MPLS Traffic Engineering—Configurable Path Calculation Metric for Tunnels

First Published: 12.0(18)ST
Last Updated: February 28, 2006

The MPLS Traffic Engineering—Configurable Path Calculation Metric for Tunnels feature enables the user to control the metric used in path calculation for TE tunnels on a per-tunnel basis.

History for the MPLS Traffic Engineering—Configurable Path Calculation Metric for Tunnels Feature

Release	Modification
12.0(18)ST	This feature was introduced.
12.2(14)S	This feature was integrated into Cisco IOS Release 12.2(14)S.
12.2(28)SB	This feature was integrated into Cisco IOS Release 12.2(28)SB.

Finding Support Information for Platforms and Cisco IOS Software Images

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

Contents

CISCO SYSTEMS

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

Feature Overview

When Multiprotocol Label Switching (MPLS) traffic engineering (TE) is configured in a network, the Interior Gateway Protocol (IGP) floods two metrics for every link: the normal IGP (Open Shortest Path First (OSPF) or (Intermediate System-to-Intermediate System (IS-IS)) link metric and a TE link metric. The IGP uses the IGP link metric in the normal way to compute routes for destination networks. In previous releases, MPLS TE used the TE link metric to calculate and verify paths for TE tunnels. When the traffic engineering metric was not explicitly configured, the traffic engineering metric was the IGP metric.

The current enhancement allows you to specify that the path calculation for a given tunnel be based on either of the following:

- IGP link metrics.
- TE link metrics, which you can configure so that they represent the needs of a particular application. For example, the TE link metrics can be configured to represent link transmission delay.

Benefits

When TE tunnels are used to carry two types of traffic, this enhancement allows you to tailor tunnel path selection to the requirements of each type of traffic.

For example, suppose certain tunnels are to carry voice traffic (which requires low delay) and other tunnels are to carry data. In this situation, you can use the TE link metric to represent link delay and do the following:

- Configure tunnels that carry voice to use the TE link metric set to represent link **delay** for path calculation.
- Configure tunnels that carry data to use the IGP metric for path calculation.

Restrictions

Unless explicitly configured, the TE link metric for a given link is the IGP link metric. When the TE link metric is used to represent a link property that is different from cost/distance, you must configure every network link that can be used for TE tunnels with a TE link metric that represents that property by using the **mpls traffic-eng administrative-weight** command. Failure to do so might cause tunnels to use unexpected paths.

Related Features and Technologies

The configurable path calculation metric feature is related to MPLS traffic engineering.

Prerequisites

Before you configure tunnel path calculation metrics, your network must support the following Cisco IOS features:

- MPLS traffic engineering tunnels
- IP Cisco Express Forwarding
- OSPF or IS-IS

Configuration Tasks

See the following sections for configuration tasks for the configurable path calculation metric feature. Each task in the list indicates if it is optional or required.

- Configuring a Platform to Support Traffic Engineering Tunnels, page 3 (required)
- Configuring the IGP (OSPF or IS-IS) for MPLS Traffic Engineering (required)
  - Configuring IS-IS for MPLS Traffic Engineering, page 3
  - Configuring OSPF for MPLS Traffic Engineering, page 4
- Configuring Traffic Engineering Link Metrics, page 4 (required)
- Configuring an MPLS Traffic Engineering Tunnel, page 5 (required)
- Configuring Metric Type for Tunnel Path Calculation, page 6 (required)

Configuring a Platform to Support Traffic Engineering Tunnels

To configure a platform to support traffic engineering tunnels, perform the following steps in configuration mode:

	Command	Purpose
Step 1	`Router(config)# ip cef`	Enables standard Cisco Express Forwarding operation. For information about Cisco Express Forwarding configuration and the command syntax, see the *Cisco IOS Switching Services Configuration Guide* and the *Cisco IOS IP Switching Command Reference*.
Step 2	`Router(config)# mpls traffic-eng tunnels`	Enables the MPLS traffic engineering tunnel feature on a device.

Configuring IS-IS for MPLS Traffic Engineering

To configure IS-IS for MPLS traffic engineering, perform the steps described below. For a description of the IS-IS commands, see the *Cisco IOS IP Command Reference, Volume 2 of 3: Routing Protocols,* Release 12.2 manual.

	Command	Purpose
Step 1	`Router(config)# router isis`	Enables IS-IS routing and specifies an IS-IS process for IP. This command places you in router configuration mode.
Step 2	`Router(config-router)# mpls traffic-eng level-1`	Turns on MPLS traffic engineering for IS-IS level 1.
Step 3	`Router(config-router)# mpls traffic-eng level-2`	Turns on MPLS traffic engineering for IS-IS level 2.
Step 4	`Router(config-router)# mpls traffic-eng router-id loopback0`	Specifies that the traffic engineering router identifier for the node is the IP address associated with interface loopback0.
Step 5	`Router(config-router)# metric-style wide`	Configures a router to generate and accept only new-style type, length, value objects (TLVs).

Configuring OSPF for MPLS Traffic Engineering

To configure OSPF for MPLS traffic engineering, perform the steps described below. For a description of the OSPF commands, see the *Cisco IOS IP Command Reference, Volume 2 of 3: Routing Protocols,* Release 12.2 manual.

	Command	Purpose
Step 1	`Router(config)# router ospf` *process-id*	Configures an OSPF routing process for IP. You are placed in router configuration mode. The *process-id* is an internally used identification parameter for an OSPF routing process. It is locally assigned and can be any positive integer. Assign a unique value for each OSPF routing process.
Step 2	`Router(config-router)# mpls traffic-eng area` *x*	Turns on MPLS traffic engineering for OSPF area *x*.
Step 3	`Router(config-router)# mpls traffic-eng router-id loopback0`	Specifies that the traffic engineering router identifier for the node is the IP address associated with interface loopback0.

Configuring Traffic Engineering Link Metrics

Unless explicitly configured, the TE link metric is the IGP link metric. To configure the TE link metric, perform these steps:

	Command	Purpose
Step 1	`Router(config)# interface interface`	Configures properties of the specified interface.
Step 2	`Router(config-router)# mpls traffic-eng administrative-weight weight`	Specifies the traffic engineering metric for the link.

Configuring an MPLS Traffic Engineering Tunnel

To configure an MPLS traffic engineering tunnel, perform the following steps in interface configuration mode. This tunnel has two path setup options: a preferred explicit path and a backup dynamic path. For more detailed descriptions of the commands and their arguments, see the *Cisco IOS Switching Services Command Reference*, Release 12.2 and the *Cisco IOS Switching Services Configuration Guide*, Release 12.2.

	Command	Purpose
Step 1	`Router(config)# interface tunnel-interface`	Configures a tunnel interface and enters interface configuration mode.
Step 2	`Router(config-if)# ip unnumbered loopback0`	Gives the tunnel interface an IP address. An MPLS traffic engineering tunnel interface should be unnumbered because it represents a unidirectional link.
Step 3	`Router(config-if)# tunnel destination A.B.C.D`	Specifies the destination for a tunnel. The destination must be the MPLS traffic engineering router ID of the destination device.
Step 4	`Router(config-if)# tunnel mode mpls traffic-eng`	Sets the tunnel encapsulation mode to MPLS traffic engineering.
Step 5	`Router(config-if)# tunnel mpls traffic-eng bandwidth bandwidth`	Configures the bandwidth for the MPLS traffic engineering tunnel. If automatic bandwidth is configured for the tunnel, use the **tunnel mpls traffic-eng bandwidth** command to configure the initial tunnel bandwidth, which is adjusted by the autobandwidth mechanism.
Step 6	`Router(config-if)# tunnel mpls traffic-eng path-option number {dynamic | explicit {name path-name | id path-number}} [lockdown]`	Configures the tunnel to use a named IP explicit path or a path dynamically calculated from the traffic engineering topology database. A dynamic path is used if an explicit path is currently unavailable.

Configuring Metric Type for Tunnel Path Calculation

Unless explicitly configured, the TE link metric type is used for tunnel path calculation. Two commands are provided for controlling the metric type to be used: an interface configuration command that specifies the metric type to be used for a particular TE tunnel and a global configuration command that specifies the metric type to be used for TE tunnels for which a metric type has not been specified by the interface configuration command.

To configure the metric type for tunnel path calculation, perform one or both of the following steps:

Command	Purpose
`Router(config-if)# tunnel mpls traffic-eng path-selection metric {igp | te}`	Specifies the metric type to use for path calculation when you are determining a tunnel's path.
`Router(config)# mpls traffic-eng path-selection metric {igp | te}`	Specifies the metric type to use if a metric type was not explicitly configured for a given tunnel.

Note If you do not enter either of the above commands, the traffic engineering (te) metric is used.

Verifying the Configuration

Use the **show mpls traffic-eng topology** command, which displays TE and IGP metrics for each link, to verify that link metrics have been correctly configured for a network.

```
Router# show mpls traffic-eng topology

My_System_id: 1440.0000.0044.00 (isis  level-1)
IGP Id: 0090.0000.0009.00, MPLS TE Id:192.168.9.9 Router Node  (isis  level-1)
      link[0 ]:Nbr IGP Id: 0090.0000.0009.03, gen:7
          frag_id 0, Intf Address:10.0.0.99
          TE metric:100, IGP metric:48, attribute_flags:0x0     !!Note TE and IGP metrics
          physical_bw: 10000 (kbps), max_reservable_bw_global: 0 (kbps)
          max_reservable_bw_sub: 0 (kbps)
   .
   .
   .
      link[1 ]:Nbr IGP Id: 0055.0000.0055.00, gen:7
          frag_id 0, Intf Address:10.205.0.9, Nbr Intf Address:10.205.0.55
          TE metric:120, IGP metric:10, attribute_flags:0x0     !!Note TE and IGP metrics
          physical_bw: 155000 (kbps), max_reservable_bw_global: 500000 (kbps)
          max_reservable_bw_sub: 0 (kbps)
   .
   .
   .
```

Use the **show mpls traffic-eng tunnels** command, which displays the link metric used for tunnel path calculation, to verify that the desired link metrics are being used for each tunnel.

```
Router# show mpls traffic-eng tunnels

Name: te3640-17-c_t221              (Tunnel22) Destination: 192.168.100.22
  Status:
    Admin: up         Oper: up     Path: valid       Signalling: connected
```

```
    path option 1, type dynamic (Basis for Setup, path weight 10)

  Config Parameters:
    Bandwidth: 400 kps (Global)    Priority: 1  1   Affinity: 0x0/0xFFFF
    Metric Type: IGP                                        !!Note metric type
    AutoRoute:  enabled   LockDown: disabled  Loadshare: 0   bw-based
    auto-bw: disabled(0/115) 0  Bandwidth Requested: 0
  .
  .
  .
Name: te3640-17-c_t222              (Tunnel33) Destination: 192.168.100.22
  Status:
    Admin: up         Oper: up     Path: valid       Signalling: connected
    path option 1, type dynamic (Basis for Setup, path weight 10)

  Config Parameters:
    Bandwidth: 200 kbps (Global)   Priority: 1  1   Affinity: 0x0/0xFFFF
    Metric Type: TE                                         !!Note metric type
    AutoRoute:  enabled   LockDown: disabled  Loadshare: 0   bw-based
    auto-bw: disabled(0/115) 0  Bandwidth Requested: 0
  .
  .
  .
```

Configuration Examples

The section illustrates how to configure the link metric type to be used for tunnel path selection, and how to configure the link metrics themselves. The configuration commands included below focus on specifying the metric type for path calculation and assigning metrics to links. Additional commands are required to fully configure the example scenario; for example, the IGP commands for traffic engineering and the link interface commands for enabling traffic engineering and specifying available bandwidth.

Consider the simple network topology shown in Figure 1.

Figure 1 Network Topology

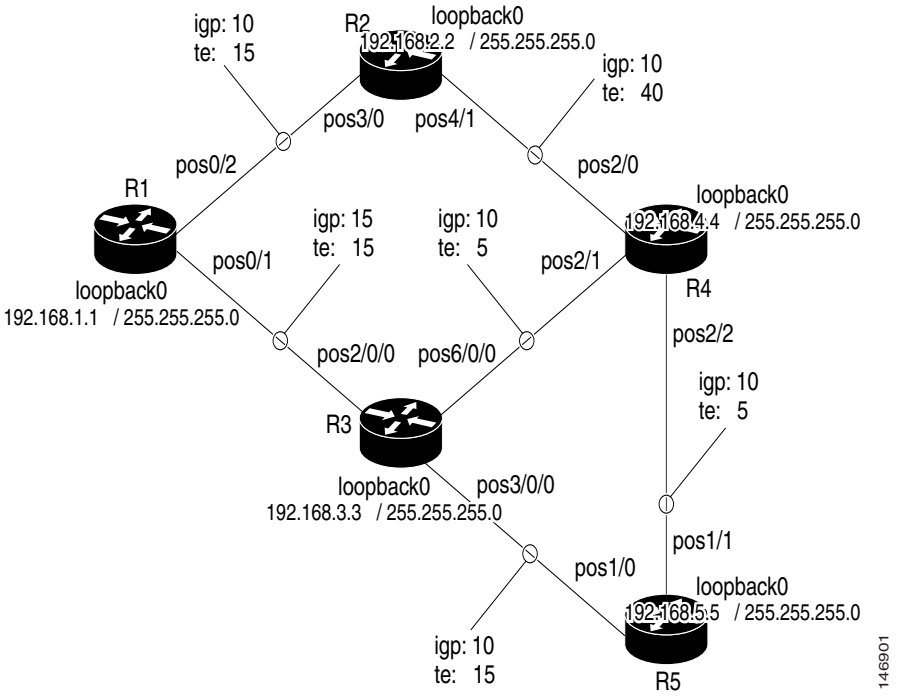

In Figure 1

- Tunnel1 and Tunnel2 run from R1 (headend) to R4 (tailend).
- Tunnel3 runs from R1 to R5.
- Path calculation for Tunnel1 and Tunnel3 should use a metric that represents link delay because these tunnels carry voice traffic.
- Path calculation for Tunnel2 should use IGP metrics because it carries data traffic with no delay requirement.

Below are configuration fragments for each of the routers that illustrate the configuration relating to link metrics and their use in tunnel path calculation. TE metrics that represent link delay must be configured for the network links on each of the routers, and the three tunnels must be configured on R1.

These configuration fragments force Tunnel1 to take path R1-R3-R4, Tunnel2 to take path R1-R2-R4, and Tunnel3 to take path R1-R3-R4-R5 (assuming the links have sufficient bandwidth to accommodate the tunnels).

R1 Configuration

```
interface pos0/1
mpls traffic-eng administrative-weight 15       !TE metric different from IGP metric
interface pos0/2
mpls traffic-eng administrative-weight 15       !TE metric different from IGP metric

interface Tunnel1                               !Tunnel1 uses TE metric (default)
                                                !for path selection
ip unnumbered loopback0
tunnel destination 192.168.4.4 255.255.255.0
tunnel mode mpls traffic-eng
tunnel mpls traffic-eng bandwidth 1000
tunnel mpls traffic-eng path-option 1 dynamic

interface Tunnel2                                !Tunnel2 uses IGP metric
```

```
                                                    !for path selection
ip unnumbered loopback0
tunnel destination 192.168.4.4 255.255.255.0
tunnel mode mpls traffic-eng
tunnel mpls traffic-eng bandwidth 1000
tunnel mpls traffic-eng path-option 1 dynamic
tunnel mpls traffic-eng path-selection-metric igp !Use IGP cost for path selection.

interface Tunnel3                                    !Tunnel3 uses TE metric (default)
                                                     !for path selection
ip unnumbered loopback0
tunnel destination 192.168.5.5 255.255.255.0
tunnel mode mpls traffic-eng
tunnel mpls traffic-eng bandwidth 1000
tunnel mpls traffic-eng path-option 1 dynamic
```

R2 Configuration

```
interface pos3/0
mpls traffic-eng administrative-weight 15       !TE metric different from IGP metric
interface pos4/1
mpls traffic-eng administrative-weight 40       !TE metric different from IGP metric
```

R3 Configuration

```
interface pos2/0/0
mpls traffic-eng administrative-weight 15       !TE metric different from IGP metric
interface pos3/0/0
mpls traffic-eng administrative-weight 15       !TE metric different from IGP metric
interface pos6/0/0
mpls traffic-eng administrative-weight 5        !TE metric different from IGP metric
```

R4 Configuration

```
interface pos2/0
mpls traffic-eng administrative-weight 40       !TE metric different from IGP metric
interface pos2/1
mpls traffic-eng administrative-weight 15       !TE metric different from IGP metric
interface pos2/2
mpls traffic-eng administrative-weight 5        !TE metric different from IGP metric
```

R5 Configuration

```
interface pos1/0
mpls traffic-eng administrative-weight 15       !TE metric different from IGP metric
interface pos1/1
mpls traffic-eng administrative-weight 5        !TE metric different from IGP metric
```

Additional References

The following sections provide references related to MPLS Traffic Engineering—Configurable Path Calculation Metric for Tunnels.

Related Documents

Related Topic	Document Title
IP commands	*Cisco IOS IP Command Reference, Volume 2 of 3: Routing Protocols*, Release 12.2
Switching services commands	*Cisco IOS IP Switching Command Reference*, Release 12.4T
Switching services	*Cisco IOS Switching Services Configuration Guide*, Release 12.2

Standards

Standard	Title
None	—

MIBs

MIB	MIBs Link
None	To locate and download MIBs for selected platforms, Cisco IOS releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

RFCs

RFC	Title
None	—

Technical Assistance

Description	Link
The Cisco Technical Support & Documentation website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport

Command Reference

This section documents modified commands only.

- **mpls traffic-eng path-selection metric**
- **tunnel mpls traffic-eng path-selection metric**

mpls traffic-eng path-selection metric

To specify the metric type to use for path selection for tunnels for which the metric type has not been explicitly configured, use the **mpls traffic-eng path-selection metric** command in global configuration mode. To remove the specified metric type, use the **no** form of this command.

mpls traffic-eng path-selection metric {**igp** | **te**}

no mpls traffic-eng path-selection metric

Syntax Description

igp	Use the Interior Gateway Protocol (IGP) metric.
te	Use the traffic engineering metric.

Defaults

The default is the **te** metric.

Command Modes

Global configuration

Command History

Release	Modification
12.0(18)ST	This command was introduced.
12.2(14)S	This command was integrated into Cisco IOS Release 12.2(14)S.
12.2(28)SB	This command was integrated into Cisco IOS Release 12.2(28)SB.

Usage Guidelines

Use this command to specify the metric type to be used for traffic engineering (TE) tunnels for which the **tunnel mpls traffic-eng path-selection metric** command has not been specified.

The metric type to be used for path calculation for a given tunnel is determined as follows:

- If the **tunnel mpls traffic-eng path-selection metric** command was entered to specify a metric type for the tunnel, use that metric type.
- Otherwise, if the **mpls traffic-eng path-selection metric** was entered to specify a metric type, use that metric type.
- Otherwise, use the default (**te**) metric.

Examples

The following command specifies that if a metric type was not specified for a given TE tunnel, the igp metric should be used for tunnel path calculation:

```
Router(config)# mpls traffic-eng path-selection metric igp
```

Related Commands

Command	Description
tunnel mpls traffic-eng path-selection metric	Specifies the metric type to use when calculating a tunnel's path.

tunnel mpls traffic-eng path-selection metric

To specify the metric type to use for path calculation for a tunnel, use the **tunnel mpls traffic-eng path-selection metric** command in interface configuration mode. To remove the specified metric type, use the **no** form of this command.

tunnel mpls traffic-eng path-selection metric {**igp** | **te**}

no tunnel mpls traffic-eng path-selection metric

Syntax Description

igp	Use the Interior Gateway Protocol (IGP) metric.
te	Use the traffic engineering (TE) metric.

Defaults

The default is the **te** metric.

Command Modes

Interface configuration

Command History

Release	Modification
12.0(18)ST	This command was introduced.
12.2(14)S	This command was integrated into Cisco IOS Release 12.2(14)S.
12.2(28)SB	This command was integrated into Cisco IOS Release 12.2(28)SB.

Usage Guidelines

The metric type to be used for path calculation for a given tunnel is determined as follows:

- If the **tunnel mpls traffic-eng path-selection metric** command was entered to specify a metric type for the tunnel, use that metric type.
- Otherwise, if the **mpls traffic-eng path-selection metric** was entered to specify a metric type, use that metric type.
- Otherwise, use the default (**te**) metric.

Examples

The following commands specify that the igp metric should be used when you are calculating the path for Tunnel102:

```
Router(config)# interface tunnel102
Router(config-if)# tunnel mpls traffic-eng path-selection metric igp
```

Related Commands

Command	Description
mpls traffic-eng path-selection metric	Specifies the metric type to use for path calculation for TE tunnels for which no metric has been explicitly configured.

Glossary

Cisco Express Forwarding—An advanced Layer 3 IP switching technology. CEF optimizes network performance and scalability for networks that have large and dynamic traffic patterns, such as the Internet, as well as for networks characterized by intensive Web-based applications or interactive sessions. CEF uses a Forwarding Information Base (FIB) to make IP destination prefix-based switching decisions. The FIB is conceptually similar to a routing table or information base. When routing or topology changes occur in the network, the IP routing table is updated, and those changes are reflected in the FIB. The FIB maintains next-hop address information based on the information in the IP routing table.

IGP—Internet protocol used to exchange routing information within an autonomous system. Examples of common Internet IGPs include OSPF and RIP.

Interior Gateway Protocol—See IGP.

metric—A metric is the method by which a routing algorithm determines that one route is better than another. The information is stored in routing tables. Metrics include bandwidth, communication cost, delay, hop count, load, and reliability.

MPLS—Multiprotocol Label Switching. A method for forwarding packets (frames) through a network. It enables routers at the edge of a network to apply labels to packets (frames). ATM switches or existing routers in the network core can switch packets according to the labels with minimal lookup overhead.

Multiprotocol Label Switching—See MPLS.

Open Shortest Path First—See OSPF.

OSPF—Link-state hierarchical IGP routing algorithm proposed as a successor to RIP in the Internet community. OSPF features include least-cost routing, multipath routing, and load balancing.

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

Note See *Internetworking Terms and Acronyms* for terms not included in this glossary.