



MPLS Traffic Engineering--Configurable Path Calculation Metric for Tunnels

Last Updated: November 30, 2011

The MPLS Traffic Engineering--Configurable Path Calculation Metric for Tunnels feature enables the user to control the metric used in path calculation for traffic engineering (TE) tunnels on a per-tunnel basis. Certain tunnels are used to carry voice traffic, which requires low delay, and other tunnels are used to carry data. A TE link metric can be used to represent link delay and configure tunnels that carry voice traffic for path calculation and configure tunnels that carry data to use the Interior Gateway Protocol (IGP) metric for path calculation.

- [Finding Feature Information, page 1](#)
- [Feature Overview, page 1](#)
- [Prerequisites, page 2](#)
- [Configuration Tasks, page 3](#)
- [Configuration Examples, page 8](#)
- [Additional References, page 10](#)
- [Feature Information for MPLS Traffic Engineering--Configurable Path Calculation Metric for Tunnels, page 11](#)

Finding Feature Information

Your software release may not support all the features documented in this module. For the latest feature information and caveats, see the release notes for your platform and software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the Feature Information Table at the end of this document.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Feature Overview

When Multiprotocol Label Switching (MPLS) 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



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

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 releases previous to Cisco IOS Release 12.2(18)ST, 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 enhancement as of release Cisco IOS Release 12.2(18)ST 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, page 2](#)
- [Restrictions, page 2](#)
- [Related Features and Technologies, page 2](#)

Benefits

When TE tunnels are used to carry two types of traffic, the Configurable Path Calculation Metric for Tunnels 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.
- You cannot configure MPLS Traffic Engineering over the logical GRE tunnel interface.
- MPLS traffic engineering supports only a single IGP process/instance. Multiple IGP processes/instances are not supported and MPLS traffic engineering should not be configured in more than one IGP process/instance.

Related Features and Technologies

The configurable path calculation metric feature is related to MPLS TE.

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.

- [Configuring a Platform to Support Traffic Engineering Tunnels, page 3](#)
- [Configuring IS-IS for MPLS Traffic Engineering, page 3](#)
- [Configuring OSPF for MPLS Traffic Engineering, page 4](#)
- [Configuring Traffic Engineering Link Metrics, page 5](#)
- [Configuring an MPLS Traffic Engineering Tunnel, page 6](#)
- [Configuring the Metric Type for Tunnel Path Calculation, page 7](#)
- [Verifying the Configuration, page 7](#)

Configuring a Platform to Support Traffic Engineering Tunnels

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

SUMMARY STEPS

1. Router(config)# **ip cef**
2. Router(config)# **mpls traffic-eng tunnels**

DETAILED STEPS

| | Command or Action | Purpose |
|---------------|---|--|
| Step 1 | Router(config)# ip cef | Enables standard Cisco Express Forwarding operation. |
| 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 following steps.



Note

MPLS traffic engineering supports only a single IGP process/instance. Multiple IGP processes/instances are not supported and MPLS traffic engineering should not be configured in more than one IGP process/instance.

SUMMARY STEPS

1. Router(config)# **router isis**
2. Router(config-router)# **mpls traffic-eng level-1**
3. Router(config-router)# **mpls traffic-eng level-2**
4. Router(config-router)# **mpls traffic-eng router-id loopback 0**
5. Router(config-router)# **metric-style wide**

DETAILED STEPS

| | Command or Action | Purpose |
|---------------|---|--|
| Step 1 | Router(config)# router isis | Enables IS-IS routing and specifies an IS-IS process for IP. The router is placed in 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 loopback 0 | 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**SUMMARY STEPS**

1. **enable**
2. **configure terminal**
3. **router ospf *process-id***
4. **mpls traffic-eng area *number***
5. **mpls traffic-eng router-id loopback0**
6. **exit**
7. **exit**

DETAILED STEPS

| | Command or Action | Purpose |
|---------------|---|--|
| Step 1 | enable Example: Router> enable | Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted. |

| Command or Action | Purpose |
|--|---|
| Step 2 <code>configure terminal</code> Example: <pre>Router# configure terminal</pre> | Enters global configuration mode. |
| Step 3 <code>router ospf process-id</code> Example: <pre>Router(config)# router ospf 200</pre> | Configures an OSPF routing process for IP and enters router configuration mode. <ul style="list-style-type: none"> The value for the <i>process-id</i> argument 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 4 <code>mpls traffic-eng area number</code> Example: <pre>Router(config-router)# mpls traffic-eng area 0</pre> | Turns on MPLS TE for the indicated OSPF area. |
| Step 5 <code>mpls traffic-eng router-id loopback0</code> Example: <pre>Router(config-router)# mpls traffic-eng router-id loopback0</pre> | Specifies that the TE router identifier for the node is the IP address associated with interface loopback0. |
| Step 6 <code>exit</code> Example: <pre>Router(config-router)# exit</pre> | Exits to global configuration mode. |
| Step 7 <code>exit</code> Example: <pre>Router(config)# exit</pre> | Exits to privileged EXEC mode. |

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:

SUMMARY STEPS

1. Router(config)# **interface** *interface*
2. Router(config-router)# **mpls traffic-eng administrative-weight** *weight*

DETAILED STEPS

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | Router(config)# interface <i>interface</i> | Configures properties of the specified interface. |
| Step 2 | Router(config-router)# mpls traffic-eng administrative-weight <i>weight</i> | 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.

SUMMARY STEPS

1. Router(config)# **interface** *tunnel-interface*
2. Router(config-if)# **ip unnumbered loopback0**
3. Router(config-if)# **tunnel destination** *ip-address*
4. Router(config-if)# **tunnel mode mpls traffic-eng**
5. Router(config-if)# **tunnel mpls traffic-eng bandwidth** *bandwidth*
6. Router(config-if)# **tunnel mpls traffic-eng path-option** *number* { **dynamic** | **explicit** { **name** *path-name* | **id** *path-number* } } [**lockdown**]

DETAILED STEPS

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | Router(config)# interface <i>tunnel-interface</i> | 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 <i>ip-address</i> | 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 <i>bandwidth</i> | 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. |

| Command or Action | Purpose |
|---|---|
| Step 6 Router(config-if)# tunnel mpls traffic-eng path-option <i>number</i> { dynamic explicit { name <i>path-name</i> id <i>path-number</i> }} [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 the 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 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.

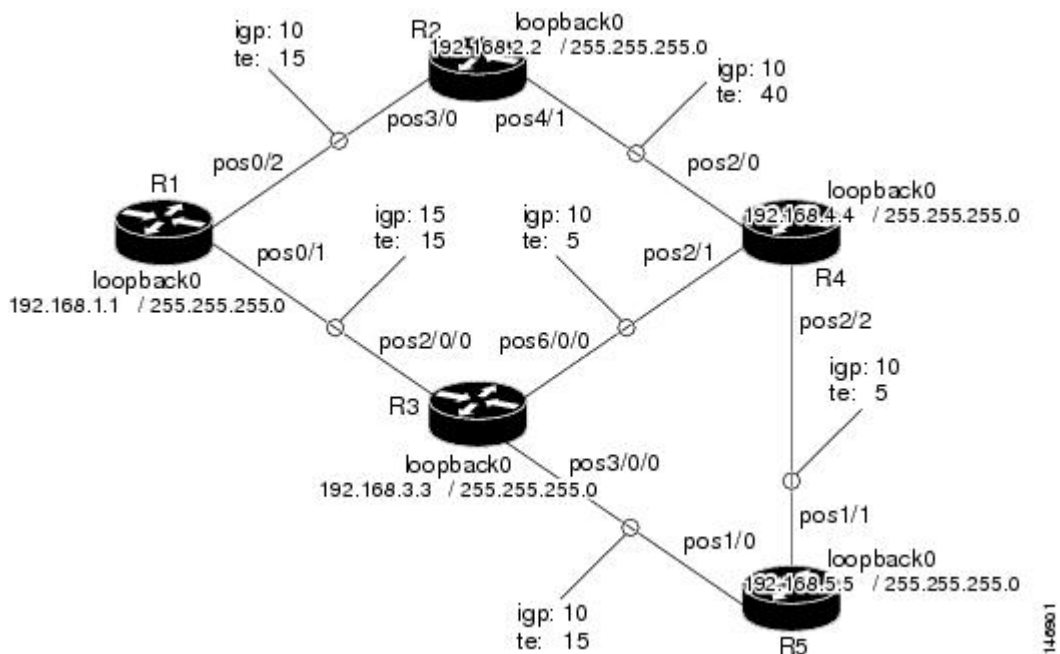
```
Router# show mpls traffic-eng tunnels
Name: te3640-17-c_t221                (Tunnel22) Destination: 192.168.100.22
Status:
  Admin: up          Oper: up          Path: valid          Signalling: connected
  path option 1, type dynamic (Basis for Setup, path weight 10)
Config Parameters:
  Bandwidth: 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 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.

The examples in this section support the simple network topology shown in the figure below.

Figure 1 Network Topology



In the figure above:

- 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 MPLS TE carries data traffic with no delay requirement.

Configuration fragments follow 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 15      !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

Related Document

| Related Topic | Document Title |
|---|---|
| MPLS Traffic Engineering Path Protection | MPLS Traffic Engineering Path Protection |
| MPLS Traffic Engineering--Fast Reroute Link and Node Protection | MPLS Traffic Engineering--Fast Reroute Link and Node Protection |

Standards

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

MIBs

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

RFCs

| RFC | Title |
|---|-------|
| No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified. | - |

Technical Assistance

| Description | Link |
|---|--|
| <p>The Cisco Support website provides extensive online resources, including documentation and tools for troubleshooting and resolving technical issues with Cisco products and technologies.</p> <p>To receive security and technical information about your products, you can subscribe to various services, such as the Product Alert Tool (accessed from Field Notices), the Cisco Technical Services Newsletter, and Really Simple Syndication (RSS) Feeds.</p> <p>Access to most tools on the Cisco Support website requires a Cisco.com user ID and password.</p> | <p>http://www.cisco.com/techsupport</p> |

Feature Information for MPLS Traffic Engineering--Configurable Path Calculation Metric for Tunnels

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Table 1 Feature Information for MPLS Traffic Engineering--Configurable Path Calculation Metric for Tunnels

| Feature Name | Releases | Feature Information |
|--|--|--|
| MPLS Traffic Engineering--Configurable Path Calculation Metric for Tunnels | 12.0(18)ST 12.2(11)S 12.2(14)S 12.2(28)SB 12.4(20)T | <p>This feature was introduced.</p> <p>This feature was integrated into Cisco IOS Release 12.2(11)S.</p> <p>This feature was integrated into Cisco IOS Release 12.2(14)S.</p> <p>This feature was integrated into Cisco IOS Release 12.2(28)SB.</p> <p>This feature was integrated into Cisco IOS Release 12.4(20)T.</p> |

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks.

Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

© 2011 Cisco Systems, Inc. All rights reserved.