



MPLS Traffic Engineering (TE)—Scalability Enhancements

First Published: February 23, 2002

Last Updated: February 28, 2006

The Multiprotocol Label Switching (MPLS) Traffic Engineering (TE) Scalability Enhancement feature improves scalability performance for large numbers of traffic engineering tunnels.

These improvements allow an increase in the number of TE tunnels a router can support when the router is configured as a tunnel headend. Additionally, when the router is configured as a tunnel midpoint, the enhancements reduce the time required to establish large numbers of TE tunnels.

This feature module describes the MPLS traffic engineering scalability enhancements.

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

This document contains the following sections:

- [Prerequisites for MPLS Traffic Engineering \(TE\)—Scalability Enhancements, page 2](#)
- [Restrictions for MPLS Traffic Engineering \(TE\)—Scalability Enhancements, page 2](#)
- [Information About MPLS Traffic Engineering \(TE\)—Scalability Enhancements, page 2](#)
- [How to Configure MPLS Traffic Engineering \(TE\)—Scalability Enhancements, page 4](#)
- [Additional References, page 6](#)
- [Command Reference, page 8](#)
- [Glossary, page 23](#)
- [Feature Information for Multiprotocol Label Switching \(MPLS\) Traffic Engineering \(TE\) Scalability Enhancement, page 24](#)



Corporate Headquarters:

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

© 2006 Cisco Systems, Inc. All rights reserved.

Prerequisites for MPLS Traffic Engineering (TE)—Scalability Enhancements

Your network must support the following Cisco IOS features before you enable MPLS traffic engineering:

- MPLS
- IP Cisco Express Forwarding
- Intermediate System-to-Intermediate System (IS-IS) or Open Shortest Path First (OSPF)

Restrictions for MPLS Traffic Engineering (TE)—Scalability Enhancements

The number of tunnels that a particular platform can support can vary depending on:

- The types of interfaces that the tunnels traverse
- The manner in which the Resource Reservation Protocol (RSVP) message pacing feature is configured

Information About MPLS Traffic Engineering (TE)—Scalability Enhancements

Scalability performance is improved for large numbers of traffic engineering tunnels, and includes the following enhancements:

- Increase the number of traffic engineering tunnels a router can support when configured as a tunnel headend and when configured as a tunnel midpoint
- Reduce the time required to establish large numbers of traffic engineering tunnels

The following sections describe user-observable scalability enhancements:

- [Pacing for RSVP Messages, page 2](#)
- [Signaling and Management for MPLS Traffic Engineering Tunnels, page 3](#)
- [Controlling IS-IS and MPLS Traffic Engineering Topology Database Interactions, page 3](#)
- [Improved Diagnostic Capabilities for MPLS Traffic Engineering and RSVP Signaling, page 3](#)
- [Benefits of MPLS Traffic Engineering \(TE\)—Scalability Enhancements, page 3](#)

Pacing for RSVP Messages

A burst of RSVP traffic engineering signaling messages can overflow the input queue of a receiving router, causing some messages to be dropped. Dropped messages cause a substantial delay in completing label-switched path (LSP) signaling.

This feature provides an enhancement mechanism that controls the transmission rate for RSVP messages and reduces the likelihood of input drops on the receiving router. The default transmission rate is 200 RSVP messages per second to a given neighbor. The rate is configurable.

Signaling and Management for MPLS Traffic Engineering Tunnels

LSP recovery responsiveness is improved when a link used by an LSP fails:

- When the upstream end of a failed link detects the failure, it generates an RSVP No Route path error message. This enables the LSP headend to detect the link failure and initiate recovery, even when the Interior Gateway Protocol (IGP) update announcing the link failure is delayed.
- The LSP headend marks the link in question so that subsequent constraint-based shortest path first (SPF) calculations ignore the link until either a new IGP update arrives or a configurable timeout occurs. This ensures that resignaling to restore the LSP avoids the failed link.

Controlling IS-IS and MPLS Traffic Engineering Topology Database Interactions

This feature reduces the interval between when the IS-IS protocol receives an IGP update and when it delivers the update to the MPLS traffic engineering topology database.

Before this feature was introduced, when IS-IS received a new LSP that contained traffic engineering type, length, and value (TLV) objects, a delay of several seconds could occur before IS-IS passed the traffic engineering TLVs to the traffic engineering database. The purpose of the delay was to provide better scalability during periods of network instability and to give the router an opportunity to receive more fragments of the LSP before passing the information to the traffic engineering database. However, this delay increased the convergence time for the traffic engineering database.

With this feature, IS-IS extracts traffic engineering TLVs from received LSPs and passes them to the traffic engineering database immediately. The exception to this occurs when there are large numbers of LSPs to process and it is important to limit CPU consumption, such as during periods of network instability. The parameters that control IS-IS delivery of traffic engineering TLVs to the traffic engineering topology database are configurable.

Improved Diagnostic Capabilities for MPLS Traffic Engineering and RSVP Signaling

With this feature, diagnostic and troubleshooting capabilities for MPLS traffic engineering and RSVP are improved:

- Counters record tunnel headend error events such as no route (link down), preemption, and insufficient bandwidth on a per-tunnel basis.
- Counters record RSVP messages. The counters are per-interface and record the number of RSVP messages of each type sent and received on the interface.

Benefits of MPLS Traffic Engineering (TE)—Scalability Enhancements

This feature provide the following benefits:

- Increased scalability—Up to 600 MPLS traffic engineering tunnel headends are supported. Up to 10,000 traffic engineering tunnel midpoints are supported, with up to 5000 midpoints per interface.
- Faster recovery after failure conditions—Message pacing provides a mechanism to throttle RSVP control messages so that they are less likely to be dropped. This results in a faster recovery from failure conditions when many MPLS traffic engineering tunnels are being set up.
- Improved reroute time—When a traffic engineering tunnel is down, the headend router needs to be notified so that it can signal for a new LSP for the tunnel along an alternate path. The headend router does not have to wait for an IGP update to signal for a new LSP for the tunnel along an alternate path.
- Improved tunnel setup time—Fewer control messages and tunnel setup messages are dropped. This reduces the average time required to set up tunnels.

How to Configure MPLS Traffic Engineering (TE)—Scalability Enhancements

This section describes the following tasks:

- [Enabling RSVP Message Pacing, page 4](#)
- [Monitoring and Maintaining Scalability Enhancements, page 5](#)

Enabling RSVP Message Pacing

RSVP message pacing maintains, on an outgoing interface basis, a count of messages that were dropped because the output queue for the interface used for message pacing was full.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **ip rsvp msg-pacing [period *ms* [burst *msgs* [maxsize *qsize*]]]**
4. **end**
5. **show ip rsvp neighbor**

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.
Step 2	configure terminal Example: Router# configure terminal	Enters global configuration mode.

	Command or Action	Purpose
Step 3	ip rsvp msg-pacing [<i>period ms</i> [<i>burst msgs</i> [<i>maxsize qsize</i>]]] Example: Router(config)# ip rsvp msg-pacing period 2 burst 5 maxsize 3	Enables RSVP message pacing.
Step 4	end Example: Router(config)# end	Returns to privileged EXEC mode.
Step 5	show ip rsvp neighbor Example: Router# show ip rsvp neighbor	Verifies that RSVP message pacing is enabled.

Examples

The following command enables RSVP message pacing:

```
Router(config)# ip rsvp msg-pacing
```

Enter the following command to verify that RSVP message pacing is enabled:

```
Router# show ip rsvp neighbor
```

The following is sample output that traffic engineering displays when RSVP message pacing is enabled:

```
RSVP Msg-Pacing:Period 20, Burst 4, Rate 200, MaxSize 500
```

Interface	Neighbor	Encapsulation	OQueue	OQueue-drops
AT0.1	10.12.9.12	RSVP	0	0
Hs0	10.13.9.13	RSVP	0	0

If RSVP message pacing is not enabled, the OQueue or OQueue-drops column do not display.

Monitoring and Maintaining Scalability Enhancements

To monitor and maintain the scalability enhancements, use the commands shown below. For detailed command descriptions, see the “[Command Reference](#)” section.

Command	Purpose
show ip rsvp neighbor	Verifies that RSVP message pacing is turned on.
show ip rsvp counters	Displays the counts of RSVP messages that were sent and received.
clear ip rsvp counters	Clears (sets to zero) all IP RSVP counters that are being maintained.

Additional References

Command	Purpose
<code>clear ip rsvp msg-pacing</code>	Clears (sets to zero) counts of the messages that message pacing was forced to drop because the output queue for the interface used for message pacing was full.
<code>show mpls traffic-eng tunnels statistics</code>	Displays event counters for one or more MPLS traffic engineering tunnels.
<code>clear mpls traffic-eng tunnel counters</code>	Clears the counters for all MPLS traffic engineering tunnels.

Additional References

The following sections provide references related to the MPLS Traffic Engineering (TE)—Scalability Enhancements feature.

Related Documents

Related Topic	Document Title
Quality of service	<ul style="list-style-type: none"> • Cisco IOS Quality of Service Solutions Command Reference, Release 12.4 • Cisco IOS Quality of Service Solutions Configuration Guide, Release 12.4
MPLS	<ul style="list-style-type: none"> • Cisco IOS Multiprotocol Label Switching Command Reference, Release 12.4 • Cisco IOS Multiprotocol Label Switching Configuration Guide, Release 12.4

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 by this feature.	—

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 the following modified commands only:

- **clear ip rsvp counters**
- **clear ip rsvp msg-pacing**
- **clear mpls traffic-eng tunnel counters**
- **ip rsvp msg-pacing**
- **mpls traffic-eng scanner**
- **mpls traffic-eng topology holddown sigerr**
- **show ip rsvp counters**
- **show mpls traffic-eng tunnels statistics**

clear ip rsvp counters

To clear (set to zero) all IP Resource Reservation Protocol (RSVP) counters that are being maintained, use the **clear ip rsvp counters** command in EXEC mode.

clear ip rsvp counters [confirm]

Syntax Description	confirm	(Optional) Requests a confirmation that all IP RSVP counters were cleared.
--------------------	---------	--

Command Modes	EXEC
---------------	------

Command History	Release	Modification
	12.0(14)ST	This command was introduced.
	12.2(14)S	This command was integrated into Cisco IOS Release 12.2(14)S.
	12.0(22)S	This command was integrated into Cisco IOS Release 12.0(22)S.
	12.2(28)SB	This command was integrated into Cisco IOS Release 12.2(28)SB.

Usage Guidelines	This command allows you to set all IP RSVP counters to zero so that you can see changes easily.
------------------	---

Examples	In the following example, all IP RSVP counters that are being maintained are cleared:
----------	---

```
Router# clear ip rsvp counters
Clear rsvp counters [confirm]
```

Related Commands	Command	Description
	show ip rsvp counters	Displays counts of RSVP messages that were sent and received.

clear ip rsvp msg-pacing

To clear the Resource Reservation Protocol (RSVP) message pacing output from the **show ip rsvp neighbor** command, use the **clear ip rsvp msg-pacing** command in EXEC mode.

clear ip rsvp msg-pacing

Syntax Description This command has no arguments or keywords.

Command Modes EXEC

Command History

Release	Modification
12.0(14)ST	This command was introduced.
12.2(14)S	This command was integrated into Cisco IOS Release 12.2(14)S.
12.0(22)S	This command was integrated into Cisco IOS Release 12.0(22)S.
12.2(28)SB	This command was integrated into Cisco IOS Release 12.2(28)SB.

Examples

The following example clears the RSVP message pacing output:

```
Router# clear ip rsvp msg-pacing
```

Related Commands

Command	Description
show ip rsvp counters	Displays counts of RSVP messages that were sent and received.
show ip rsvp neighbor	Displays the current RSVP neighbors and indicates whether the neighbor is using IP or UDP encapsulation for a specified interface or for all interfaces.

clear mpls traffic-eng tunnel counters

To clear the counters for all Multiprotocol Label Switching (MPLS) traffic engineering tunnels, use the **clear mpls traffic-eng tunnel counters** command in EXEC mode.

clear mpls traffic-eng tunnel counters

Syntax Description

This command has no optional parameters or keywords.

Command Modes

EXEC

Command History

Release	Modification
12.0(14)ST	This command was introduced.
12.2(14)S	This command was integrated into Cisco IOS Release 12.2(14)S.
12.0(22)S	This command was integrated into Cisco IOS Release 12.0(22)S.
12.2(28)SB	This command was integrated into Cisco IOS Release 12.2(28)SB.

Usage Guidelines

This command allows you to set the MPLS traffic engineering tunnel counters to zero so that you can see changes to the counters easily.

Examples

In the following example, the counters for all MPLS traffic engineering tunnels are cleared and a request is made for confirmation that the specified action occurred:

```
Router# clear mpls traffic-eng tunnel counters

Clear traffic engineering tunnel counters [confirm]
```

Related Commands

Command	Description
show mpls traffic-eng tunnels statistics	Displays event counters for one or more MPLS traffic engineering tunnels.

ip rsvp msg-pacing

To set up message pacing (that is, to control the transmission rate for Resource Reservation Protocol (RSVP) messages), use the **ip rsvp msg-pacing** command in global configuration mode. To disable this feature, use the **no** form of this command.

```
ip rsvp msg-pacing [period ms [burst msgs [maxsize qsize]]]
```

```
no rsvp msg-pacing
```

Syntax Description

period <i>ms</i>	(Optional) Length of the interval, in milliseconds, during which a router can send the number of RSVP messages specified in the <i>burst</i> keyword. The value can be from 1 to 1000 milliseconds.
burst <i>msgs</i>	(Optional) Maximum number of RSVP messages that a router can send to an output interface during each interval specified in the <i>period</i> keyword. The value can be from 1 to 2000.
maxsize <i>qsize</i>	(Optional) Size of per-interface output queues in the sending router. Valid values are from 1 to 2000.

Defaults

RSVP messages are not paced.

If you enter the command without the optional arguments, the transmission rate for RSVP messages is limited to 200 messages per second per outgoing interface.

The default output queue size, specified in the **maxsize** keyword, is 500.

Command Modes

Global configuration

Command History

Release	Modification
12.0(14)ST	This command was introduced.
12.2(11)S	This command was integrated into Cisco IOS Release 12.2(11)S.
12.0(22)S	This command was integrated into Cisco IOS Release 12.0(22)S.
12.2(28)SB	This command was integrated into Cisco IOS Release 12.2(28)SB.

Usage Guidelines

You can use this command to prevent a burst of RSVP traffic engineering signaling messages from overflowing the input queue of a receiving router, which would cause the router to drop some messages. Dropped messages substantially delay the completion of signaling for LSPs for which messages have been dropped.

Examples

In the following example, a router can send a maximum of 150 RSVP traffic engineering signaling messages in 1 second to a neighbor, and the size of the output queue is 750:

```
Router(config)# ip rsvp msg-pacing period 1 burst 150 maxsize 750
```

Related Commands	Command	Description
	clear ip rsvp msg-pacing	Clears the RSVP message pacing output from the show ip rsvp neighbor command.

mpls traffic-eng scanner

To specify how often Intermediate System-to-Intermediate System (IS-IS) extracts traffic engineering type, length, and value (TLVs) from flagged label-switched path (LSPs) and passes them to the traffic engineering topology database, and the maximum number of LSPs that the router can process immediately, use the **mpls traffic-eng scanner** command router IS-IS configuration subcommand. To disable this feature, use the **no** form of this command.

mpls traffic-eng scanner [*interval sec*] [**max-flash** *LSPs*]

no mpls traffic-eng scanner

Syntax Description

interval <i>seconds</i>	(Optional) Frequency, in seconds, at which IS-IS sends traffic engineering TLVs into the traffic engineering database. The value can be from 1 to 60.
max-flash <i>LSPs</i>	(Optional) Maximum number of LSPs that the router can process immediately without incurring a delay. The value can be from 0 to 200.

Defaults

The default **interval** is 5 seconds.

The default **max-flash** value is 15 LSPs.

The first 15 LSPs are sent without a delay into the traffic engineering database. If more LSPs are received, the default delay of 5 seconds applies.

If you specify the **no** form of this command, there is a delay of 5 seconds before IS-IS scans its database and passes traffic engineering TLVs associated with flagged LSPs to the traffic engineering database.

Command Modes

Router IS-IS configuration

Command History

Release	Modification
12.0(14)ST	This command was introduced.
12.2(11)S	This command was integrated into Cisco IOS Release 12.2(11)S.
12.0(22)S	This command was integrated into Cisco IOS Release 12.0(22)S.
12.2(28)SB	This command was integrated into Cisco IOS Release 12.2(28)SB.

Usage Guidelines

When IS-IS receives a new LSP, it inserts it into the IS-IS database. If the LSP contains traffic engineering TLVs, IS-IS flags the LSPs for transmission to the traffic engineering database. At the default or user-specified interval, traffic engineering TLVs are extracted and sent to the traffic engineering database. Users can also specify the maximum number of LSPs that the router can process immediately. Processing entails checking for traffic engineering TLVs, extracting them, and passing them to the traffic engineering database. If more than 50 LSPs need to be processed, there is a delay of 5 seconds for subsequent LSPs.

Examples

In the following example, the router is allowed to process up to 50 IS-IS LSPs without any delay.

```
Router(isis)# mpls traffic-eng scanner interval 5 max-flash 50
```

Related Commands

Command	Description
mpls traffic-eng	Configures a router running IS-IS so that it floods MPLS traffic engineering link information into the indicated IS-IS level.
mpls traffic-eng router-id	Specifies that the traffic engineering router identifier for the node is the IP address associated with a given interface.

mpls traffic-eng topology holddown sigerr

To specify the amount of time that a router ignores a link in its traffic engineering topology database in tunnel path Constrained Shortest Path First (CSPF) computations following a traffic engineering tunnel error on the link, use the **mpls traffic-eng topology holddown sigerr** command in global configuration mode. To disable this feature, use the **no** form of this command.

mpls traffic-eng topology holddown sigerr *seconds*

no mpls traffic-eng topology holddown sigerr

Syntax Description

<i>seconds</i>	Length of time (in seconds) a router should ignore a link during tunnel path calculations following a traffic engineering tunnel error on the link. The value can be from 0 to 300.
----------------	---

Defaults

If you do not specify this command, tunnel path calculations ignore a link on which there is a traffic engineering error until either 10 seconds have elapsed or a topology update is received from the Interior Gateway Protocol (IGP).

Command Modes

Global configuration

Command History

Release	Modification
12.0(14)ST	This command was introduced.
12.2(11)S	This command was integrated into Cisco IOS Release 12.2(11)S.
12.0(22)S	This command was integrated into Cisco IOS Release 12.0(22)S.
12.2(28)SB	This command was integrated into Cisco IOS Release 12.2(28)SB.

Usage Guidelines

A router that is at the headend for traffic engineering tunnels might receive a Resource Reservation Protocol (RSVP) No Route error message for an existing tunnel or for one being signaled due to the failure of a link the tunnel traverses before the router receives a topology update from the IGP routing protocol announcing that the link is down. In such a case, the headend router ignores the link in subsequent tunnel path calculations to avoid generating paths that include the link and are likely to fail when signaled. The link is ignored until the router receives a topology update from its IGP or a link hold-down timeout occurs. You can use the **mpls traffic-eng topology holddown sigerr** command to change the link hold-down time from its 10 second default value.

Examples

In the following example, the link hold-down time for signaling errors is set at 15 seconds:

```
Router(config)# mpls traffic-eng topology holddown sigerr 15
```


Related Commands	Command	Description
	show mpls traffic-eng topology	Displays the MPLS traffic engineering global topology as currently known at the node.

show ip rsvp counters

To display the counts of Resource Reservation Protocol (RSVP) messages that were sent and received, use the **show ip rsvp counters** command in EXEC mode.

show ip rsvp counters [**interface** *interface-name*] **summary**]

Syntax Description

interface <i>interface-name</i>	(Optional) Displays the number of RSVP messages sent and received for the specified interface name.
summary	(Optional) Displays the cumulative number of RSVP messages sent and received by the platform.

Defaults

If you enter this command without a keyword, the command displays the number of RSVP messages that were sent and received for each interface for which RSVP is configured.

Command Modes

EXEC

Command History

Release	Modification
12.0(14)ST	This command was introduced.
12.2(11)S	This command was integrated into Cisco IOS Release 12.2(11)S.
12.0(22)S	This command was integrated into Cisco IOS Release 12.0(22)S.
12.2(28)SB	This command was integrated into Cisco IOS Release 12.2(28)SB.

Examples

In the following example, values are shown for the number of RSVP messages of each type (such as Path and Resv) that were sent and received on POS2/2:

```
Router# show ip rsvp counters interface pos2/2
```

```

POS2/2                Recv      Xmit
  Path                 0      112876  Resv                 74414      0
  PathError            0         0  ResvError             0         0
  PathTear              0         7  ResvTear               4         0
  ResvConfirm           0         0  ResvTearConfirm        0         4
  UnknownMsg           0         0  Errors                 0         0

```

In the following example, values are shown for the number of RSVP messages of each type that were sent and received by the router over all interfaces:

```
Router# show ip rsvp counters summary
```

```

All Interfaces        Recv      Xmit
  Path                 512     224962  Resv                 148446     385
  PathError            0         2  ResvError             0         0
  PathTear              8        12  ResvTear               4         4
  ResvConfirm           0         0  ResvTearConfirm        4         4
  UnknownMsg           0         0  Errors                 0         0

```

Table 1 describes the significant fields shown in the display.

Table 1 *show ip rsvp counters Field Descriptions*

Field	Description
Recv	Number of messages received at the specified interface or at all interfaces.
Xmit	Number of messages sent from the specified interface or from all interfaces.
All Interfaces	Specifies that the number of messages displayed is for all interfaces.

Related Commands

Command	Description
clear ip rsvp counters	Clears (sets to zero) all IP RSVP counters that are being maintained.

show mpls traffic-eng tunnels statistics

To display event counters for one or more Multiprotocol Label Switching (MPLS) traffic engineering tunnels, use the **show mpls traffic-eng tunnels statistics** command in EXEC mode.

show mpls traffic-eng tunnels [**tunnel** *tunnel-name*] **statistics** [**summary**]

Syntax Description

tunnel <i>tunnel-name</i>	(Optional) Displays event counters accumulated for the specified tunnel.
summary	(Optional) Displays event counters accumulated for all tunnels.

Defaults

If you enter the command without any keywords, the command displays the event counters for every MPLS traffic engineering tunnel interface configured on the router.

Command Modes

EXEC

Command History

Release	Modification
12.0(14)ST	This command was introduced.
12.2(11)S	This command was integrated into Cisco IOS Release 12.2(11)S.
12.0(22)S	This command was integrated into Cisco IOS Release 12.0(22)S.
12.2(28)SB	This command was integrated into Cisco IOS Release 12.2(28)SB.

Usage Guidelines

A label switching router (LSR) maintains counters for each MPLS traffic engineering tunnel headend that counts significant events for the tunnel, such as state transitions for the tunnel, changes to the tunnel path, and various signaling failures. You can use the **show mpls traffic-eng tunnels statistics** command to display these counters for a single tunnel, for every tunnel, or for all tunnels (accumulated values). Displaying the counters is often useful for troubleshooting tunnel problems.

Examples

The following are examples of output from the **show mpls traffic-eng tunnels statistics** command:

```
Router# show mpls traffic-eng tunnels tunnel tunnel1001 statistics

Tunnel1001 (Destination 10.8.8.8; Name Router_t1001)
  Management statistics:
    Path:25 no path, 1 path no longer valid, 0 missing ip exp path
    5 path changes
    State:3 transitions, 0 admin down, 1 oper down
  Signalling statistics:
    Opens:2 succeeded, 0 timed out, 0 bad path spec
    0 other aborts
    Errors:0 no b/w, 0 no route, 0 admin
    0 bad exp route, 0 rec route loop, 0 other
```

```

Router# show mpls traffic-eng tunnels statistics

Tunnel1001 (Destination 10.8.8.8; Name Router_t1001)
  Management statistics:
    Path:25 no path, 1 path no longer valid, 0 missing ip exp path
    5 path changes
    State:3 transitions, 0 admin down, 1 oper down
  Signalling statistics:
    Opens:2 succeeded, 0 timed out, 0 bad path spec
    0 other aborts
    Errors:0 no b/w, 0 no route, 0 admin
    0 bad exp route, 0 rec route loop, 0 other
  .
  .
  .

Tunnel7050 (Destination 10.8.8.8; Name Router_t7050)
  Management statistics:
    Path: 19 no path, 1 path no longer valid, 0 missing ip exp path
    3 path changes
    State: 3 transitions, 0 admin down, 1 oper down
  Signalling statistics:
    Opens: 2 succeeded, 0 timed out, 0 bad path spec
    0 other aborts
    Errors:0 no b/w, 0 no route, 0 admin
    0 bad exp route, 0 rec route loop, 0 other

Router# show mpls traffic-eng tunnels statistics summary

```

```

Management statistics:
  Path:2304 no path, 73 path no longer valid, 0 missing ip exp path
  432 path changes
  State:300 transitions, 0 admin down, 100 oper down
Signalling statistics:
  Opens:200 succeeded, 0 timed out, 0 bad path spec
  0 other aborts
  Errors:0 no b/w, 18 no route, 0 admin
  0 bad exp route, 0 rec route loop, 0 other

```

Table 2 describes the significant fields shown in the display.

Table 2 *show mpls traffic-eng tunnels statistics Field Descriptions*

Field	Description
Tunnel <i>tunnel-name</i>	Name of the tunnel interface.
Destination	IP address of the tunnel tailend.
Name	Internal name for the tunnel, composed of the router name and the tunnel interface number.

Table 2 *show mpls traffic-eng tunnels statistics Field Descriptions (continued)*

Field	Description
Path	Heading for counters for tunnel path events, including: <ul style="list-style-type: none"> • no path—Number of unsuccessful attempts to calculate a path for the tunnel. • path no longer valid—Number of times a previously valid path for the tunnel became invalid. • missing ip exp path—Number of times that attempts to use “obtain a path for the tunnel” failed because no path was configured (and there was no dynamic path option for the tunnel). • path changes—Number of times the tunnel path changed.
State	Heading for counters for tunnel state transitions.
Opens	Heading for counters for tunnel open attempt events.
Errors	Heading for various tunnel signaling errors, such as no bandwidth, no route, admin (preemption), a bad explicit route, a loop in the explicit route, and so forth.

Related Commands

Command	Description
clear mpls traffic-eng tunnel counters	Clears the counters for all MPLS traffic engineering tunnels.

Glossary

Cisco Express Forwarding—A means for accelerating the forwarding of packets within a router, by storing route lookup information in several data structures instead of in a route cache.

enterprise network—A large and diverse network connecting most major points in a company or other organization.

headend—The endpoint of a broadband network. All stations send toward the headend; the headend then sends toward the destination stations.

IGP—Interior Gateway Protocol. An Internet protocol used to exchange routing information within an autonomous system. Examples of common Internet IGPs include IGRP, OSPF, and RIP.

interface—A network connection.

IS-IS—Intermediate System-to-Intermediate System. OSI link-state hierarchical routing protocol based on DECnet Phase V routing, where ISs (routers) exchange routing information based on a single metric, to determine the network topology.

label-switched path (LSP)—A sequence of hops (R0...Rn) in which a packet travels from R0 to Rn through label switching mechanisms. A label-switched path can be chosen dynamically, based on normal routing mechanisms, or through configuration.

MPLS—Multiprotocol Label Switching (formerly known as tag switching). A method for directing packets primarily through Layer 2 switching rather than Layer 3 routing. In MPLS, packets are assigned short fixed-length labels at the ingress to an MPLS cloud by using the concept of forwarding equivalence classes. Within the MPLS domain, the labels are used to make forwarding decisions mostly without recourse to the original packet headers.

router—A network layer device that uses one or more metrics to determine the optimal path along which network traffic should be forwarded. Routers forward packets from one network to another based on network layer information.

RSVP—Resource Reservation Protocol. A protocol that supports the reservation of resources across an IP network.

scalability—An indicator showing how quickly some measure of resource usage increases as a network gets larger.

topology—The physical arrangement of network nodes and media within an enterprise networking structure.

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

traffic engineering tunnel—A label-switched tunnel that is used for traffic engineering. Such a tunnel is set up through means other than normal Layer 3 routing; it is used to direct traffic over a path different from the one that Layer 3 routing would cause the tunnel to take.

Feature Information for Multiprotocol Label Switching (MPLS) Traffic Engineering (TE) Scalability Enhancement

Table 3 lists the release history for this feature.

Not all commands may be available in your Cisco IOS software release. For release information about a specific command, see the command reference documentation.

Cisco IOS software images are specific to a Cisco IOS software release, a feature set, and a platform. 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.

**Note**

Table 3 lists only the Cisco IOS software release that introduced support for a given feature in a given Cisco IOS software release train. Unless noted otherwise, subsequent releases of that Cisco IOS software release train also support that feature.

Table 3 *Feature Information for Multiprotocol Label Switching (MPLS) Traffic Engineering (TE) Scalability Enhancement*

Feature Name	Releases	Feature Information
Multiprotocol Label Switching (MPLS) Traffic Engineering (TE) Scalability Enhancement	12.0(14)ST 12.2(14)S 12.0(22)S 12.2(28)SB	<p>The Multiprotocol Label Switching (MPLS) Traffic Engineering (TE) Scalability Enhancement feature improves scalability performance for large numbers of traffic engineering tunnels.</p> <p>These improvements allow an increase in the number of TE tunnels a router can support when the router is configured as a tunnel headend. Additionally, when the router is configured as a tunnel midpoint, the enhancements reduce the time required to establish large numbers of TE tunnels.</p> <p>In 12.0(14)ST, this feature was introduced.</p> <p>In 12.2(14)S, this feature was integrated into Cisco IOS Release 12.2(14)S.</p> <p>In 12.0(22)S, this feature was integrated into Cisco IOS Release 12.0(22)S.</p> <p>In 12.2(28)SB, this feature was integrated into Cisco IOS Release 12.2(28)SB.</p> <p>The following sections provide information about this feature:</p> <ul style="list-style-type: none"> • Prerequisites for MPLS Traffic Engineering (TE)—Scalability Enhancements, page 2 • Restrictions for MPLS Traffic Engineering (TE)—Scalability Enhancements, page 2 • Information About MPLS Traffic Engineering (TE)—Scalability Enhancements, page 2 • How to Configure MPLS Traffic Engineering (TE)—Scalability Enhancements, page 4 • Additional References, page 6 • Command Reference, page 8 • Glossary, page 23 • Feature Information for Multiprotocol Label Switching (MPLS) Traffic Engineering (TE) Scalability Enhancement, page 24

CCVP, the Cisco logo, and Welcome to the Human Network are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn is a service mark of Cisco Systems, Inc.; and Access Registrar, Aironet, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, iQuick Study, LightStream, Linksys, MeetingPlace, MGX, Networkers, Networking Academy, Network Registrar, PIX, ProConnect, ScriptShare, SMARTnet, StackWise, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website 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. (0711R)

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

© 2006 Cisco Systems, Inc. All rights reserved.