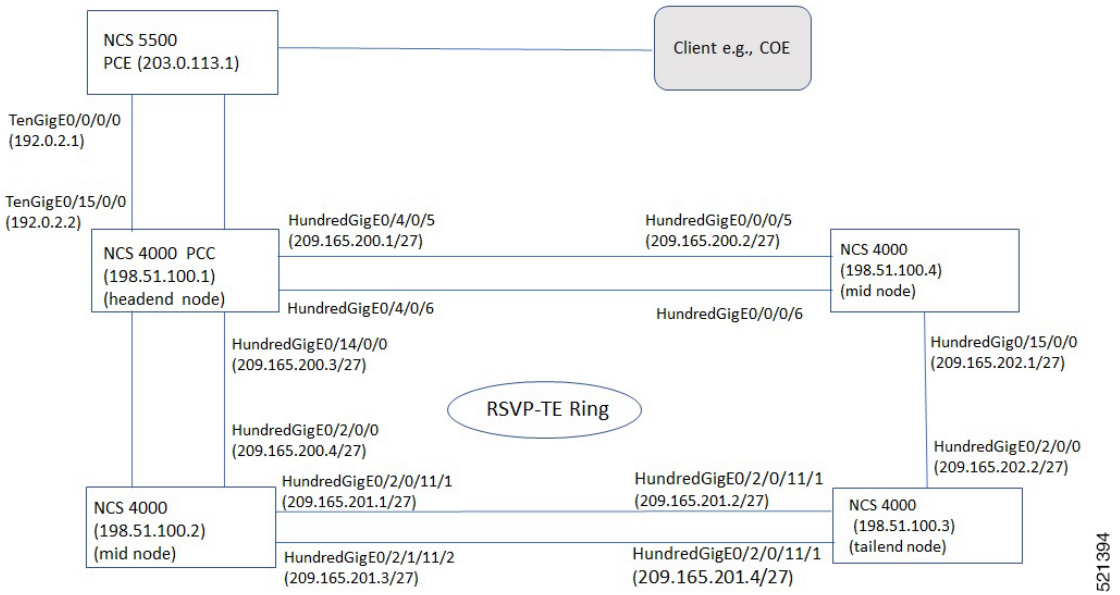
PCC runs on NCS 4000 and the PCE runs on the platforms such as ASR 9000, XRv9000, and NCS 5500 where the software should be 6.6.3+optima1.1SMU or higher, up where the PCE support for optima1.1 exists.

Limitations

PCE supports low latency, low cost, disjoint path computation with affinity and bandwidth constraints. Affinity and disjoint constraints are not supported.

Use Case - PCC-Initiated RSVP-TE for BWoD

The following topology explains the workflow for initiating delegated RSVP-TE tunnel for the BWoD application:



The topology has four NCS 4000 nodes for redundancy and one NCS 5500 node (PCE). The headend node (198.51.100.1) is connected to PCE (203.0.113.1) through the interface. Perform the following steps to create an RSVP-TE tunnel between the headend node (198.51.100.1) and the tailend node (198.51.100.3).

Procedure

Step 1 Check whether the IS-IS interfaces of the NCS 4000 nodes (headend, mid node, tailend) are up and running using the following command:

```
RP/0/RP0:NCS4016-1#show ip interface brief
Tue Feb 9 12:14:34.807 IST

Interface                  IP-Address       Status      Protocol      Vrf-Name
Bundle-Ether12             unassigned       Down        Down          default
Loopback5000               198.51.100.1          Up          Up            default
HundredGigE0/0/0/5         unassigned       Down        Down          default
HundredGigE0/0/0/5.100     85.1.1.1         Down        Down          default
HundredGigE0/0/0/10/1      unassigned       Down        Down          default
HundredGigE0/2/0/5         unassigned       Up          Up            default
HundredGigE0/2/0/5.100     6198.51.100.1         Up          Up            default
HundredGigE0/4/0/5         20.20.20.2       Up          Up            default
HundredGigE0/4/0/5.100     14.1.1.1         Up          Up            default
HundredGigE0/4/0/6         unassigned       Down        Down          default
HundredGigE0/13/0/10/1     unassigned       Up          Up            default
HundredGigE0/14/0/0        17.0.0.1         Up          Up            default
HundredGigE0/14/0/0.100    12.1.2.1         Up          Up            default
HundredGigE0/14/0/11/1     13.1.1.1         Down        Down          default
HundredGigE0/15/0/0        unassigned       Shutdown    Down          default
FortyGigE0/15/0/8          24.0.0.1         Up          Up            default
TenGigE0/15/0/5/1          192.168.1.1      Up          Up            default
TenGigE0/15/0/5/2          unassigned       Shutdown    Down          default
TenGigE0/15/0/5/3          unassigned       Shutdown    Down          default
TenGigE0/15/0/5/4          unassigned       Shutdown    Down          default
MgmtEth0/RP0/CPU0/0        10.58.230.71     Up          Up            default
MgmtEth0/RP0/EMS/0         unassigned       Up          Up            default
MgmtEth0/RP0/CRAFT/0       unassigned       Shutdown    Down          default
MgmtEth0/RP1/CPU0/0        10.58.230.69     Shutdown    Down          default
MgmtEth0/RP1/EMS/0         unassigned       Shutdown    Down          default
MgmtEth0/RP1/CRAFT/0       unassigned       Shutdown    Down          default
RP/0/RP0:NCS4016-1#
```

Step 2 If any interface is down, use the following command to make the interface up and running:

```
RP/0/RP1:NCS4016-1#configure
RP/0/RP1:NCS4016-1(config)#controller optics 0/15/0/1
RP/0/RP1:NCS4016-1(config-Optics)#no shutdown
RP/0/RP1:NCS4016-1(config-Optics)#commit
```

Step 3 Configure the NCS 5500 for PCE:

Table 5:

Configuration	Commands
1. Assign IP address to the interface and configure static routing with the NCS 4000 headend node.	interface Loopback0 ipv4 address 203.0.113.1 255.255.255.255 ! interface MgmtEth0/RP0/CPU0/0 ipv4 address 10.58.230.130 255.255.0.interface TenGigE0/0/0/0 ipv4 address 192.0.2.1 255.255.255.0 ! router static address-family ipv4 unicast 0.0.0.0/0 10.58.228.1 198.51.100.1/32 192.0.2.2
2. Configure iBGP protocol	router bgp 1 bgp router-id 203.0.113.1 address-family ipv4 unicast network 203.0.113.1/32 ! address-family link-state link-state ! neighbor 198.51.100.1 remote-as 1 update-source Loopback0 address-family ipv4 unicast ! address-family link-state link-state
3.Configure PCE	pce address ipv4 203.0.113.1 api user cisco password encrypted 00071A15075 ! ! timers minimum-peer-keepalive 0

Step 4 Configure the NCS 4000 headend and tailend nodes:

Table 6:

Configuration	NCS 4000 (Headend node - PCC)	NCS 4000 (Tailend node)

1. Configure the Interface	Loopback Configuration
	`interface Loopback5000`
	` ipv4 address 198.51.100.1 255.255.255.255`
	Headend to mid node1 configuration
	`interface HundredGigE0/14/0/0`
	` ipv4 address 209.165.200.3 255.255.255.0`
	` load-interval 30`
	Headend to PCE configuration
	`interface TenGigE0/15/0/5/1`
	` ipv4 address 192.0.2.2 255.255.255.0`
	Headend to mid node2 configuration
	`interface HundredGigE0/4/0/5`
	` ipv4 address 209.165.200.1 255.255.255.0`
	` load-interval 30`
		Loopback Configuration
		`interface Loopback9000`
		` ipv4 address 198.51.100.3 255.255.255.255`
		Tailend to mid node2 configuration
		`interface HundredGigE0/2/0/11/1`
		` mtu 9600`
		` ipv4 address 209.165.201.2 0.3255.255.255.252`
		` load-interval 30`
		Tailend to mid node2 configuration
		`interface HundredGigE0/2/0/0`
		` mtu 9600`
		` ipv4 address 209.165.202.2 255.255.255.0`
		` load-interval 30`

2. Configure IGP (IS-IS)	`router isis 100`
	` is-type level-2-only`
	` net 49.2001.1000.0100.1001.00`
	` nsr`
	` distribute link-state`
	` nsf cisco`
	` log adjacency changes`
	` address-family ipv4 unicast`
	`  metric-style wide`
	`  mpls traffic-eng level-2-only`
	`  mpls traffic-eng router-id Loopback5000`
	`  router-id 198.51.100.1`
	` !`
	` interface Loopback5000`
	`  passive`
	`  address-family ipv4 unicast`
	`  !`
	` !`
	`interface HundredGigE0/4/0/5`
	`  circuit-type level-2-only`
	`  point-to-point`
	`  address-family ipv4 unicast`
	`  !`
	` !`
	` interface HundredGigE0/14/0/0`
	`  point-to-point`
	`  address-family ipv4 unicast`
	`  !`
	` !`
	` interface HundredGigE0/14/0/0`
	`  point-to-point`
	`  address-family ipv4 unicast`
		`router isis 100`
		` is-type level-2-only`
		` net 47.0001.0000.0000.0009.00`
		` nsr`
		` distribute link-state level 2`
		` nsf cisco`
		` log adjacency changes`
		` address-family ipv4 unicast`
		`  metric-style wide`
		`  mpls traffic-eng level-2-only`
		`  mpls traffic-eng router-id Loopback9000`
		` !`
		` interface Loopback9000`
		`  address-family ipv4 unicast`
		`  !`
		` !`
		` interface HundredGigE0/2/0/0`
		`  point-to-point`
		`  address-family ipv4 unicast`
		`!`
		` interface HundredGigE0/2/0/11/1`
		`  point-to-point`
		`  address-family ipv4 unicast`
		`  !`
		`  !`
		`!`

Configuration	NCS 4000 (Headend node - PCC)	NCS 4000 (Tailend node)
3. Configure BGP (iBGP)	router bgp 1 bgp router-id 198.51.100.1 address-family ipv4 unicast ! address-family link-state link-state ! neighbor 198.51.100.3 remote-as 1 update-source Loopback5000 address-family ipv4 labeled-unicast route-reflector-client next-hop-self ! ! neighbor 203.0.113.1 remote-as 1 update-source Loopback5000 address-family ipv4 unicast ! address-family link-state link-state ! ! neighbor 11.11.11.11 remote-as 1 update-source Loopback5000 address-family ipv4 unicast ! address-family link-state link-state ! ! !	router bgp 1 bgp router-id 198.51.100.3 ibgp policy out enforce-modifications address-family ipv4 unicast allocate-label all ! address-family link-state link-state ! neighbor 198.51.100.1 remote-as 1 update-source Loopback9000 address-family ipv4 labeled-unicast route-reflector-client next-hop-self ! address-family link-state link-state ! ! !
4. Configure RSVP	rsvp interface HundredGigE0/0/0/5 bandwidth percentage 99 ! interface HundredGigE0/4/0/5 bandwidth percentage 99 ! interface HundredGigE0/14/0/0 bandwidth percentage 99 ! latency threshold 100 !	rsvp interface HundredGigE0/0/0/5 bandwidth percentage 99 ! interface HundredGigE0/15/0/0 bandwidth percentage 99 ! !
5. Configure MPLS-TE	mpls traffic-eng interface HundredGigE0/0/0/5 ! interface HundredGigE0/4/0/5 auto-tunnel backup ! ! interface HundredGigE0/14/0/0 auto-tunnel backup ! !	mpls traffic-eng interface HundredGigE0/2/0/0 ! interface HundredGigE0/2/0/11/1 ! fault-oam signalling advertise explicit-null path-selection metric igp ! !

Configuration	NCS 4000 (Headend node - PCC)	NCS 4000 (Tailend node)
6. Establish PCEP session with PCE	pce peer source ipv4 198.51.100.1 peer ipv4 203.0.113.1 precedence 10 ! peer ipv4 11.11.11.11 precedence 20 ! logging events peer-status stateful-client instantiation report timers state-timeout 600 redundancy pcc-centric ! ! auto-tunnel pcc tunnel-id min 5000 max 7000 ! fault-oam signalling advertise explicit-null path-selection metric igp ! !	Not applicable for tail end configuration.
7. MPLS-TE TUNNEL Configuration	interface tunnel-te300 description PCEP-TEST bandwidth 100 destination 198.51.100.3 fast-reroute protect bandwidth path-protection protection-mode non-revertive ! path-option 1 dynamic pce delegation ! ! !	Not applicable for tail end configuration.

Step 5 Configure the NCS 4000 mid nodes:

Table 7:

Configuration	NCS 4000 (mid node 1)	NCS 4000 (mid node 2)
Interface	interface Loopback8000 ipv4 address 198.51.100.2 255.255.255.255 ! interface HundredGigE0/2/0/0 mtu 9600 ipv4 address 209.165.200.4 255.255.255.0 load-interval 30 ! interface HundredGigE0/2/0/11/1 mtu 9600 ipv4 address 209.165.201.1 255.255.255.0 load-interval 30 !	interface Loopback7000 ipv4 address 198.51.100.4 255.255.255.255 ! interface HundredGigE0/0/0/5 ipv4 address 209.165.200.2 255.255.255.0 load-interval 30 ! interface HundredGigE0/15/0/0 mtu 9600 ipv4 address 209.165.202.1 255.255.255.0 !
IGP (IS-IS)	router isis 100 is-type level-2-only net 47.0001.0000.0000.0008.00 nsr distribute link-state level 2 nsf cisco log adjacency changes address-family ipv4 unicast metric-style wide mpls traffic-eng level-2-only mpls traffic-eng router-id Loopback8000 ! interface Loopback8000 address-family ipv4 unicast ! ! interface HundredGigE0/2/0/0 point-to-point address-family ipv4 unicast	router isis 100 is-type level-2-only net 47.0001.0000.0000.0010.00 nsr nsf cisco log adjacency changes address-family ipv4 unicast metric-style wide mpls traffic-eng level-2-only mpls traffic-eng router-id Loopback7000 ! interface Loopback7000 address-family ipv4 unicast ! ! interface HundredGigE0/0/0/5 point-to-point address-family ipv4 unicast ! ! interface HundredGigE0/15/0/0 point-to-point address-family ipv4 unicast ! ! !
RSVP	rsvp interface HundredGigE0/2/0/0 bandwidth percentage 99 ! interface HundredGigE0/2/0/11/1 bandwidth percentage 99 !	rsvp interface HundredGigE0/0/0/5 bandwidth percentage 99 ! interface HundredGigE0/15/0/0 bandwidth percentage 99 ! !

Configuration	NCS 4000 (mid node 1)	NCS 4000 (mid node 2)
MPLS TE	mpls traffic-eng interface HundredGigE0/2/0/0 ! interface HundredGigE0/2/0/11/1 ! fault-oam signalling advertise explicit-null path-selection metric igp !	mpls traffic-eng interface HundredGigE0/0/0/5 ! interface HundredGigE0/15/0/0 ! fault-oam signalling advertise explicit-null path-selection metric igp ! !

Step 6 Log in to a server installed with CURL and execute the following command, so that PCE initiates the RSVP-TE tunnel provisioning:

```
bash-4.2$ curl --raw -vN "http://cisco:cisco@10.77.142.23:8080/lsp/create/
simple?allow-xtc-reoptimization=1&name=l&source=198.51.100.1&destination=198.51.100.3&peer=198.51.100.1&metric-latency=20&type=rsvp"
* About to connect() to 10.77.142.23 port 8080 (#0)
* Trying 10.77.142.23...
* Connected to 10.77.142.23 (10.77.142.23) port 8080 (#0)
* Server auth using Basic with user 'cisco'
> GET
/lsp/create/simple?allow-xtc-reoptimization=1&name=l&source=198.51.100.1&destination=198.51.100.3
&peer=198.51.100.1&metric-latency=20&type=rsvp HTTP/1.1
> Authorization: Basic Y2lzY286Y2lzY28=
> User-Agent: curl/7.29.0
Host: 10.77.142.23:8080
Accept: */*
>
< HTTP/1.1 200 OK
< Cache-Control: no-cache, no-store
< Content-Type: text/json; charset=utf-8
< Expires: -1
< Transfer-Encoding: chunked
< Connection: keep-alive
<
30
create-lsp "l" (rsvp) on peer 198.51.100.1 (Success)
0
* Connection #0 to host 10.77.142.23 left intact
```

Step 7 Verify if the RSVP-TE tunnel is created, using the **show mpls traffic-eng** command:

```
RP/0/RP1:NCS4016-1#show mpls traffic-eng tunnels tabular
Tunnel             LSP    Destination     Source        Tun      FRR         LSP    Path
Name               ID     Address         Address       State    State       Role   Prot
--------------     -----  ------------ --------------- ------ ------ ---- -----
tunnel-te1          2     198.51.100.3    198.51.100.1    up     Ready    Headend   Inact
*tunnel-te8240      4     198.51.100.2    198.51.100.1    up     Inact    Headend   Inact
*tunnel-te8260      0     198.51.100.3    0.0.0.0         down   Inact    Headend   Inact
NCS4016-3 tl        6     198.51.100.1    198.51.100.3    up     Inact    Tailend
Autob NCS4009-2 t   2     198.51.100.1    198.51.100.2    up     Inact    Tailend
*= automatically created backup tunnel
```

Step 8 Verify the detailed information of the RSVP-TE tunnel, using the **show mpls traffic-eng** command:

```
   RP/0/RP0:NCS4016-1#show mpls traffic-eng tunnels 300
Name: tunnel-te300  Destination: 198.51.100.3  Ifhandle:0x8800584
  Signalled-Name: PCEP-TEST
```

```
  Status:
    Admin:    up Oper:   up   Path:  valid   Signalling: connected
    path option 10, (verbatim) type explicit (autopcc_te300) (Basis for Setup)
    G-PID: 0x0800 (derived from egress interface properties)
    Bandwidth Requested: 0 kbps  CT0
    Creation Time: Thu Jul  2 12:28:37 2020 (4w0d ago)
  Config Parameters:
    Bandwidth:        0 kbps (CT0) Priority:  7  7 Affinity: 0x0/0xffff
    Metric Type: IGP (interface)
    Path Selection:
      Tiebreaker: Min-fill (default)
    Hop-limit: disabled
    Cost-limit: disabled
    Delay-limit: disabled
    Path-invalidation timeout: 10000 msec (default), Action: Tear (default)
    AutoRoute: disabled  LockDown: disabled   Policy class: not set
    Forward class: 0 (not enabled)
    Forwarding-Adjacency: disabled
    Autoroute Destinations: 0
    Loadshare:          0 equal loadshares
    Auto-bw: disabled
    Auto-Capacity: Disabled:
    Fast Reroute: Enabled, Protection Desired: Bandwidth
    Path Protection: Enabled
      Non-revertive
    BFD Fast Detection: Disabled
    Reoptimization after affinity failure: Enabled
    Soft Preemption: Disabled
PCE Delegation:
    Symbolic name: "PCEP-TEST"
    PCEP ID: 301
    Delegated to: 203.0.113.1
  History:
    Tunnel has been up for: 2d13h (since Mon Jul 27 23:30:10 IST 2020)
    Current LSP:
      Uptime: 1d03h (since Wed Jul 29 09:36:50 IST 2020)
    Prior LSP:
      ID: 124 Path Option: 10
      Removal Trigger: reoptimization completed
  Path info (PCE controlled):
  Hop0: 209.165.200.4
  Hop1: 51.0.0.2
Displayed 1 (of 2004) headends, 0 (of 0) midpoints, 0 (of 7) tailends
Displayed 1 up, 0 down, 0 recovering, 0 recovered headends
```

Step 9 Verify the PCE node peer address and state using the **show mpls traffic-eng pce peer** command:

```
RP/0/RP0:NCS4016-1#show mpls traffic-eng pce peer
Address          Precedence    State         Learned From
---------------- ------------ ------------ --------------------
203.0.113.1        10           Up           Static config
RP/0/RP0:NCS4016-1#show mpls tr pce lsp-database brief
PCE ID Tun ID LSP ID Symbolic-name  Destination     State Type DLG
------ ------ ------ -------------------- --------------- ----- ---- ---
301    300    130    PCEP-TEST     198.51.100.3        Up    Conf yes *Manual + PCE Delegated
5001    5000   8      m1        198.51.100.3        Up    Init yes . .Curl or PCE Initiated
• CURL COMMAND INITIATED TUNNEL
*Manually CONFIGURED under HEADEND Node (Tunnel-te 300)\
```

Step 10 Check the LSP database of the tunnel using the **show mpls traffic-eng pce lsp-database** command:

```
RP/0/RP0:NCS4016-1#show mpls traffic-eng pce lsp-database symbolic-name PCEP-TEST detail
Thu Jul 30 16:50:05.121 IST
Symbolic name: PCEP-TEST
Session internal LSP ID: 301
```

```
Stateful Request Parameters ID: 0
Path Setup Type: 0 - (RSVP)
Request queue size: 0
Create: FALSE
    Created by: Not set
Delegatable: TRUE
    Delegation status: Delegated
    Delegated to: Speaker-entity-id: Not set ip: 203.0.113.1
Destination: 198.51.100.3    Source: 198.51.100.1
LSP Object:
    Administrative: Up
    Operational state: Up
    Identifiers:
        Sender Address: 198.51.100.1
        TE LSP ID: 141
        Tunnel ID: 300
        Extended tunnel ID: 0x3030303
    Binding SID: 24012
LSP Path Object:
    Explicit Route Object:
        Cost: 0
        1.  ipv4: 209.165.200.4/32 (strict)
        2.  ipv4: 51.0.0.2/32 (strict)
LSP Attributes:
        Exclude any: 0
        Include any: 0
        Include all: 0
        Setup priority: 7
        Hold priority: 7
        Local Protection Bit: TRUE
    Reported Route Object:
        Cost: 0
        1.  ipv4: 198.51.100.2/32
        2. label: 26004 (global)
        3.  ipv4: 209.165.200.4/32
        4. label: 26004 (global)
        5.  ipv4: 198.51.100.3/32
        6. label: 0 (global)
        7.  ipv4: 51.0.0.2/32
        8. label: 0 (global)
    Bandwidth:  0 Bps (0 kbps)
    Reoptimized bandwidth: Not set
    Applied bandwidth: Not set
    Metric:
        Cost: 20        Type: IGP
    Vendor Specific Information:
        Forward-Class: Not set
        Load Share: Not set
        Backup path: Not set
```

Step 11 Verify PCEP session details using the **show pce ipv4 peer** command:

```
RP/0/RP0/CPU0:NCS5500-10#show pce ipv4 peer
PCE's peer database:
--------------------
Peer address: 198.51.100.1
  State: Up
  Capabilities: Stateful, Update, Instantiation
RP/0/RP0/CPU0:NCS5500-10#show pce lsp tabular
PCC               Tunnel Name   Color   Source          Destination    TunID  LSPID  Admin
  Oper
198.51.100.1       PCEP-TEST       0     198.51.100.1     198.51.100.3   00     141    up
    up  Manual
```

```
198.51.100.1       m1              0     198.51.100.1     198.51.100.3   5000   8      up
    up  □ PCE Initiated (CURL)
```

Step 12 View the summary of the PCE topology information using the **show pce ipv4 topology summary** command:

```
RP/0/RP0/CPU0:NCS5500-10#show pce ipv4 topology summary
PCE's topology database summary:
--------------------------------
Topology nodes:                4
Prefixes:                      4
Prefix SIDs:
  Total:                       0
  Regular:                     0
  Strict:                      0
Links:
  Total:                       8
  EPE:                         0
Adjacency SIDs:
  Total:                       0
  Unprotected:                 0
  Protected:                   0
  EPE:                         0
Private Information:
Lookup Nodes                   4
Consistent                   yes
Update Stats (from IGP and/or BGP):
  Nodes added:                 4
  Nodes deleted:               0
  Links added:                11
  Links deleted:               3
  Prefix added:               12
  Prefix deleted:              0
Topology Ready Summary:
  Ready:                     yes
  PCEP allowed:              yes
  Last HA case:          startup
  Timer value (sec):         300
  Timer:
    Running: no
```

Step 13 View the detailed information of an LSP present in the PCE's LSP database, in table format using the **show pce lsp tabular** command:

```
RP/0/RP0/CPU0:NCS5500-10#show pce lsp tabular
Tue Feb 9 11:14:08.858 UTC
PCC           TunnelName      Color  Source       Destination  TunID  LSPID  Admin Oper
198.51.100.1  NCS4016-1_t1000   0   198.51.100.1 198.51.100.2 1000   10     up    up
198.51.100.1  NCS4016-1_t300    0   198.51.100.1 198.51.100.2 300    6      up    up
198.51.100.1  m                 0   198.51.100.1 198.51.100.2 5000   3      up    up
198.51.100.1  mapm1             0   198.51.100.1 198.51.100.2 5003   3      up    up
198.51.100.1  te99              0   198.51.100.1 198.51.100.2 5002   4      up    up
198.51.100.1  tunnel-te500      0   198.51.100.1 198.51.100.2 5001   3      up    up
```