Configuring MPLS Traffic Engineering—Bundled Interface Support

- Prerequisites for MPLS TE—Bundled Interface Support, on page 1
- Restrictions for MPLS TE—Bundled Interface Support, on page 1
- Information About MPLS TE—Bundled Interface Support, on page 2
- How to Configure MPLS TE—Bundled Interface Support, on page 3
- Configuration Examples for MPLS Traffic Engineering—Bundled Interface Support, on page 4
- Additional References for MPLS Traffic Engineering—Bundled Interface Support, on page 7
- Feature History for MPLS Traffic Engineering—Bundled Interface Support, on page 7

Prerequisites for MPLS TE—Bundled Interface Support

- Configure Multiprotocol Label Switching (MPLS) traffic engineering (TE) tunnels.
- Enable Cisco Express Forwarding in global configuration mode.
- Enable Resource Reservation Protocol (RSVP) feature.
- Configure EtherChannel.
- Configure Gigabit EtherChannel.

Restrictions for MPLS TE—Bundled Interface Support

- Traffic engineering over switch virtual interfaces (SVIs) is not supported unless the SVI consists of a bundle of links that represent a single point-to-point interface.
- There must be a valid IP address configuration on the bundled interface and there must not be an IP address configuration on the member links.

Information About MPLS TE—Bundled Interface Support

The MPLS Traffic Engineering—Bundled Interface Support feature enables Multiprotocol Label Switching (MPLS) traffic engineering (TE) tunnels over the bundled interfaces—EtherChannel and Gigabit EtherChannel (GEC).

The Resource Reservation Protocol (RSVP) notifies TE about bandwidth changes that occur when member links are added or deleted, or when links become active or inactive. TE notifies other nodes in the network via Interior Gateway Protocol (IGP) flooding. By default, the bandwidth available to TE Label-Switched Paths (LSPs) is 75 percent of the interface bandwidth. You can change the percentage of the global bandwidth available for TE LSPs by using an RSVP command on the bundled interface. The feature supports bandwidth reservation and preemption.

The following section provides information about Bundled Interface Support for MPLS Traffic Engineering.

Cisco EtherChannel Overview

Cisco EtherChannel technology builds upon standards-based 802.3 full-duplex Fast Ethernet to provide network managers with a reliable, high-speed solution for the campus network backbone. EtherChannel technology provides bandwidth scalability within the campus. It provides up to 800 Mbps, 8 Gbps, or 80 Gbps of aggregate bandwidth for a Fast EtherChannel, Gigabit EtherChannel, or 10 Gigabit EtherChannel connection, respectively. Each of these connection speeds can vary in amounts equal to the speed of the links used (100 Mbps, 1 Gbps, or 10 Gbps). Even in the most bandwidth-demanding situations, EtherChannel technology helps to aggregate traffic, keeps oversubscription to a minimum, and provides effective link-resiliency mechanisms.

Cisco EtherChannel Benefits

Cisco EtherChannel technology allows network managers to provide higher bandwidth among servers, routers, and switches than a single-link Ethernet technology can provide.

Cisco EtherChannel technology provides incremental scalable bandwidth and the following benefits:

- Standards-based—Cisco EtherChannel technology builds upon IEEE 802.3-compliant Ethernet by grouping multiple, full-duplex point-to-point links. EtherChannel technology uses IEEE 802.3 mechanisms for full-duplex autonegotiation and autosensing, when applicable.
- Flexible incremental bandwidth—Cisco EtherChannel technology provides bandwidth aggregation in multiples of 100 Mbps, 1 Gbps, or 10 Gbps. This depends on the speed of the aggregated links. For example, network managers can deploy EtherChannel technology that consists of pairs of full-duplex Fast Ethernet links to provide more than 400 Mbps between the wiring closet and the data center. In the data center, bandwidths of up to 800 Mbps can be provided between servers and the network backbone to provide large amounts of scalable incremental bandwidth.
- Load balancing—Cisco EtherChannel technology comprises several Fast Ethernet links. It is capable of load balancing traffic across those links. Unicast, broadcast, and multicast traffic is evenly distributed across the links, providing improved performance and redundant parallel paths. When a link fails, traffic is redirected to the remaining links within the channel without user intervention and with minimal packet loss.
- Resiliency and fast convergence—When a link fails, Cisco EtherChannel technology provides automatic recovery by redistributing the load across the remaining links. When a link fails, Cisco EtherChannel

technology redirects traffic from the failed link to the remaining links in less than one second. This convergence is transparent to the end user—no host protocol timers expire and no sessions are dropped.

Cisco Gigabit EtherChannel Overview

Cisco Gigabit EtherChannel (GEC) is a high-performance Ethernet technology that provides transmission rates in Gigabit per second (Gbps). A Gigabit EtherChannel bundles individual ethernet links (Gigabit Ethernet and 10 Gigabit Ethernet) into a single logical link. This single link provides the aggregate bandwidth of up to four physical links. All LAN ports in each EtherChannel must be of the same speed and must be configured as either Layer 2 or Layer 3 LAN ports. Broadcast and multicast packets which are inbound on one link in an EtherChannel are blocked from returning on any other link in the EtherChannel.

Load Balancing in EtherChannel

Load balancing affects the actual and practical bandwidth that is in use for TE. Multilink load balancing uses a per-packet load balancing method. The entire bundle interface bandwidth is available. EtherChannel load balancing has various load balancing methods, depending on the traffic pattern and the load balancing configuration. The total bandwidth available for TE may be limited to the bandwidth of a single member link.

How to Configure MPLS TE—Bundled Interface Support

The following section provides information about how to configure Bundled Interface Support for MPLS Traffic Engineering.

Configuring MPLS Traffic Engineering on an EtherChannel Interface

To configure MPLS Traffic Engineering on an etherchannel interface, perform the following procedure.

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.

	Command or Action	Purpose
Step 3	**interface** *type number* [*name-tag*] **Example:** `Device(config)# interface port-channel 1`	Creates an EtherChannel bundle, assigns a group number to the bundle, and enters interface configuration mode.
Step 4	**ip address** *ip-address mask* [**secondary**] **Example:** `Device(config-if)# ip address 10.0.0.4 255.255.255.0`	Specifies an IP address for the EtherChannel group.
Step 5	**mpls traffic-eng tunnels** **Example:** `Device(config-if)# mpls traffic-eng tunnels`	Enables MPLS TE tunnel signaling on an interface. • Enable MPLS TE tunnel on the device before enabling the signaling.
Step 6	**ip rsvp bandwidth** [*interface-kbps*] [*single-flow-kbps*] **Example:** `Device(config-if)# ip rsvp bandwidth 100`	Enables RSVP for IP on an interface. Specifies a percentage of the total interface bandwidth as available in the RSVP bandwidth pool.
Step 7	**end** **Example:** `Device(config-if)# end`	Exits interface configuration mode and returns to privileged EXEC mode.

Configuration Examples for MPLS Traffic Engineering—Bundled Interface Support

The following section provides configuration examples for MPLS Traffic Engineering—Bundled Interface Support.

Example: Configuring MPLS TE on an EtherChannel Interface

The following example shows how to configure MPLS TE on an EtherChannel interface.

```
Device> enable
Device# configure terminal
Device(config)# interface port-channel 1
Device(config-if)# ip address 10.0.0.4 255.255.255.0
Device(config-if)# mpls traffic-eng tunnels
```

```
Device(config-if)# ip rsvp bandwidth 100
Device(config-if)# end
```

Example: Configuring MPLS Traffic Engineering—Bundled Interface Support over Gigabit Etherchannel

The following example shows how to enable MPLS Traffic Engineering—Bundled Interface Support over GEC on Cisco devices:

```
Device> enable
Device# configure terminal

! Enable global MPLS TE on routers
Device(config)# router ospf 100
Device(config-router)# network 10.0.0.1 0.0.0.255 area 0
Device(config-router)# mpls traffic-eng area 0
Device(config-router)# mpls traffic-eng router-id Loopback 0
Device(config-router)# exit

! Configure  GEC  interface and enable MPLS TE and RSVP on interface
Device(config)# interface Port-channel 1
Device(config-if)# ip address 10.0.0.1 255.255.255.0
Device(config-if)# mpls traffic-eng tunnels
Device(config-if)# ip rsvp bandwidth
Device(config-if)# exit

! Define explicit path
Device(config)# ip explicit-path name primary enable
Device(cfg-ip-expl-path)# next-address 172.12.1.2
Device(cfg-ip-expl-path)# next-address 172.23.1.2
Device(cfg-ip-expl-path)# next-address 172.34.1.2
Device(cfg-ip-expl-path)# next-address 10.4.4.4
Device(cfg-ip-expl-path)# exit

! Configure primary tunnel on head-end device
Device(config)# interface Tunnel 14
Device(config-if)# ip unnumbered Loopback 0
Device(config-if)# tunnel mode mpls traffic-eng
Device(config-if)# tunnel destination 10.10.10.0
Device(config-if)# tunnel mpls traffic-eng autoroute announce
Device(config-if)# tunnel mpls traffic-eng path-option 10 explicit name primary
Device(config-if)# exit

! Configure GEC interface
Device(config)# interface GigabitEthernet 0/0/1
Device(config-if)# no ip address
Device(config-if)# channel-group 1 mode active
Device(config-if)# exit

! Configure GEC interface
Device(config)# interface GigabitEthernet 0/0/2
Device(config-if)# no ip address
Device(config-if)# channel-group 1 mode active
Device(config-if)# exit
```

The **show mpls traffic-eng tunnels** command output displays information about a tunnel or one–line information about all tunnels configured on the device:

```
Device# show mpls traffic-eng tunnels tunnel 14

Name: Cat9k_t14                           (Tunnel10) Destination: 10.4.4.4
  Status:
    Admin: up         Oper: up     Path: valid       Signalling: connected
    path option 1, type explicit toR4overR3R3 (Basis for Setup, path weight 3)

  Config Parameters:
    Bandwidth: 0        kbps (Global)  Priority: 7  7   Affinity: 0x0/0xFFFF
    Metric Type: TE (default)
    AutoRoute: enabled  LockDown: disabled Loadshare: 0 [0] bw-based
    auto-bw: disabled
  Active Path Option Parameters:
    State: explicit path option 1 is active
    BandwidthOverride: disabled  LockDown: disabled  Verbatim: disabled


      InLabel  :  -
  OutLabel : Port-channel1, 1608
  Next Hop : 172.16.1.2
  RSVP Signalling Info:
       Src 10.1.1.1, Dst 10.4.4.4, Tun_Id 14, Tun_Instance 35
    RSVP Path Info:
      My Address: 172.12.1.1
      Explicit Route: 172.12.1.2 172.23.1.1 172.23.1.2 172.34.1.1
                      172.34.1.2 10.4.4.4

  History:
    Tunnel:
      Time since created: 17 hours
      Time since path change: 18 minutes, 22 seconds
      Number of LSP IDs (Tun_Instances) used: 35
    Current LSP: [ID: 35]
      Uptime: 18 minutes, 22 seconds
      Selection: reoptimization
    Prior LSP: [ID: 32]
      ID: path option unknown
      Removal Trigger: signalling shutdown


Device# show mpls traffic-eng tunnels brief

show mpls traffic-eng tunnels brief
Signalling Summary:
    LSP Tunnels Process:            running
    Passive LSP Listener:           running
    RSVP Process:                   running
    Forwarding:                     enabled
    Periodic reoptimization:        every 3600 seconds, next in 3299 seconds
    Periodic FRR Promotion:         Not Running
    Periodic auto-bw collection:    every 300 seconds, next in 299 seconds

P2P TUNNELS/LSPs:
TUNNEL NAME                      DESTINATION      UP IF     DOWN IF   STATE/PROT^M
Cat9k_t14                      10.4.1.1                        -         Po12      up/up
On Mid Router:
P2P TUNNELS/LSPs:
TUNNEL NAME                       DESTINATION      UP IF     DOWN IF   STATE/PROT
Cat9k_t14                        10.4.1.1                  Po12       Po23       up/up
Cat9k_t23                        10.2.1.1                  Po25        -         up/up
```

Additional References for MPLS Traffic Engineering—Bundled Interface Support

Related Documents

Related Topic	Document Title
Cisco IOS commands	Cisco IOS Master Commands List, All Releases
MPLS traffic engineering commands	Cisco IOS Multiprotocol Label Switching Command Reference
IPv6 commands	IPv6 Command Reference

Technical Assistance

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

Feature History for MPLS Traffic Engineering—Bundled Interface Support

This table provides release and related information for the features explained in this module.

These features are available in all the releases subsequent to the one they were introduced in, unless noted otherwise.

Release	Feature	Feature Information
Cisco IOS XE Bengaluru 17.6.1	MPLS Traffic Engineering—Bundled Interface Support	The MPLS Traffic Engineering—Bundled Interface Support feature enables MPLS Traffic Engineering tunnels over the bundled interfaces EtherChannel and Gigabit EtherChannel (GEC).

Use the Cisco Feature Navigator to find information about platform and software image support. To access Cisco Feature Navigator, go to https://cfnng.cisco.com/