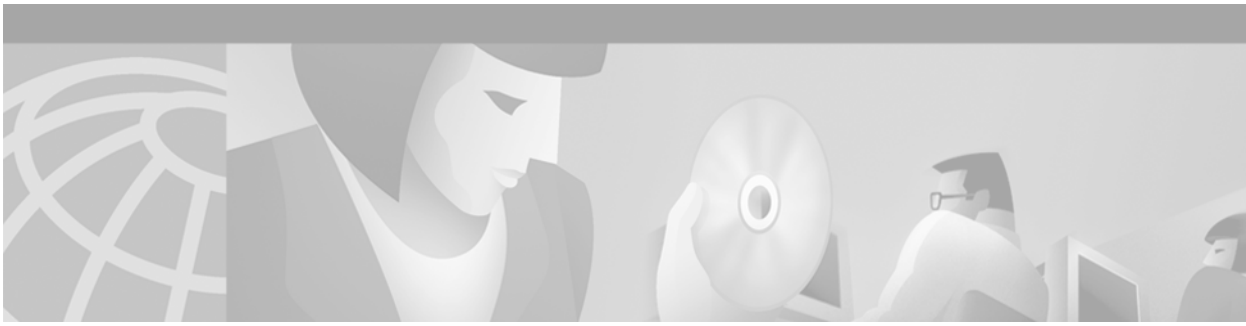

MPLS Traffic Engineering (TE)—Link and Node Protection, with RSVP Hellos Support

Feature History

Release	Modification
12.0(10)ST	Fast ReRoute Link Protection feature was introduced.
12.0(16)ST	Link Protection for Cisco Series 7200 and 7500 platforms was added.
12.0(22)S	Fast ReRoute enhancements were added.

This document describes the Fast ReRoute (FRR) enhancements, including Node Protection, in Cisco IOS Release 12.0(22)S. Node Protection can be viewed as a superset of (that is, an enhancement to) FRR Link Protection. The document includes the following sections:

Note If you plan to use or already use MPLS Traffic Engineering Fast ReRoute Link Protection before Release 12.0(22)S, contact Cisco TAC Support for important deployment and upgrade information.

Feature Overview

Fast ReRoute (FRR) is a mechanism for protecting MPLS Traffic Engineering (TE) label-switched paths (LSPs) from link and node failures by locally repairing the LSPs at the point of failure, allowing data to continue to flow on them while their headend routers attempt to establish new end-to-end LSPs to replace them. FRR locally repairs the protected LSPs by rerouting them over backup tunnels that bypass failed links or nodes.

Backup tunnels that bypass only a single link of the LSP's path provide Link Protection. They protect LSPs if a link along their path fails by rerouting the LSP's traffic to the next hop (bypassing the failed link). These are referred to as next-hop (NHOP) backup tunnels because they terminate at the LSP's next hop beyond the point of failure. Figure 1 illustrates a next-hop backup tunnel.

Figure 1 Next-Hop Backup Tunnel

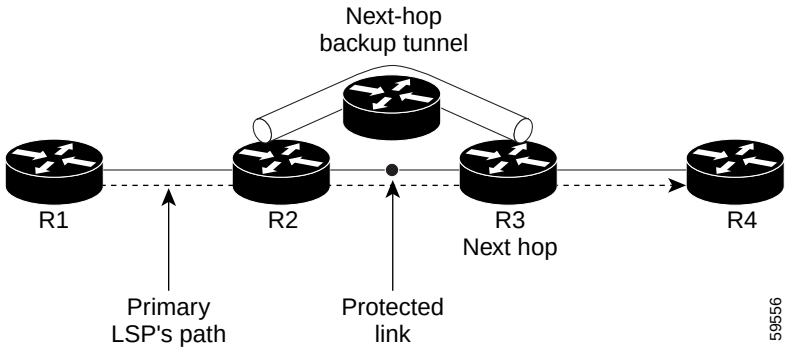

FRR provides Node Protection for LSPs. Backup tunnels that bypass next-hop nodes along LSP paths are called next-next-hop (NNH) backup tunnels because they terminate at the node following the next-hop node of the LSP paths, thereby bypassing the next-hop node. They protect LSPs if a node along their path fails by enabling the node upstream of the failure to reroute the LSPs and their traffic around the failed node to the next-next hop. FRR supports the use of RSVP Hellos to accelerate the detection of node failures. NNHOP backup tunnels also provide protection from link failures, because they bypass the failed link as well as the node.

Figure 2 illustrates a next-next hop backup tunnel.

Figure 2 Next-Next Hop Backup Tunnel

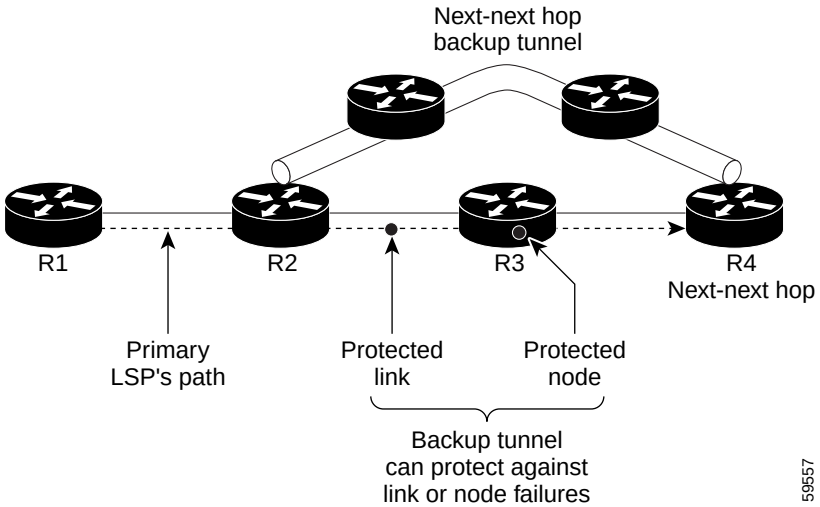

If an LSP is using a backup tunnel and something changes so that the LSP is no longer appropriate for the backup tunnel, the LSP is torn down. Such changes include the following:

- Backup bandwidth of the backup tunnel is reduced.
- Backup bandwidth type of backup tunnel is changed to a type that is incompatible with the primary LSP.
- Primary LSP is modified so that Fast ReRoute is disabled. (The **no mpls traffic-eng fast-reroute** command is entered.)

The Fast ReRoute enhancements include the following:

- **Backup tunnel support**—Backup tunnels can terminate at the next-next hop to support FRR.
- **Multiple backup tunnels**—There no longer is a limit (except memory limitations) to the number of backup tunnels that can protect a given interface. In many topologies, support for Node Protection requires supporting multiple backup tunnels per protected interface. These backup tunnels can terminate at the same destination or at different destinations. That is, for a given protected interface, you can configure multiple NHOP or NNHOP backup tunnels. This allows redundancy and load balancing (see "Benefits").
- **Bandwidth protection on backup tunnels**—NHOP and NNHOP backup tunnels can be used to provide bandwidth protection for rerouted LSPs. This is referred to as backup-bandwidth. You can associate backup-bandwidth with NHOP or NNHOP backup tunnels. This informs the router of the amount of backup-bandwidth a particular backup tunnel can protect. When a router maps LSPs to backup tunnels, bandwidth protection ensures that an LSP uses a given backup tunnel only if there is sufficient backup-bandwidth. The router selects which LSPs use which backup tunnels in order to provide maximum bandwidth protection. That is, the router determines the best way to map LSPs onto backup tunnels in order to maximize the number of LSPs that can be protected. For information about mapping tunnels and assigning backup-bandwidth, see "Backup Tunnel Selection Procedure".
- **Bandwidth pool restrictions for backup tunnels**—You can restrict the types of LSPs that can use a given backup tunnel. Backup tunnels can be restricted so that only LSPs using sub-pool bandwidth can use them or only LSPs that use global-pool bandwidth can use them. This allows different backup tunnels to be used for voice and data. Example: The backup tunnel used for voice could provide bandwidth protection, and the backup tunnel used for data could (optionally) not provide bandwidth protection.
- **Semi-dynamic backup tunnel paths**—The path of a backup tunnel can be configured to be determined dynamically. This can be done by using the IP explicit address exclusion feature that was added in Release 12.0(14)ST. Using this feature, semi-dynamic NHOP backup tunnel paths can be specified simply by excluding the protected link; semi-dynamic NNHOP backup tunnel paths can be configured simply by excluding the protected node.
- **RSVP Hello**—RSVP Hello enables RSVP nodes to detect when a neighboring node is not reachable. This feature is useful when next-hop node failure is not detectable by link layer mechanisms, or when notification of link-layer failures is not available (for example, Gigabit Ethernet).

Benefits

Node Protection

Backup tunnels that terminate at the next-next hop protect both the downstream link and node. This provides protection for link and node failures.

Multiple Backup Tunnels Can Protect the Same Interface

In addition to being required for Node Protection, this enhancement provides the following benefits:

- Redundancy—If one backup tunnel is down, other backup tunnels protect LSPs.
- Increased backup capacity—If the protected interface is a high-capacity link and no single backup path exists with an equal capacity, multiple backup tunnels can protect that one high-capacity link. The LSPs using this link will fail over to different backup tunnels, allowing all of the LSPs to have adequate bandwidth protection during failure (rerouting). If bandwidth protection is not desired, the router spreads LSPs across all available backup tunnels (that is, there is load balancing across backup tunnels). For a more detailed explanation, see "Backup Tunnel Selection Procedure".

Bandwidth Protection

Rerouted LSPs not only have their packets delivered during a failure, but the quality of service can also be maintained.

Scalability

A backup tunnel can protect multiple LSPs. Furthermore, a backup tunnel can protect multiple interfaces. This is called many-to-one (N:1) protection. N:1 protection has significant scalability advantages over one-to-one (1:1) protection, where a separate backup tunnel must be used for each LSP needing protection. N:1 protection is not new with Node Protection; it existed with Link Protection.

Example of 1:1 protection: When 5,000 backup tunnels protect 5,000 LSPs, each router along the backup path must maintain state for an additional 5,000 tunnels.

Example of N:1 protection: When one backup tunnel protects 5,000 LSPs, each router along the backup path maintains one additional tunnel.

RSVP Hello

RSVP Hello allows a router to detect when its neighbor has gone down but its interface to that neighbor is still operational. When Layer 2 link protocols are unable to detect that the neighbor is unreachable, Hellos provide the detection mechanism; this allows the router to switch LSPs onto its backup tunnels and avoid packet loss.

Fast ReRoute Operation

This section illustrates and describes the following:

Fast ReRoute Activation

Two mechanisms cause routers to switch LSPs onto their backup tunnels:

- Interface down notification
- RSVP Hello neighbor down notification

When a router's link or neighboring node fails, the router often detects this failure by an interface down notification. On a GSR Packet Over SONET (POS) interface, this notification is very fast. When a router notices that an interface has gone down, it switches LPSs going out that interface onto their respective backup tunnels (if any).

RSVP Hellos can also be used to trigger Fast ReRoute. If RSVP Hellos are configured on an interface, messages are periodically sent to the neighboring router. If no response is received, Hellos declare that the neighbor is down. This causes any LSPs going out that interface to be switched to their respective backup tunnels.

Backup Tunnels Terminating at Different Destinations

Figure 3 illustrates an interface that has multiple backup tunnels terminating at different destinations and demonstrates why, in many topologies, support for Node Protection requires supporting multiple backup tunnels per protected interface.

Figure 3 Backup Tunnels that Terminate at Different Destinations

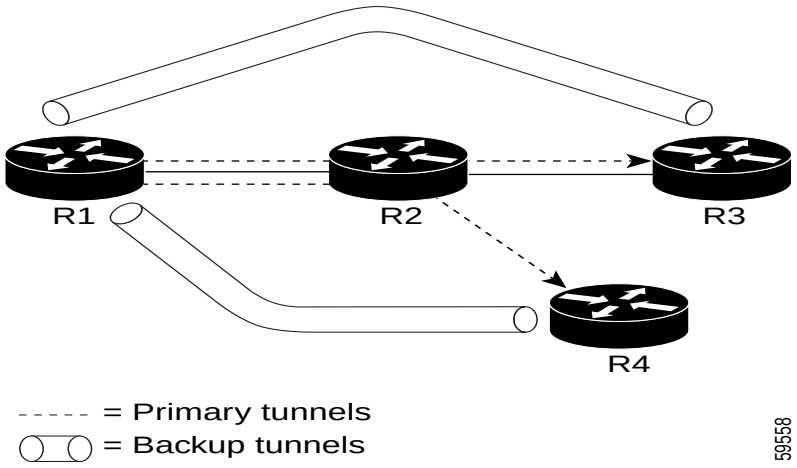

In this illustration, a single interface on R1 requires multiple backup tunnels. LSPs traverse the following routes:

- R1, R2, R3
- R1, R2, R4

To provide protection if node R2 fails, two NNHOP backup tunnels are required: one terminating at R3 and one terminating at R4.

Backup Tunnels Terminating at the Same Destination

Figure 4 shows how backup tunnels terminating at the same location can be used for redundancy and load balancing. Redundancy and load balancing work for both NHOP and NNHOP backup tunnels.

Figure 4 Backup Tunnels that Terminate at the Same Destination

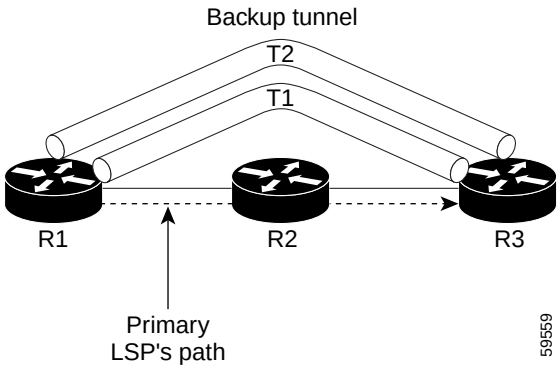

In this illustration, there are three routers: R1, R2, and R3. At R1, there are two NNHOP backup tunnels (T1 and T2) that go from R1 to R3 without traversing R2.

Redundancy—If R2 fails or the link from R1 to R2 fails, either backup tunnel can be used. If one backup tunnel is down, the other can be used. LSPs are assigned to backup tunnels when the LSPs are first established. This is done before a failure.

Load balancing—If neither backup tunnel has enough bandwidth to back up all LSPs, both tunnels can be used. Some LSPs will use one backup tunnel, other LSPs will use the other backup tunnel. The router decides the best way to fit the LSPs onto the backup tunnels.

Backup Tunnel Selection Procedure

When an LSP is signaled, each node along the LSP path that provides FRR protection for the LSP selects a backup tunnel for the LSP to use if either of the following events occurs:

- The link to the next hop fails.
- The next hop fails.

By having the node select the backup tunnel for an LSP before a failure occurs, the LSP can be rerouted onto the backup tunnel quickly if there is a failure.

For an LSP to be mapped to a backup tunnel, all of the following conditions must exist:

- The LSP is protected by FRR; that is, the LSP is configured with the **tunnel mpls traffic-eng fast-reroute** command.
- The backup tunnel is up.
- The backup tunnel is configured to have an IP address, typically a loopback address.
- The backup tunnel is configured to protect this LSP’s outgoing interface; that is, the interface is configured with the **mpls traffic-eng backup-path** command.
- The backup tunnel does not traverse the LSP’s protected interface.
- The backup tunnel terminates at the LSP’s NHOP or NNHOP. If it is an NNHOP tunnel, it does not traverse the LSP’s NHOP.
- The bandwidth protection requirements and constraints, if any, for the LSP and backup tunnel are met. For information about bandwidth protection considerations, see “Bandwidth Protection”.

Bandwidth Protection

A backup tunnel can be configured to protect two types of backup-bandwidth:

- Limited backup-bandwidth—A backup tunnel provides bandwidth protection. The sum of the bandwidth of all LSPs using this backup tunnel cannot exceed the backup tunnel's backup-bandwidth. When assigning LSPs to this type of backup tunnel, sufficient backup-bandwidth must exist.
- Unlimited backup-bandwidth—The backup tunnel does not provide any bandwidth protection (that is, best-effort protection exists). There is no limit to the amount of bandwidth used by the LSPs that are mapped to this backup tunnel. LSPs that allocate zero bandwidth can only use backup tunnels that have unlimited backup-bandwidth.

Load-balancing on Limited-bandwidth Backup Tunnels

There may be more than one backup tunnel that has sufficient backup-bandwidth to protect a given LSP. In this case, the router chooses the one that has the least amount of backup-bandwidth available. This algorithm limits fragmentation, maintaining the largest amount of backup-bandwidth available.

Specifying limited backup bandwidth does not "guarantee" bandwidth protection if there is a link or node failure. For example, the set of NHOP and NNHOP backup tunnels that gets triggered when an interface fails may all share some link on the network topology, and this link may not have sufficient bandwidth to support all LSPs using this set of backup tunnels.

In Figure 5, both backup tunnels traverse the same links and hop. When the link between routers R1 and R4 fails, backup tunnels for primary tunnel 1 and primary tunnel 2 are triggered simultaneously. The two backup tunnels may share a link in the network.

Figure 5 Backup Tunnels Share a Link

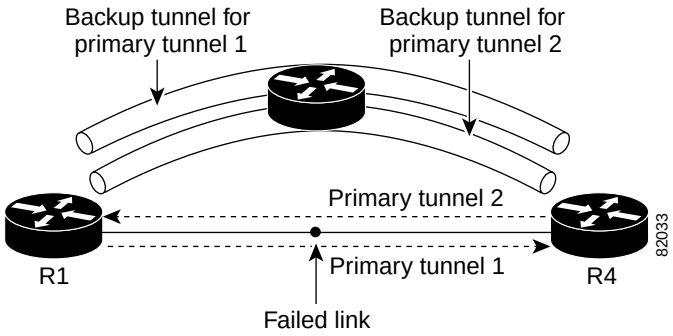

In Figure 6, the backup tunnel for primary tunnel 1 may traverse routers R1-R2-R3-R4, and the backup tunnel for primary tunnel 2 may traverse routers R4-R2-R3-R1. In this case, the link R2-R3 may get overloaded if R1-R4 fails.

Figure 6 Overloaded Link

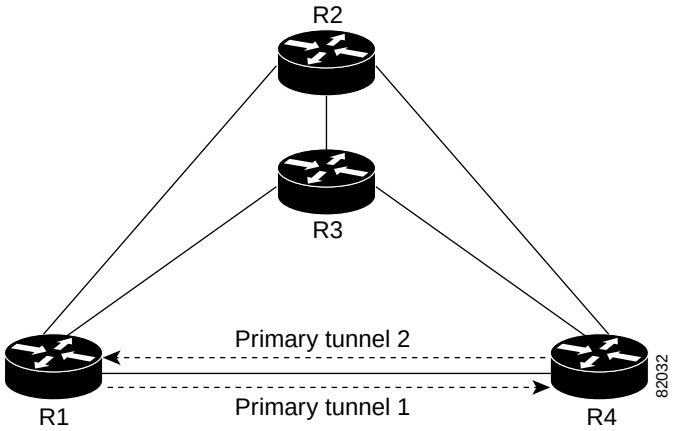

Load-balancing on Unlimited-bandwidth Backup Tunnels

More than one backup tunnel, each having unlimited backup-bandwidth, can protect a given interface. In this case, when choosing a backup tunnel for a given LSP, the router chooses the backup tunnel that has the least amount of backup-bandwidth in use. This algorithm evenly distributes the LSPs across backup tunnels based on LSP's bandwidth. If an LSP is requesting zero bandwidth, the router chooses the backup tunnel that is currently protecting the fewest LSPs.

Pool Type and Backup Tunnels

By default, a backup tunnel provides protection for LSPs that allocate from any pool (that is, global or sub-pool). However, a backup tunnel can be configured to protect only LSPs that use global-pool bandwidth, or only those that use sub-pool bandwidth.

Next-hop Versus Next-next Hop Backup Tunnels

More than one backup tunnel can protect a given LSP, where one backup tunnel terminates at the LSP's NNHOP, and the other terminates at the LSP's NHOP. In this case, the router chooses the backup tunnel that terminates at the NNHOP (that is, Fast ReRoute prefers NNHOP over NHOP backup tunnels).

Table 1 lists the tunnel selection priorities. The first choice is an NNHOP backup tunnel that acquires its bandwidth from a sub-pool or global-pool, and has limited bandwidth. If there is no such backup tunnel, the next choice (2) is a next-next hop backup tunnel that acquires a limited amount of bandwidth from any pool. The preferences go from 1 (best) to 8 (worst), where choice 3 is for an NNHOP backup tunnel with an unlimited amount of sub-pool or global-pool bandwidth.

Table 1 Tunnel Selection Priorities

Preference	Backup Tunnel Destination	Bandwidth Pool	Bandwidth Amount
1 (Best)	Next-next hop	Sub-pool or global-pool	Limited
2	Next-next hop	Any	Limited
3	Next-next hop	Sub-pool or global-pool	Unlimited
4	Next-next hop	Any	Unlimited
5	Next-hop	Sub-pool or global-pool	Limited
6	Next-hop	Any	Limited
7	Next-hop	Sub-pool or global-pool	Unlimited
8 (Worst)	Next-hop	Any	Unlimited

Figure 7 shows an example of the backup tunnel selection procedure based on the designated amount of global-pool and sub-pool bandwidth currently available.

Figure 7 Choosing from Among Multiple Backup Tunnels

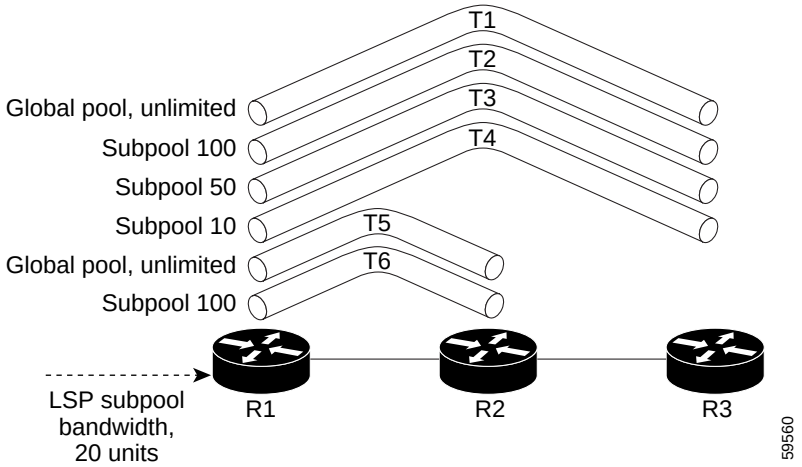

In this example, an LSP requires 20 units (kilobits per second) of sub-pool backup-bandwidth. The best backup tunnel is selected as follows:

1. Backup tunnels T1 through T4 are considered first because they terminate at the NNHOP.
2. Tunnel T4 is eliminated because it only has 10 units of sub-pool backup-bandwidth.
3. Tunnel T1 is eliminated because it protects only LSPs using global-pool bandwidth.
4. Tunnel T3 is chosen over T2 because, although both have sufficient backup-bandwidth, T3 has the least backup-bandwidth available (leaving the most backup-bandwidth available on T2).
5. Tunnels T5 and T6 need not be considered because they terminate at an NHOP, and therefore are less desirable than T3, which terminates at an NNHOP.

Promotion

After a backup tunnel has been chosen for an LSP, conditions may change that will cause us to reevaluate this choice. This reevaluation, if successful, is called promotion. Such conditions may include:

1. A new backup tunnel comes up.
2. The currently chosen backup tunnel for this LSP goes down.
3. A backup tunnel's available backup-bandwidth increases. For example, an LSP protected by the tunnel has been reoptimized by the headend to use another path.

For cases 1 and 2, above, the LSP's backup tunnel is evaluated immediately. Case 3 is addressed by periodically reevaluating LSP-to-backup tunnel mappings. By default, background reevaluation is performed every 5 minutes. This interval is configurable via the **mpls traffic-eng fast-reroute timers** command.

RSVP Hello Operation

RSVP Hello enables RSVP nodes to detect when a neighboring node is not reachable. This provides node-to-node failure detection. When such a failure is detected, it is handled in a similar manner as a link-layer communication failure.

RSVP Hello can be used by FRR when notification of link-layer failures is not available (for example, with Ethernet), or when the failure detection mechanisms provided by the link layer are not sufficient for the timely detection of node failures.

A node running Hello sends a Hello Request to a neighboring node every interval. If the receiving node is running Hello, it responds with Hello Ack. If four intervals pass and the sending node has not received an Ack or it receives a bad message, the sending node declares that the neighbor is down and notifies FRR.

There are two configurable parameters:

- Hello interval, by using the **ip rsvp signalling hello refresh interval** command
- Number of acknowledgment messages that are missed before the sending node declares that the neighbor is down, by using the **ip rsvp signalling hello refresh misses** command

Hello Instance

A Hello instance implements RSVP Hello for a given router interface address and remote IP address. A Hello instance is expensive because of the large number of Hello requests that are sent and the strains they put on the router resources. Therefore, create a Hello instance only when it is necessary and delete it when it is no longer needed.

There are two types of Hello instances:

- Active Hello Instances
- Passive Hello Instances

Active Hello Instances

If a neighbor is unreachable when an LSP is ready to be fast rerouted, an active Hello instance is needed. Create an active Hello instance for each neighbor with at least one LSP in this state.

Active Hello instances periodically send Hello Request messages, and expect Hello Ack messages in response. If the expected Ack message is not received, the active Hello instance declares that the neighbor (remote IP address) is unreachable (lost). LSPs traversing that neighbor may be fast rerouted.

If there is a Hello instance with no LSPs for an unreachable neighbor, do not delete the Hello instance. Convert the active Hello instance to a passive Hello instance because there may be an active instance on the neighboring router that is sending Hello requests to this instance.

Passive Hello Instances

Passive Hello instances respond to Hello Request messages (sending Ack messages), but do not initiate Hello Request messages and do not cause LSPs to be fast rerouted. A router with multiple interfaces can run multiple Hello instances to different neighbors or to the same neighbor.

A passive Hello instance is created when a Hello Request is received from a neighbor with a source IP address/destination IP address pair in the IP header for which a Hello instance does not exist.

Delete passive instances if no Hello messages are received for this instance within 10 minutes.

Hello Commands

RSVP Hello comprises the following commands:

- **RSVP Hello Configuration Commands**
- **RSVP Hello Statistics Commands**
- **RSVP Hello Show Commands**
- **RSVP Hello Debug Commands**

RSVP Hello Configuration Commands

- **ip rsvp signalling hello (configuration)**—Enables Hello globally on the router.
- **ip rsvp signalling hello (interface)**—Enables Hello on an interface where you need Fast ReRoute protection.
- **ip rsvp signalling hello dscp**—Sets the DSCP value that is in the IP header of the Hello message sent out from an interface.
- **ip rsvp signalling hello refresh interval**—Configures the Hello request interval.
- **ip rsvp signalling hello refresh misses**—Specifies how many Hello acknowledgments a node can miss in a row before the node considers that communication with its neighbor is down.
- **ip rsvp signalling hello statistics**—Enables Hello statistics on the router.

RSVP Hello Statistics Commands

- **clear ip rsvp hello instance counters**—Clears (refreshes) the values for Hello instance counters.
- **clear ip rsvp hello instance statistics**—Clears Hello statistics for an instance.
- **clear ip rsvp hello statistics**—Globally clears Hello statistics.

RSVP Hello Show Commands

- **show ip rsvp hello**—Shows if Hello is enabled globally on the router and if Hello statistics are enabled.
- **show ip rsvp hello instance detail**—Shows detailed information about a Hello instance.
- **show ip rsvp hello instance summary**—Shows summary information about a Hello instance.
- **show ip rsvp hello statistics**—Shows how long Hello packets have been in the Hello input queue.
- **show ip rsvp interface detail**—Shows the interface configuration for Hello.

RSVP Hello Debug Commands

- **debug ip rsvp hello**—Verifies that a Hello instance has been created, a Hello instance has been deleted, and when communication with a neighbor has been lost.

Restrictions

- Interfaces must use MPLS Global Label Allocation.
- Backup tunnel headend and tailend routers must implement Fast ReRoute as described in this document and in draft-pan-rsvp-fastreroute-00.txt.
- Backup tunnels are not protected. If an LSP is actively using a backup tunnel and the backup tunnel fails, the LSP is torn down.
- LSPs that are actively using backup tunnels are not considered for promotion. So, if an LSP is actively using a backup tunnel and a better backup tunnel becomes available, the active LSP is not switched to the better backup tunnel.

Related Features and Technologies

- Intermediate System-to-Intermediate-System (IS-IS)
- MPLS
- MPLS Traffic Engineering Exclude Node/Link
- Open Shortest Path First (OSPF)
- RSVP

Related Documents

For IS-IS:

- *Cisco IOS IP Configuration Guide*, Release 12.2
- *Cisco IOS IP Command Reference, Volume 2 of 3: Routing Protocols*, Release 12.2

For Link Protection:

- *Cisco IOS Switching Services Command Reference*, Release 12.2
- *Cisco IOS Switching Services Configuration Guide*, Release 12.2

For MPLS Traffic Engineering:

- *Cisco IOS Switching Services Command Reference*, Release 12.2
- *Cisco IOS Switching Services Configuration Guide*, Release 12.2
- *Cisco IOS IP Command Reference, Volume 2 of 3: Routing Protocols*, Release 12.2
- *Diff-Serv-aware MPLS Traffic Engineering*, Release 12.2(4)T
- *MPLS Traffic Engineering (TE)—Interarea Tunnels*, Release 12.0(22)S
- *MPLS Traffic Engineering (TE)—IP Explicit Address Exclusion*, Release 12.0(22)S

For OSPF:

- *Cisco IOS IP Command Reference, Volume 2 of 3: Routing Protocols*, Release 12.2
- *Cisco IOS IP Configuration Guide*, Release 12.2

For RSVP:

- *Cisco IOS Quality of Service Solutions Command Reference*, Release 12.2
- *Cisco IOS Quality of Service Solutions Configuration Guide*, Release 12.2

Supported Platforms and Interfaces

Supported Platforms

- Cisco 7200 series (including the Cisco 7202, Cisco 7204, Cisco 7204 VXR, Cisco 7206, and Cisco 7206 VXR)
- Cisco 7500 series (including the Cisco 7505, Cisco 7507, Cisco 7513, and Cisco 7576)
- Cisco GSR 12000 series (including the Cisco 12008, Cisco 12012, Cisco 12016, Cisco 12404, Cisco 12406, Cisco 12410, and Cisco 12416)

Determining Platform Support Through Cisco Feature Navigator

Cisco IOS software is packaged in feature sets that support specific platforms. To get updated information regarding platform support for this feature, access Cisco Feature Navigator. Cisco Feature Navigator dynamically updates the list of supported platforms as new platform support is added for the feature.

Cisco Feature Navigator is a web-based tool that enables you to determine which Cisco IOS software images support a specific set of features and which features are supported in a specific Cisco IOS image. You can search by feature or release. Under the release section, you can compare releases side by side to display both the features unique to each software release and the features in common.

To access Cisco Feature Navigator, you must have an account on Cisco.com. If you have forgotten or lost your account information, send a blank e-mail to cco-locksmith@cisco.com. An automatic check will verify that your e-mail address is registered with Cisco.com. If the check is successful, account details with a new random password will be e-mailed to you. Qualified users can establish an account on Cisco.com by following the directions at http://www.cisco.com/register.

Cisco Feature Navigator is updated regularly when major Cisco IOS software releases and technology releases occur. For the most current information, go to the Cisco Feature Navigator home page at the following URL:

http://www.cisco.com/go/fn

Availability of Cisco IOS Software Images

Platform support for particular Cisco IOS software releases is dependent on the availability of the software images for those platforms. Software images for some platforms may be deferred, delayed, or changed without prior notice. For updated information about platform support and availability of software images for each Cisco IOS software release, refer to the online release notes or, if supported, Cisco Feature Navigator.

Supported Interfaces

- Fast Ethernet
- Gigabit Ethernet
- Packet over SONET (POS)

Supported Standards, MIBs, RFCs, and Drafts

Standards

- draft-ietf-mpls-rsvp-lsp-tunnel-09.txt
- draft-pan-rsvp-fastreroute-00.txt

MIBs

No new or modified MIBs are supported by this feature.

To obtain lists of supported MIBs by platform and Cisco IOS release, and to download MIB modules, go to the Cisco MIB website on Cisco.com at the following URL:

http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs and Drafts

This feature complies with draft-swallow-rsvp-bypass-label-03.txt.

Prerequisites

Your network must support the following Cisco IOS features in order to support features described in this document:

- IP Cisco Express Forwarding (CEF)
- MPLS

Your network must support at least one of the following protocols:

- IS-IS
- OSPF

Configuration Tasks

This section assumes that you want to add Fast ReRoute protection to a network in which MPLS TE LSPs are configured.

Before performing the configuration tasks, it is assumed that you have done the following tasks but you do not have to already have configured MPLS TE tunnels:

- Enabled MPLS TE on all relevant routers and interfaces
- Configured MPLS TE tunnels

To review how to configure MPLS TE tunnels, see the *Cisco IOS Switching Services Configuration Guide*, Release 12.2.

The following sections describe how to use FRR to protect LSPs in your network from link or node failures. Each task is identified as either required or optional.

- Enabling Fast ReRoute on LSPs (required)
- Creating a Backup Tunnel to the Next Hop or to the Next-Next Hop (required)
- Assigning Backup Tunnels to a Protected Interface (required)
- Associating Backup-Bandwidth and Pool Type with a Backup Tunnel (optional)
- Configuring an Interface for Fast Link and Node Failure Detection (optional)

Note You can perform the configuration tasks in any order.

Note An NNHOP backup tunnel must *not* go via the NHOP.

Enabling Fast ReRoute on LSPs

LSPs can use backup tunnels only if they have been configured as fast reroutable. To do this, enter the following commands, beginning in global configuration mode, at the headend of each LSP:

	Command	Purpose
Step 1	`Router(config)# interface tunnel number`	Enters interface configuration mode for the specified tunnel.
Step 2	`Router(config-if)# tunnel mpls traffic-eng fast-reroute`	Enables the tunnel to use a backup tunnel if there is a link or node failure.

Creating a Backup Tunnel to the Next Hop or to the Next-Next Hop

To create a backup tunnel to the next hop or to the next-next hop, enter the following commands on the node that will be the headend of the backup tunnel (that is, the node whose downstream link or node may fail). The node on which you enter these commands must be a supported platform. See "Supported Platforms and Interfaces".

Creating a backup tunnel is basically no different from creating any other tunnel. None of the commands below is new.

Note When using the **exclude-address** command to specify the path for a backup tunnel, the exclude-address must exclude an interface address to avoid a link (for creating an NHOP backup tunnel), or a router-ID address to avoid a node (for creating an NNHOP backup tunnel).

	Command	Purpose
Step 1	`Router(config)# interface tunnel number`	Creates a new tunnel interface and enters interface configuration mode.
Step 2	`Router(config-if)# ip unnumbered loopback0`	Gives the tunnel interface an IP address which is the same as that of interface Loopback0. **Note** This command is not effective until Lookback0 has been configured with an IP address.
Step 3	`Router(config-if)# tunnel destination A.B.C.D.`	Specifies the IP address of the device where the tunnel will terminate. This address should be the router-id of the device which is the NHOP or NNHOP of LSPs to be protected.
Step 4	`Router(config-if)# tunnel mode mpls traffic-eng`	Sets encapsulation mode of the tunnel to MPLS Traffic Engineering.
Step 5	`Router(config-if)# tunnel mpls traffic-eng path-option number explicit name path-name`	The backup tunnel's path will be computed by using the topology database to find the shortest path to the destination (*A.B.C.D*) that meets the constraints. The only constraint we have configured (via the explicit-path, below) is that the path should avoid the protected link or node.
Step 6	`Router(config)# ip explicit-path name name`	Enters the subcommand mode for IP explicit paths to create the named path.
Step 7	`Router(cfg-ip-expl-path)# exclude-address address`	For Link Protection, specify the IP address of the link to be protected. For Node Protection, specify the router-ID of the node to be protected. **Note** Backup tunnel paths can be dynamic or explicit and they do not have to use exclude-address. Because backup tunnels must avoid the protected link or node, it is convenient to use an exclude-address.

Assigning Backup Tunnels to a Protected Interface

To assign one or more backup tunnels to a protected interface, enter the following commands on the node that will be the headend of the backup tunnel (that is, the node whose downstream link or node may fail). The node on which you enter these commands must be a supported platform. See "Supported Platforms and Interfaces".

Note You must configure the interface to have an IP address and to enable the MPLS Traffic Engineering tunnel feature.

	Command	Purpose
Step 1	`Router(config)# interface type slot/port`	Moves configuration to the physical interface level, directing subsequent configuration commands to the specific physical interface identified by the *type*. The *slot* and *port* identify the slot and port being configured. The interface must be a supported interface. See "Supported Platforms and Interfaces".
Step 2	`Router(config-if)# mpls traffic-eng backup-path tunnel tunnel-id`	Allows LSPs going out this interface to use this backup tunnel if there is a link or node failure. **Note** You can enter this command multiple times to associate multiple backup tunnels with the same protected interface.

Associating Backup-Bandwidth and Pool Type with a Backup Tunnel

To associate backup-bandwidth with a backup tunnel and designate the type of LSP that can use a backup tunnel, enter the following command:

Command	Purpose
`Router(config-if)# tunnel mpls traffic-eng backup-bw {bandwidth | [sub-pool {bandwidth | Unlimited}] [global-pool {bandwidth | Unlimited}] } [any {bandwidth | Unlimited}] }`	Associates bandwidth with a backup tunnel and designates whether LSPs that allocate bandwidth from the specified pool can use the tunnel.

Configuring an Interface for Fast Link and Node Failure Detection

To configure pos ais-shut, enter the following commands:

```
interface pos0/0
pos ais-shut
```

To configure pos report lrdi on OS interfaces, enter the following commands:

```
interface pos0/0
 pos  report lrdi
```

Verifying That Fast ReRoute Is In Place

To ensure that Fast ReRoute can function, do the following:

- Determine if Fast ReRoute has been configured correctly.
- Verify that certain conditions exist so that backup tunnels can be operational.
- Enter the **show mpls traffic-eng fast-reroute database** command.
- Enter the **show mpls traffic tunnel backup** command.

- Enter the **show ip rsvp sender** command, with the **detail** keyword specified.
- Enter the **show ip rsvp reservation** command, with the **detail** keyword specified.

Fast ReRoute Configuration

To determine if Fast ReRoute has been configured correctly, do the following:

- Verify that backup tunnels are up. To do so, enter the **show mpls traffic-eng tunnels brief** command.
- Verify that LSPs are protected by the appropriate backup tunnels. To do so, enter the **show ip rsvp sender** command with the **detail** keyword.

Conditions that Must Exist for Backup Tunnels to be Operational

If you created LSPs and performed the required configuration tasks but do not have operational backup tunnels (that is, the backup tunnels are not up or the LSPs are not associated with those backup tunnels), verify that all the following conditions exist:

- **LSP is reroutable**—At the headend of the LSP, enter the **show run int tunnel** *tunnel-number* command. The output should include the **tunnel mpls traffic-eng fast-reroute** command. If it does not, enter this command for the tunnel.

On the router where the backup tunnels originate, enter the **show mpls traffic-eng tunnels backup** command. The command output will allow you to verify the following:

- **Backup tunnel exists**—Verify that there is a backup tunnel that terminates at this LSP's NHOP or NNHOP. Look for the LSP's NHOP or NNHOP in the Dest field.
- **Backup tunnel is up**—To verify that the backup tunnel is up, look for "Up" in the State field.
- **Backup tunnel is associated with LSP's I/F**—Verify that the interface for the LSP is allowed to use this backup tunnel. Look for the LSP's output interface in the "protects" field list.
- **Backup tunnel has sufficient bandwidth**—If you restricted the amount of bandwidth a backup tunnel can hold, verify that the backup tunnel has sufficient bandwidth to hold the LSPs that would use this backup tunnel if there is a failure. The bandwidth of an LSP is defined by the line **tunnel mpls traffic-eng bandwidth** at the headend of the LSP. To determine the available bandwidth on a backup tunnel, look at the "cfg" and "inuse" fields. If there is insufficient backup-bandwidth to accommodate the LSPs that would use this backup tunnel in the event of a failure, create an additional backup tunnel or increase the backup-bandwidth of the existing tunnel by using the **tunnel mpls traffic-eng backup-bw** command.

Note To determine how much bandwidth is sufficient, offline capacity planning may be required.

- **Backup tunnel has appropriate bandwidth type**—If you restricted the type of LSPs (sub-pool or global-pool) that can use this backup tunnel, verify that the LSP is the appropriate type for the backup tunnel. The type of the LSP is defined by the line **tunnel mpls traffic-eng bandwidth** at the headend of this LSP. If this line contains the word "sub-pool", then it uses sub-pool bandwidth; otherwise, it uses global-pool bandwidth. Verify that the type matches the type the backup tunnel can hold by looking in the output of the above command.

If none of the above actions works, enable debug by entering the **debug ip rsvp fast-reroute** command and the **debug mpls traffic-eng fast-reroute** command on the router that is the headend of the backup tunnel. Then do the following:

1. Enter the **shutdown** command for the primary tunnel.
2. Enter the **no shutdown** command for the primary tunnel.
3. View the debug output.

show mpls traffic-eng fast-reroute database command

Enter the **clear ip rsvp hello instance counters** command to verify the following:

- MPLS Traffic Engineering Fast ReRoute Node Protection has been enabled.
- A certain type of LSP can use a backup tunnel.

The following command output displays the LSPs that are protected:

```
Router# show mpls traffic-eng fast-reroute database

Tunnel head fast reroute information:
Prefix              Tunnel     In-label Out intf/label   FRR intf/label   Status
10.0.0.11/32        Tu1        Tun hd   pos0/1:Untagged   Tu2000:16        ready
LSP midpoint frr information:
LSP identifier                 In-label Out intf/label   FRR intf/label   Status
10.0.0.12 1 [459]              16       pos0/1:17         Tu2000:19        ready
```

show mpls traffic tunnel backup command

The following **show ip rsvp sender** command output verifies that protection has been enabled.

```
Router# show mpls traffic-eng tunnel backup

Tunnel1000           Dest: 12.0.0.10       State: Up
glb-pool cfg 100 inuse 0 num_lsps 0
     protects: POS0/0
     protects: POS0/1
```

The command shows the following information for a given backup tunnel:

- Tunnel ID
- Tunnel destination
- Tunnel state—Up designates the status of the backup tunnel
- Backup-bandwidth configured for each pool this tunnel protects
- Backup-bandwidth in use for each pool
- Number of LSPs currently using this backup tunnel
- Protected interfaces that are using the backup tunnel

show ip rsvp sender command

Following is sample output from the **show ip rsvp sender detail** command when the command is entered at the Point of Local Repair (PLR) before a failure. For a detailed explanation of the output, see the **show ip rsvp sender** command.

```
Router# show ip rsvp sender detail

PATH:
   Tun Dest: 24.1.1.1 Tun ID: 1 Ext Tun ID: 23.1.1.1
   Tun Sender: 23.1.1.1, LSP ID: 126
   Path refreshes arriving on POS1/0 from PHOP 11.1.1.1
   Path refreshes being sent to NHOP 12.1.1.2 on POS1/1
   Session Attr::
     Setup Prio: 0, Holding Prio: 0
     Flags: Local Prot desired, Label Recording, SE Style
     Session Name:tagsw4500-23_t1
   ERO:
     12.1.1.2 (Strict IPv4 Prefix, 8 bytes, /32)
     14.1.1.1 (Strict IPv4 Prefix, 8 bytes, /32)
     14.1.1.2 (Strict IPv4 Prefix, 8 bytes, /32)
     24.1.1.1 (Strict IPv4 Prefix, 8 bytes, /32)
```

```
  Traffic params - Rate: 0G bits/sec, Max. burst: 1K bytes
  Fast-Reroute Backup info:
    Inbound  FRR: Not active
    Outbound FRR: Ready -- backup tunnel selected
      Backup Tunnel: Tu2        (label 0)
      Bkup Sender Template:
        Tun Sender: 15.1.1.1, LSP ID: 126
      Bkup FilerSpec:
        Tun Sender: 15.1.1.1, LSP ID 126
```

show ip rsvp reservation command

Following is sample output from the **show ip rsvp reservation** command entered at the head-end of a primary LSP. Entering the command at the head-end of the primary LSP shows, among other things, the status of FRR (that is, local protection) at each hop this LSP traverses. The per-hop information is collected in the Record Route Object (RRO) that travels with the Resv message from the tail to the head.

```
Router# sho ip rsvp res det
Reservation:
  Tun Dest:   24.1.1.1  Tun ID: 1  Ext Tun ID: 23.1.1.1
  Tun Sender: 23.1.1.1  LSP ID: 104
  Next Hop: 11.1.1.2 on POS1/0
  Label: 18 (outgoing)
  Reservation Style is Shared-Explicit, QoS Service is Controlled-Load
  Average Bitrate is 0 bits/sec, Maximum Burst is 1K bytes
  Min Policed Unit: 0 bytes, Max Pkt Size: 0 bytes
  RRO:
    12.1.1.1/32, Flags:0x1 (Local Prot Avail/to NHOP)
      Label subobject: Flags 0x1, C-Type 1, Label 18
    14.1.1.1/32, Flags:0x0 (Local Prot Avail/In Use/Has BW/to NHOP)
      Label subobject: Flags 0x1, C-Type 1, Label 16
    14.1.1.2/32, Flags:0x0 (No Local Protection)
      Label subobject: Flags 0x1, C-Type 1, Label 0
  Resv ID handle: CD000404.
  Policy: Accepted. Policy source(s): MPLS/TE
```

Notice the following about the primary LSP:

- It has protection that uses a NHOP backup tunnel at its first hop.
- It has protection and is actively using an NHOP backup tunnel at its second hop.
- It has no local protection at its third hop.

The RRO display shows the following information, for each hop:

- Whether local protection is available (that is, whether the LSP has selected a backup tunnel)
- Whether local protection is in use (that is, whether the LSP is currently using its selected backup tunnel)
- Whether the selected backup tunnel is an NHOP or NNHOP backup tunnel
- Whether the backup tunnel used at this hop provides bandwidth protection

Troubleshooting Tips

This section describes the following:

LSPs Do Not Become Active; They Remain Ready

At a Point of Local Repair (PLR), LSPs transition from Ready to Active if one of the following events occurs:

- Primary interface goes down—If the primary interface (LSP's outbound interface) goes down and the LSP is ready to use a backup tunnel, the LSP will transition to the active state causing its data to flow over the backup tunnel. On some platforms and interface types (for example, GSR POS interfaces), fast interface-down logic has been added to detect this event very quickly. On other platforms where this logic does not exist, detection time is slower. On such platforms, it may be desirable to enable RSVP Hello (see the next bulleted item, "Hellos detect next hop is down").
- Hellos detect next hop is down—If Hellos are enabled on the primary interface (LSP's outbound interface), and the LSP's next hop is no longer reachable, the next hop is declared down. This event will cause the LSP to begin actively using its backup tunnel. Notice that a next hop will be declared down even if the primary interface does not go down. For example, if the next hop stops responding due to a reboot or software/hardware problem, Hellos will trigger the LSPs using this next hop to switch to their backup tunnels. Hellos can also help trigger Fast ReRoute on interfaces such as Gigabit Ethernet where the interface remains up but is unusable (due to lack of link-layer liveness detection mechanisms).

Primary Tunnel Does Not Select Backup Tunnel That is Up

If a backup tunnel is up, but it is not selected as a backup tunnel by the primary tunnel (LSP), enter the following commands for the backup tunnel:

- **shutdown**
- **no shutdown**

Note If you change the status of a backup tunnel, the backup tunnel selection algorithm is rerun for the backup tunnel. LSPs that have currently selected (that is, are ready to use) that backup tunnel will be disassociated from it, and then reassociated with that backup tunnel or another backup tunnel. This is generally harmless and usually results in mapping the same LSPs to that backup tunnel. However, if any LSPs are actively using that backup tunnel, shutting down the backup tunnel will tear down those LSPs.

Enhanced RSVP Commands

The following RSVP commands have been enhanced to display information that can be helpful when examining Fast ReRoute state or when troubleshooting Fast ReRoute:

- **show ip rsvp request**—Displays upstream reservation state (that is, information related to the Resv messages that this node will send upstream).
- **show ip rsvp reservation**—Displays information about Resv messages received.
- **show ip rsvp sender**—Displays information about Path messages being received.

These commands show control plane state; they do not show data state. That is, they show information about RSVP messages (Path and Resv) used to signal LSPs. For information about the data packets being forwarded along LSPs, use the **show mpls forwarding** command.

RSVP Hello

The RSVP Hello feature enables RSVP nodes to detect when a neighboring node is not reachable. Use this feature when notification of link-layer failures is not available and unnumbered links are not used, or when the failure detection mechanisms provided by the link layer are not sufficient for timely node failure detection. Hello must be configured both globally on the router and on the specific interface to be operational.

The RSVP Hello commands are described briefly in the "RSVP Hello Operation" section and in more detail in the "Command Reference" section.

Hello Instances Have Not Been Created

If Hello instances have not been created, do the following:

- Determine if RSVP Hello has been enabled globally on the router. See the "ip rsvp signalling hello (configuration)" command.
- Determine if RSVP Hello has been enabled on an interface that the LSPs traverse. See the "ip rsvp signalling hello (interface)" command.
- Verify that at least one LSP has a backup tunnel by viewing the output of the "show ip rsvp sender" command. A value of "Ready" indicates that a backup tunnel has been selected.

"No entry at index (error may self-correct, RRO may not yet have propagated from downstream node of interest)" Error Message is Printed at the Point of Local Repair

Fast ReRoute relies on a Record Route Object (RRO) in Resv messages arriving from downstream. Routers receiving Path messages with the SESSION_ATTRIBUTE bit indicating that the LSP is fast-reroutable should include an RRO in the corresponding Resv messages.

If an LSP is configured for Fast ReRoute, but the Resv arriving from a downstream router contains an incomplete RRO, the "No entry at index (error may self-correct, RRO may not yet have propagated from downstream node of interest)" message is printed. An incomplete RRO is one in which the NHOP or the NNHOP did not include an entry in the RRO.

This error typically means that backup tunnels to the NHOP or the NNHOP cannot be selected for this LSP because there is insufficient information about the NHOP or NNHOP due to the lack of an RRO entry.

Occasionally there are valid circumstances in which this situation occurs temporarily and the problem is self-corrected. If subsequent Resv messages arrive with a complete RRO, ignore the error message.

To determine whether the error has been corrected, view the RRO in Resv messages by entering the **clear ip rsvp hello instance counters** command. Use an output filter keyword to view only the LSP of interest.

"Couldn't get rsbs (error may self-correct when Resv arrives)" Error Message is Printed at the Point of Local Repair

The PLR cannot select a backup tunnel for an LSP until a Resv message has arrived from downstream.

When this error occurs, it typically means that something is truly wrong. For example, no reservation exists for this LSP. You can troubleshoot this problem by using the **debug ip rsvp reservation** command to enable debug.

Occasionally there are valid circumstances in which this error message occurs and there is no need for concern. One such circumstance is when an LSP experiences a change before any Resv message has arrived from downstream. Changes can cause a PLR to try to select a backup tunnel for an LSP, and the selection will fail (causing this error message) if no Resv message has arrived for this LSP.

Configuration Examples

This section provides the following configuration examples:

- Enabling Fast ReRoute for all Tunnels
- Creating an NHOP Backup Tunnel
- Creating an NNHOP Backup Tunnel
- Assigning Backup Tunnels to a Protected Interface
- Associating Backup-Bandwidth and Pool Type with Backup Tunnels
- Configuring RSVP Hello and POS Signals

The examples relate to the illustration shown in Figure 8.

Figure 8 Backup Tunnels

Enabling Fast ReRoute for all Tunnels

On router R1, enter interface configuration mode for each tunnel to be protected (Tunnel 1000 and Tunnel 2000). Enable these tunnels to use a backup tunnel in case of a link or node failure along their paths. Tunnel 1000 will use 10 units of bandwidth from the sub-pool; Tunnel 2000 will use 5 units of bandwidth from the global-pool.

```
interface Tunnel1000
 tunnel mpls traffic-eng fast-reroute
 tunnel mpls traffic-eng bandwidth sub-pool 10

interface Tunnel2000
 tunnel mpls traffic-eng fast-reroute
     tunnel mpls traffic-eng bandwidth 5
```

Creating an NHOP Backup Tunnel

On router R2, create an NHOP backup tunnel to R3. This backup tunnel should avoid using the link 12.1.1.2.

```
Router(config)# ip explicit-path name avoid-protected-link
Router(cfg-ip-expl-path)# exclude-address 12.1.1.2
Explicit Path name avoid-protected-link:
____1: exclude-address 12.1.1.2
Router(cfg-ip_expl-path)# end

Router(config)# interface Tunnel1
Router(config-if)# ip unnumbered loopback0
Router(config-if)# tunnel destination 3.3.3.3
Router(config-if)# tunnel mode mpls traffic-eng0
Router(config-if)# tunnel mpls traffic-eng path-option explicit avoid-protected-link
```

Creating an NNHOP Backup Tunnel

On router R2, create an NNHOP backup tunnel to R4. This backup tunnel should avoid R3.

```
Router(config)# ip explicit-path name avoid-protected-node
Router(cfg-ip-expl-path)# exclude-address 3.3.3.3
Explicit Path name avoid-protected-node:
____1: exclude-address 3.3.3.3
Router(cfg-ip_expl-path)# end

Router(config)# interface Tunnel2
Router(config-if)# ip unnumbered loopback0
Router(config-if)# tunnel destination 4.4.4.4
Router(config-if)# tunnel mode mpls traffic-eng0
Router(config-if)# tunnel mpls traffic-eng path-option explicit avoid-protected-node
```

Assigning Backup Tunnels to a Protected Interface

On router R2, associate both backup tunnels with interface POS5/0.

```
Router(config)# interface POS5/0
Router(config-if)# mpls traffic-eng backup-path tunnel1
Router(config-if)# mpls traffic-eng backup-path tunnel2
```

Associating Backup-Bandwidth and Pool Type with Backup Tunnels

Backup tunnel 1 is to be used only by LSPs that take their bandwidth from the global pool. It does not provide bandwidth protection. Backup tunnel 2 is to be used only by LSPs that take their bandwidth from the sub-pool. Backup tunnel 2 provides bandwidth protection for up to 1000 units.

```
Router(config)# interface Tunnel1
Router(config-if)# tunnel mpls traffic-eng backup-bw global-pool Unlimited

Router(config)# interface Tunnel2
Router(config-if)# tunnel mpls traffic-eng backup-bw sub-pool 1000
```

Configuring RSVP Hello and POS Signals

Hello must be configured both globally on the router and on the specific interface on which you need Fast ReRoute protection. To configure Hello, use the following configuration commands:

- **ip rsvp signalling hello (configuration)**—Enables Hello globally on the router.
- **ip rsvp signalling hello (interface)**—Enables Hello on an interface where you need Fast ReRoute protection.

The following configuration commands are optional:

- **ip rsvp signalling hello dscp**—Sets the DSCP value that is in the IP header of the Hello message.
- **ip rsvp signalling hello refresh misses**—Specifies how many acknowledgments a node can miss in a row before the node considers that communication with its neighbor is down.
- **ip rsvp signalling hello refresh interval**—Configures the Hello request interval.
- **ip rsvp signalling hello statistics**—Enables Hello statistics on the router.

For configuration examples, see the Hello command descriptions in the "Command Reference" section.

To configure POS signalling for detecting Fast ReRoute failures, enter **pos report all** or enter the following commands to request individual reports:

```
pos ais-shut
pos report rdool
pos report lais
pos report lrdi
pos report pais
pos report prdi
pos report sd-ber
```

Command Reference

This section documents new or modified commands. All other commands used with this feature are documented in the Cisco IOS Release 12.2 command reference publications.

New Commands—Backup Tunnels

- **mpls traffic-eng fast-reroute timers**
- **show mpls traffic tunnel backup**
- **tunnel mpls traffic-eng backup-bw**

Modified Commands—Backup Tunnels

- **mpls traffic-eng backup-path**
- **show ip rsvp request**
- **show ip rsvp reservation**
- **show ip rsvp sender**
- **show mpls traffic-eng fast-reroute database**
- **show mpls traffic-eng tunnels**
- **show mpls traffic-eng tunnels summary**

New Commands—Hello

- **clear ip rsvp hello instance counters**
- **clear ip rsvp hello instance statistics**
- **clear ip rsvp hello statistics**
- **debug ip rsvp hello**
- **ip rsvp signalling hello (configuration)**
- **ip rsvp signalling hello (interface)**
- **ip rsvp signalling hello dscp**
- **ip rsvp signalling hello refresh interval**
- **ip rsvp signalling hello refresh misses**
- **ip rsvp signalling hello statistics**
- **show ip rsvp hello**
- **show ip rsvp hello instance detail**
- **show ip rsvp hello instance summary**
- **show ip rsvp hello statistics**
- **show ip rsvp interface detail**

Note The **show ip rsvp** command examples refer to Figure 9.

Figure 9 RSVP Configuration Example

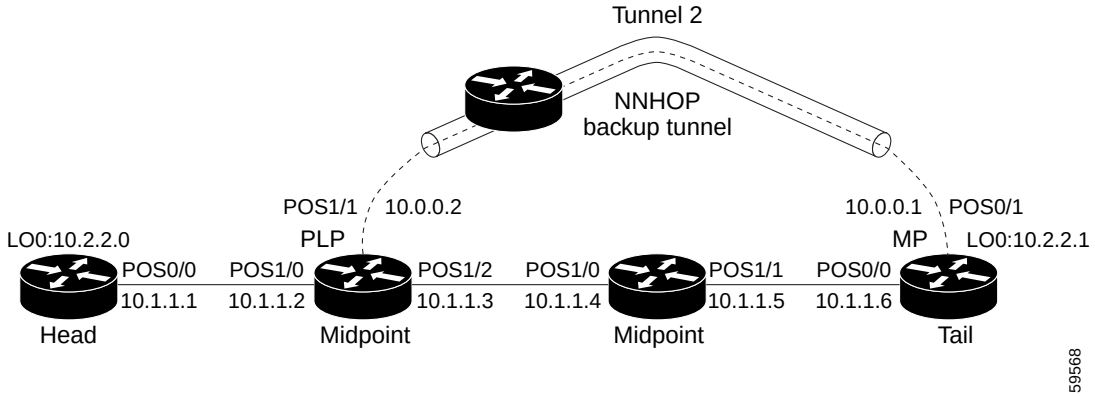

—— = Primary tunnel before failure

- - - - = Section of primary tunnel after failure

clear ip rsvp hello instance counters

To clear (refresh) the values for Hello instance counters, use the **clear ip rsvp hello instance counters** command in global EXEC mode.

clear ip rsvp hello instance counters

Syntax Description

This command has no arguments or keywords.

Defaults

This command has no default behavior or values.

Command Modes

Global EXEC

Command History

Release	Modification
12.0(22)S	This command was introduced.

Examples

Following is sample output from the **show ip rsvp hello instance detail** command and then the **clear ip rsvp hello instance counters** command. Notice that the "Statistics" fields have been cleared to zero.

```
Router# show ip rsvp hello instance detail

Neighbor 11.0.0.2  Source  11.0.0.1
    State: UP      (for 2d18h)
    Type: PASSIVE  (responding to requests)
    I/F: Et1/1
    LSPs protecting: 0
    Refresh Interval (msec) (used when ACTIVE)
      Configured: 100
      Statistics: (from 2398195 samples)
        Min:      100
        Max:      132
        Average:  100
        Waverage: 100 (Weight = 0.8)
        Current:  100
    Src_instance 0xA9F07C13, Dst_instance 0x9BBAA407
    Counters:
      Communication with neighbor lost:
        Num times: 0
        Reasons:
          Missed acks:             0
          Bad Src_Inst received:   0
          Bad Dst_Inst received:   0
          I/F went down:           0
          Neighbor disabled Hello: 0
      Msgs Received:   2398194
           Sent:       2398195
           Suppressed: 0

Router# clear ip rsvp hello instance counters
```

```
Neighbor 11.0.0.2  Source  11.0.0.1
    State: UP       (for 2d18h)
    Type: PASSIVE  (responding to requests)
    I/F: Et1/1
    LSPs protecting: 0
    Refresh Interval (msec) (used when ACTIVE)
      Configured: 100
      Statistics:
        Min:         0
        Max:         0
        Average:     0
        Waverage:    0
        Current:     0
    Src_instance 0xA9F07C13, Dst_instance 0x9BBAA407
    Counters:
      Communication with neighbor lost:
        Num times: 0
        Reasons:
          Missed acks:              0
          Bad Src_Inst received:    0
          Bad Dst_Inst received:    0
          I/F went down:            0
          Neighbor disabled Hello: 0
      Msgs Received:   2398194
           Sent:       2398195
           Suppressed: 0
```

Related Commands

Command	Description
ip rsvp signalling hello (configuration)	Enables Hello globally on the router.
ip rsvp signalling hello (interface)	Enables Hello on an interface where you need Fast ReRoute protection.
ip rsvp signalling hello statistics	Enables Hello statistics on the router.
show ip rsvp hello statistics	Shows how long Hello packets have been in the Hello input queue.

clear ip rsvp hello instance statistics

To clear Hello statistics for an instance, use the **clear ip rsvp hello instance statistics** command in global EXEC mode.

clear ip rsvp hello instance statistics

Syntax Description

This command has no arguments or keywords.

Defaults

This command has no default behavior or values.

Command Modes

Global EXEC

Command History

Release	Modification
12.0(22)S	This command was introduced.

Examples

This example shows output from the **show ip rsvp hello statistics** command and the values in those fields after you enter the **clear ip rsvp hello instance statistics** command.

```
Router# show ip rsvp hello statistics

  Status: Enabled
  Packet arrival queue:
    Wait times (msec)
      Current:0
      Average:0
      Weighted Average:0 (weight = 0.8)
      Max:4
    Current length: 0 (max:500)
  Number of samples taken: 2398525


Router# clear ip rsvp hello instance statistics

Status: Enabled
  Packet arrival queue:
    Wait times (msec)
      Current:0
      Average:0
      Weighted Average:0 (weight = 0.8)
      Max:0
    Current length: 0 (max:500)
  Number of samples taken: 0
```

Related Commands

Command	Description
ip rsvp signalling hello (configuration)	Enables Hello globally on the router.
ip rsvp signalling hello (interface)	Enables Hello on an interface where you need Fast ReRoute protection.
ip rsvp signalling hello statistics	Enables Hello statistics on the router.
show ip rsvp hello statistics	Shows how long Hello packets have been in the Hello input queue.

clear ip rsvp hello statistics

To globally clear Hello statistics, use the **clear ip rsvp hello statistics** command in global EXEC mode.

clear ip rsvp hello statistics

Syntax Description

This command has no arguments or keywords.

Defaults

This command has no default behavior or values.

Command Modes

Global EXEC

Command History

Release	Modification
12.0(22)S	This command was introduced.

Usage Guidelines

Use this command to remove all information about how long Hello packets have been in the Hello input queue.

Examples

Following is sample output from the **show ip rsvp hello statistics** command and the **clear ip rsvp hello statistics** command. Notice that the values in the "Packet arrival queue" fields have been cleared.

```
Router# show ip rsvp hello statistics

  Status: Enabled
  Packet arrival queue:
    Wait times (msec)
      Current:0
      Average:0
      Weighted Average:0 (weight = 0.8)
      Max:4
    Current length: 0 (max:500)
  Number of samples taken: 2398525

Router# clear ip rsvp hello statistics

  Status: Enabled
  Packet arrival queue:
    Wait times (msec)
      Current:0
      Average:0
      Weighted Average:0 (weight = 0.8)
      Max:0
    Current length: 0 (max:500)
  Number of samples taken: 16
```

Related Commands

Command	Description
ip rsvp signalling hello statistics	Enables Hello statistics on the router.
show ip rsvp hello statistics	Shows how long Hello packets have been in the Hello input queue.

debug ip rsvp hello

To verify that a Hello instance has been created, a Hello instance has been deleted, and that communication with a neighbor has been lost, use the **debug ip rsvp hello** command in global EXEC mode.

debug ip rsvp hello [**stats**]

Syntax Description

stats	(Optional) Indicates whether statistics are enabled or disabled.

Defaults

This command has no default behavior or values.

Command Modes

Global EXEC

Command History

Release	Modification
12.0(22)S	This command was introduced.

Usage Guidelines

When you enter the **debug ip rsvp hello** command, RSVP signalling messages are shown, but RSVP HELLO messages are excluded because of the large number of HELLO messages that are sent.

Examples

Following is sample output from the **debug ip rsvp hello** command. The first portion of the output is for interface Se2/0 when Hello is created:

```
Router# debug ip rsvp hello

00:22:03: RSVP-HELLO: rsvp_hello_inst_init: Initializing ACTIVE hello inst
12.0.0.2->12.0.0.3
00:22:03: RSVP-HELLO: rsvp_hello_create_instance_from_psb: Next hop Se2/0 is adjacent
00:22:03: RSVP-HELLO: rsvp_hello_create_instance_from_psb: Create hello instance for
12.0.0.2->12.0.0.3 on Se2/0 (psb=61BC5F60)
00:22:03: RSVP-HELLO: rsvp_hello_find_instance: psb_cnt=2 for hello inst
12.0.0.2->12.0.0.3
00:22:03: RSVP-HELLO: rsvp_hello_incoming_message: Neighbor 12.0.0.3 state changed to UP
00:22:05: %LINK-3-UPDOWN: Interface Tunnel1, changed state to up
00:22:06: %LINEPROTO-5-UPDOWN: Line protocol on Interface Tunnel1, changed state to up
rsvp-3640-2(config-if)#
rsvp-3640-2(config-if)#shut
rsvp-3640-2(config-if)#
```

The following output shows that Hello has been deleted:

```
00:25:19: RSVP-HELLO: rsvp_hello_path_delete: psb for hello inst 12.0.0.2->12.0.0.3 exited
READY state (psb_cnt=1)
00:25:19: RSVP-HELLO: rsvp_hello_path_delete: psb for hello inst 12.0.0.2->12.0.0.3 exited
READY state (psb_cnt=0)
00:25:19: RSVP-HELLO: rsvp_hello_path_delete: Last psb deleted, hello inst for
12.0.0.2->12.0.0.3 ACTIVE->PASSIVE
```

```
00:25:19: RSVP-HELLO: rsvp_hello_path_delete: psb for hello inst 13.0.0.2->13.0.0.3 exited
READY state (psb_cnt=0)
00:25:19: RSVP-HELLO: rsvp_hello_path_delete: Last psb deleted, hello inst for
13.0.0.2->13.0.0.3 ACTIVE->PASSIVE
00:25:21: %LINK-5-CHANGED: Interface Tunnel1, changed state to administratively down
00:25:22: %LINEPROTO-5-UPDOWN: Line protocol on Interface Tunnel1,
changed state to down

00:05:51: RSVP-HELLO: Communication lost with 12.0.0.2
00:05:51: RSVP-HELLO: rsvp_hello_communication_lost: Neighbor 12.0.0.2 was reset
(src_inst)
```

Following is sample output from the **debug ip rsvp hello stats** command:

```
rsvp-3640-3(config)#ip rsvp sig hel stat
rsvp-3640-3(config)#end
rsvp-3640-3#
00:32:28: RSVP-HELLO: rsvp_hello_stats_init: Hello stats is being configured
```

Related Commands

Command	Description
ip rsvp signalling hello (configuration)	Enables Hello globally on the router.
ip rsvp signalling hello (interface)	Enables Hello on an interface where you need Fast ReRoute protection.
ip rsvp signalling hello dscp	Sets the DSCP value that is in the IP header of the Hello message sent out from an interface.
ip rsvp signalling hello refresh misses	Specifies how many Hello acknowledgments a node can miss in a row before the node considers that communication with its neighbor is down.
ip rsvp signalling hello refresh interval	Configures the Hello request interval.
ip rsvp signalling hello statistics	Enables Hello statistics on the router.

ip rsvp signalling hello (configuration)

To enable Hello globally on the router, use the **ip rsvp signalling hello** command in global configuration mode.

ip rsvp signalling hello

Syntax Description

This command has no arguments or keywords.

Defaults

This command has no default behavior or values.

Command Modes

Global configuration

Command History

Release	Modification
12.0(22)S	This command was introduced.

Usage Guidelines

To enable Hello globally on the router, you must enter this command. You also must enable Hello on the interface.

Examples

In the following example, Hello is enabled globally on the router:

```
Router(config-if)# ip rsvp signalling hello
```

Related Commands

Command	Description
ip rsvp signalling hello (interface)	Enables Hello on an interface where you need Fast ReRoute protection.
ip rsvp signalling hello statistics	Enables Hello statistics on the router.

ip rsvp signalling hello (interface)

To enable Hello on an interface where you need Fast ReRoute protection, use the **ip rsvp signalling hello** command in interface configuration mode.

ip rsvp signalling hello

Syntax Description

This command has no arguments or keywords.

Defaults

This command has no default behavior or values.

Command Modes

Interface configuration

Command History

Release	Modification
12.0(22)S	This command was introduced.

Usage Guidelines

You must configure Hello globally on the router and on the specific interface.

Examples

In the following example, Hello is enabled on an interface:

```
Router(config-if)# ip rsvp signalling hello
```

Related Commands

Command	Description
ip rsvp signalling hello (configuration)	Enables Hello globally on the router.
ip rsvp signalling hello dscp	Sets the DSCP value that is in the IP header of the Hello messages sent out from the interface.
ip rsvp signalling hello refresh misses	Specifies how many Hello acknowledgments a node can miss in a row before the node considers that communication with its neighbor is down.
ip rsvp signalling hello refresh interval	Configures the Hello request interval.

ip rsvp signalling hello dscp

To set the Differentiated Services Code Point (DSCP) value that is in the IP header of the Hello message sent out from an interface, use the **ip rsvp signalling hello dscp** command in interface configuration mode. To disable this feature, use the **no** form of this command.

ip rsvp signalling hello dscp [*num*]

no ip rsvp signalling hello dscp

Syntax Description

num	(Optional) DSCP value. Valid values are from 0 to 63.

Defaults

The default value is 0.

Command Modes

Interface configuration

Command History

Release	Modification
12.0(22)S	This command was introduced.

Usage Guidelines

If a link is congested, it is recommended that you set the DSCP to a value higher than zero (0) to reduce the likelihood that Hello messages will be dropped.

You configure the DSCP per interface, not per flow.

The DSCP applies to all RSVP flows installed on a specific interface. You can configure each interface independently for DSCP.

Examples

In the following example, Hello messages sent from this interface have a DSCP value of 48:

```
Router(config-if)# ip rsvp signalling hello dscp 48
```

Related Commands

Command	Description
ip rsvp signalling hello (interface)	Enables Hello on an interface where you need Fast ReRoute protection.

ip rsvp signalling hello refresh interval

To configure the Hello request interval, use the **ip rsvp signalling hello refresh interval** command in interface configuration mode.

ip rsvp signalling hello refresh interval *num*

Syntax Description

num	Frequency, in milliseconds, at which a node sends Hello messages to a neighbor. Valid values are from 10 to 30,000.

Defaults

200 milliseconds

Command Modes

Interface configuration

Command History

Release	Modification
12.0(22)S	This command was introduced.

Usage Guidelines

You can configure the Hello request interval on a per-neighbor basis. A node periodically generates a Hello message containing a HELLO REQUEST object for each neighbor whose status is being tracked. The frequency of those Hello messages is determined by the Hello interval.

Examples

In the following example, Hello requests are sent to a neighbor every 50 milliseconds:

```
Router(config-if)# ip rsvp signalling hello refresh interval 50
```

Related Commands

Command	Description
ip rsvp signalling hello (interface)	Enables Hello on an interface where you need Fast ReRoute protection.
show ip rsvp hello statistics	Shows how long Hello packets have been in the Hello input queue.

ip rsvp signalling hello refresh misses

To specify how many Hello acknowledgments a node can miss in a row before the node considers that communication with its neighbor is down, use the **ip rsvp signalling hello refresh misses** command in interface configuration mode.

ip rsvp signalling hello refresh misses *num*

Syntax Description

num	The number of sequential Hello acknowledgments that a node can miss. Valid values are from 4 to 10.

Defaults

The default is 4.

Command Modes

Interface configuration

Command History

Release	Modification
12.0(22)S	This command was introduced.

Usage Guidelines

Hello comprises a Hello message, a HELLO REQUEST object, and a HELLO ACK object. Each request is answered by an acknowledgment. If a link is very congested or has a very heavy load, set this number to a value higher than the default value to ensure that Hello does not falsely declare that a neighbor is down.

Examples

In the following example, if the node does not receive five Hello acknowledgments in a row, the node declares that its neighbor is down:

```
Router(config-if)# ip rsvp signalling hello refresh misses 5
```

Related Commands

Command	Description
ip rsvp signalling hello (interface)	Enables Help on an interface.
ip rsvp signalling hello dscp	Sets the DSCP value that is in Hello messages sent out from an interface.

ip rsvp signalling hello statistics

To enable Hello statistics on the router, use the **ip rsvp signalling hello statistics** command in global EXEC mode.

ip rsvp signalling hello statistics

Syntax Description

This command has no arguments or keywords.

Defaults

This command has no default behavior or values.

Command Modes

Global EXEC

Command History

Release	Modification
12.0(22)S	This command was introduced.

Examples

In the following example, Hello statistics are enabled on the router.

```
Router(config)# ip rsvp signalling hello statistics
```

Related Commands

Command	Description
clear ip rsvp hello instance statistics	Clears Hello statistics for an instance.
ip rsvp signalling hello (configuration)	Enables Hello globally on the router.
show ip rsvp hello statistics	Shows how long Hello packets have been in the Hello input queue.

mpls traffic-eng backup-path

To assign one or more backup tunnels to a protected interface, use the **mpls traffic-eng backup-path** command in interface configuration mode.

mpls traffic-eng backup-path tunnel*tunnel-id*

Syntax Description

tunnel*tunnel-id*	Tunnel ID of the backup tunnel that can be used in case of a failure.

Defaults

No backup tunnels are used if this interface goes down.

Command Modes

Interface configuration

Command History

Release	Modification
12.0(10)ST	This command was introduced.
12.0(16)ST	With Link Protection, this command selected the one-and-only backup tunnel for a given protected interface. If you enter the command twice, the second occurrence overwrites the first occurrence.
12.0(22)S	You can now enter this command multiple times to select multiple backup tunnels for a given protected interface. This can be done for both Link and Node Protection. The command is supported on the Cisco 10000 series ESRs.

Usage Guidelines

Enter this command on the interface to be protected (Link Protection), or on the interface whose downstream node is being protected (Node Protection). You can enter this command multiple times to select multiple backup tunnels for a given protected interface. An unlimited number of backup tunnels can be assigned to protect an interface. The only limitation is memory. By entering this command on a physical interface, LSPs using this interface (sending data *out of* this interface) can use the indicated backup tunnels if there is a link or node failure.

Examples

The following example assigns backup tunnel 34 to interface POS5/0:

```
Router(config)# interface pos5/0
Router(config-if)# mpls traffic-eng backup-path tunnel34
```

Related Commands

Command	Description
tunnel mpls traffic-eng fast-reroute	Enables an MPLS Traffic Engineering tunnel to use a backup tunnel if there is a link or node failure (provided that a backup tunnel exists).

mpls traffic-eng fast-reroute timers

To specify how often the router considers switching an LSP to a new (better) backup tunnel if additional backup-bandwidth becomes available, use the **mpls traffic-eng fast-reroute timers** command in global configuration mode. To disable this timer, set the *frequency* to zero or use the **no** form of this command.

mpls traffic-eng fast-reroute timers [**frequency** *frequency*]

no mpls traffic-eng fast-reroute timers

Syntax Description

frequency *frequency*	(Optional) Interval, in seconds, between scans to determine if an LSP should use a new, better backup tunnel. Valid values are from 0 to 604800. A value of 0 disables promotions to a better LSP.

Defaults

By default, the timer is running and is set to a frequency of every 300 seconds (5 minutes). If you enter **no mpls traffic-eng fast-reroute timers**, the router returns to this default behavior.

Command Modes

Global configuration

Command History

Release	Modification
12.0(22)S	This command was introduced.

Examples

In the following example, LSPs are scanned every 2 minutes (120 seconds) to see if they should be promoted to a better backup tunnel.

```
Router(config)# mpls traffic-eng fast-reroute timers frequency 120
```

show ip rsvp hello

To show if Hello is enabled globally on the router and if Hello statistics are enabled, use the **show ip rsvp hello** command in global EXEC mode.

show ip rsvp hello

Syntax Description

This command has no arguments or keywords.

Defaults

This command has no default behavior or values.

Command Modes

Global EXEC

Command History

Release	Modification
12.0(22)S	This command was introduced.

Examples

The following is sample output from the **show ip rsvp hello** command:

```
Router# show ip rsvp hello

  State: Enabled
  Statistics: Enabled

Default State: Disabled
Default Statistics: Disabled
```

Table 2 describes significant fields displayed in this example.

Table 2 show ip rsvp hello Field Descriptions

Field	Description
State	Status of whether Hello is globally enabled on the router.
Statistics	Status of Hello statistics. Valid values are: • Enabled—Statistics are configured. Hello packets are time-stamped when they arrive in the Hello input queue for the purpose of recording the time it takes until they are processed. • Disabled—Hello statistics are not configured. • Shutdown—Hello statistics are configured but not operational. The input queue is too long (that is, more than 10,000 packets are queued).

Related Commands

Command	Description
ip rsvp signalling hello (configuration)	Enables Hello globally on the router.
ip rsvp signalling hello statistics	Enables Hello statistics on the router.
show ip rsvp hello statistics	Shows how long Hello packets have been in the Hello input queue.

show ip rsvp hello instance detail

To show detailed information about a Hello instance, use the **show ip rsvp hello instance detail** command in global EXEC mode.

show ip rsvp hello instance detail [**filter destination** *ip-address*]

Syntax Description

filter destination *ip-address*	(Optional) IP address of the neighbor node.

Defaults

This command has no default behavior or values.

Command Modes

Global EXEC

Command History

Release	Modification
12.0(22)S	This command was introduced.

Examples

The following is sample output from a **show ip rsvp hello instance detail** command:

```
Router# show ip rsvp hello instance detail

Neighbor 11.0.0.2  Source  11.0.0.1
    State: UP      (for 2d18h)
    Type: PASSIVE  (responding to requests)
    I/F: Et1/1
    LSPs protecting: 0
    Refresh Interval (msec) (used when ACTIVE)
      Configured: 100
      Statistics: (from 2398195 samples)
        Min:      100
        Max:      132
        Average:  100
        Waverage: 100 (Weight = 0.8)
        Current:  100
    Src_instance 0xA9F07C13, Dst_instance 0x9BBAA407
    Counters:
      Communication with neighbor lost:
        Num times: 0
        Reasons:
          Missed acks:             0
          Bad Src_Inst received:   0
          Bad Dst_Inst received:   0
          I/F went down:           0
          Neighbor disabled Hello: 0
      Msgs Received:   2398194
           Sent:       2398195
           Suppressed: 0
```

Table 3 describes the fields displayed in this example.

Table 3 show ip rsvp hello instance detail Field Descriptions

Field	Description
Neighbor	IP address of the adjacent node.
Source	IP address of the node that is sending the Hello message.
State	Status of communication. Values are UP (node is communicating with its neighbor) and LOST (communication has been lost or never was established).
Type	Values are ACTIVE (node is sending requests) and PASSIVE (node is responding to a request).
I/F	Interface type.
LSPs protecting	Number of LSPs that are being protected.
Refresh Interval Configured	The frequency with which a node generates a Hello message containing a HELLO REQUEST object for each neighbor whose status is being tracked. The frequency of these Hello messages is determined by the Hello interval specified in the **ip rsvp signalling hello refresh interval** command.
Min	Minimum refresh interval.
Max	Maximum refresh interval.
Average	Average refresh interval.
Waverage	Weighted average refresh interval.
Current	Current refresh interval.
Src_instance	Source instance field value.
Dst_instance	Destination instance field value.
Communication with neighbor lost	Subsequent fields designate the number of times that communication with the neighbor was lost and why.
Num times	Total number of times that communication with neighbor was lost.
Reasons	Subsequent fields designate why communication with the neighbor was lost.
Missed acks	Number of times that communication was lost due to missed ACKs.
Bad Src_Inst received	Number of times that communication was lost due to bad Bad Src_Inst fields.
Bad Dst_Inst received	Number of times that communication was lost due to bad Dst_Inst fields.
I/F went down	Number of times that the interface became unoperational.
Neighbor disabled Hello	Number of times that neighbor disabled Hello.

Table 3 show ip rsvp hello instance detail Field Descriptions (continued)

Field	Description
Msgs Received	Number of messages that were received.
Sent	Number of messages that were sent.
Suppressed	Number of messages that were suppressed due to optimization.

Related Commands

Command	Description
ip rsvp signalling hello (configuration)	Enables Hello globally on the router.
ip rsvp signalling hello statistics	Enables Hello statistics on the router.
show ip rsvp hello	Shows if Hello is enabled globally on the router and if Hello statistics are enabled.
show ip rsvp hello instance summary	Shows summary information about a Hello instance.

show ip rsvp hello instance summary

To show summary information about a Hello instance, use the **show ip rsvp hello instance summary** command in global EXEC mode.

show ip rsvp hello instance summary

Syntax Description

This command has no arguments or keywords.

Defaults

This command has no default behavior or values.

Command Modes

Global EXEC

Command History

Release	Modification
12.0(22)S	This command was introduced.

Examples

The following is sample output from the **show ip rsvp hello instance summary** command:

```
Router# show ip rsvp hello instance summary

I/F      Neighbor   Type     State  LostCnt
Et1/1    11.0.0.1   PASSIVE  UP     0
Se2/0    12.0.0.3   ACTIVE   UP     0
Et1/2    13.0.0.3   ACTIVE   UP     0
```

Table 4 describes the fields displayed in this example.

Table 4 show ip rsvp hello instance summary Field Descriptions

Field	Description
I/F	Interface.
Neighbor	IP address of adjacent node.
Type	Activity. Values are ACTIVE (node is sending requests) and PASSIVE (node is responding to a request).
State	Status of communication. Values are UP (node is communicating with its neighbor) and LOST (communication has been lost or never was established).
LostCnt	Number of times that communication was lost with the neighbor.

Related Commands

Command	Description
ip rsvp signalling hello (configuration)	Enables Hello globally on the router.

ip rsvp signalling hello statistics	Enables Hello statistics on the router.
show ip rsvp hello	Shows if Hello is enabled globally on the router and if Hello statistics are enabled.
show ip rsvp hello instance detail	Shows detailed information about a Hello instance.

show ip rsvp hello statistics

To show how long Hello packets have been in the Hello input queue, use the **show ip rsvp hello statistics** command in global EXEC mode.

show ip rsvp hello statistics

Syntax Description

This command has no arguments or keywords.

Defaults

This command has no default behavior or values.

Command Modes

Global EXEC

Command History

Release	Modification
12.0(22)S	This command was introduced.

Usage Guidelines

You can use this command to determine if the Hello refresh interval is too small. If the interval is too small, communication may falsely be declared as lost.

Examples

The following is sample output from the **show ip rsvp hello statistics** command:

```
Router# show ip rsvp hello statistics

Status: Enabled
  Packet arrival queue:
    Wait times (msec)
      Current:0
      Average:0
      Weighted Average:0 (weight = 0.8)
      Max:4
    Current length: 0 (max:500)
  Number of samples taken: 2398525
```

Table 5 describes the fields displayed in this example.

Table 5 show ip rsvp hello statistics Field Descriptions

Field	Description
Status	Indicator of whether Hello has been enabled globally on the router.
Current	Amount of time, in milliseconds, that the current Hello packet has been in the Hello input queue.
Average	Average amount of time, in milliseconds, that Hello packets are in the Hello input queue.

Table 5 show ip rsvp hello statistics Field Descriptions (continued)

Field	Description
Max	Maximum amount of time, in milliseconds, that Hello packets have been in the Hello input queue.
Current length	Current amount of time, in milliseconds, that Hello packets have been in the Hello input queue.
Number of samples taken	Number of packets for which these statistics were compiled.

Related Commands

Command	Description
clear ip rsvp hello instance statistics	Clears Hello statistics for an instance.
clear ip rsvp hello statistics	Globally clears Hello statistics.
ip rsvp signalling hello refresh interval	Configures the Hello request interval.
ip rsvp signalling hello statistics	Enables Hello statistics on the router.

show ip rsvp interface detail

To show the interface configuration for Hello, use the **show ip rsvp interface detail** command in global EXEC mode.

show ip rsvp interface detail [*interface*]

Syntax Description

interface	(Optional) Interface for which you want to show the Hello configuration.

Defaults

This command has no default behavior or values.

Command Modes

Global EXEC

Command History

Release	Modification
12.0(22)S	This command was introduced.

Examples

The following is sample output from the **show ip rsvp interface detail** command:

```
Router# show ip rsvp interface detail Et1/2

Et1/2:
   Bandwidth:
     Curr allocated: 0G bits/sec
     Max. allowed (total): 7500K bits/sec
     Max. allowed (per flow): 7500K bits/sec
     Max. allowed for LSP tunnels using sub-pools: 0G bits/sec
   Neighbors:
     Using IP encap: 1.  Using UDP encap: 0
     DSCP value used in RSVP msgs: 0x0
   Hello:
     State: Enabled
     Refresh Interval: 500
     Missed Acks: 4
     DSCP value used in HELLO msgs: 0
```

Table 6 describes the fields displayed in this example.

Table 6 show ip rsvp interface detail Field Descriptions

Field	Description
Curr allocated	Amount of bandwidth currently allocated.
Max. allowed (total)	Total maximum amount of bandwidth allowed.
Max. allowed (per flow)	Maximum amount of bandwidth allowed per flow.
Max. allowed for LSP tunnels using sub-pools	Maximum amount of bandwidth permitted for LSP tunnels that obtain their bandwidth from sub-pools.

Table 6 show ip rsvp interface detail Field Descriptions (continued)

Field	Description
Using IP encap	Number of neighbors using IP encapsulation.
Using UDP encap	Number of neighbors using UDP encapsulation.
DSCP value used in RSVP msgs	The Differentiated Services Code Point value that is in RSVP messages.
State	State (Enabled or Disabled) of Hello.
Refresh Interval	Frequency with which a node sends a Hello message to its neighbor.
Missed Acks	Number of sequential acknowledgments that the node did not receive.
DSCP value used in HELLO msgs	The Differentiated Services Code Point value that is in Hello messages.

Related Commands

Command	Description
ip rsvp signalling hello (interface)	Enables Hello on an interface where you need Fast ReRoute protection.
ip rsvp signalling hello dscp	Sets the DSCP value that is in the IP header of the Hello message sent out from an interface.
ip rsvp signalling hello refresh interval	Configures the Hello request interval.

show ip rsvp request

To display upstream reservation state (that is, information related to the Resv messages that this node will send upstream), use the **show ip rsvp request** command in EXEC mode.

show ip rsvp request [**detail**] {**destination** *ipaddress* | **source** *ipaddress* | **dst-port** *pnum* | **src-port** *pnum*}

Syntax Description

detail	(Optional) Displays additional request information.
destination *ipaddress*	(Optional) Displays information only for tunnels that go to a particular destination (tail) that is specified in the IP address.
source *ipaddress*	(Optional) Displays information only for tunnels originating at the headend specified in the IP address.
dst-port *pnum*	(Optional) Destination port. Displays information only for tunnels that have the tunnel ID specified in the *pnum* field.
src-port *pnum*	(Optional) Source port. Displays information only for tunnels that have the LSP ID (also called the tunnel instance) specified in the *pnum* field.

Defaults

This command has no default behavior or values.

Command Modes

EXEC

Command History

Release	Modification
11.2	This command was introduced.
12.2	The **detail** keyword was added to display additional request information.
12.0(22)S	This command has been enhanced to show some useful Fast ReRoute information for when an LSP is actively using a backup tunnel that terminates at this node (that is, when a node is the Merge Point.) If desired, information for only a single tunnel or a subset of tunnels can be displayed. The command is supported on the Cisco 10000 series ESRs.

Usage Guidelines

When hundreds or thousands of tunnels exist and you are interested in only a few, it is useful to display output for only a single tunnel or a subset of tunnels. To request a limited display, enter the **show ip rsvp request** command with the appropriate keyword (which in this case is called an output filter): **destination**, **source**, **dst-port**, and **src-port**. You can enter any or all of the output filters, and you can enter them whether or not you specify the **detail** keyword.

Examples

Following is sample output from the **show ip rsvp request detail** command when the command is entered on the Merge Point (MP) before a failure and after a failure.

Note Refer to Figure 9.

Example 1: Command is entered on the MP before a failure

```
Router# show ip rsvp request detail

RSVP Reservation. Tun Dest: 24.1.1.1 Tun Sender: 23.1.1.1,
   Tun ID: 1 LSP ID: 126
   Next Hop is 14.1.1.1 on POS0/1
   Label is 0
   Reservation Style is Shared-Explicit, QoS Service is Controlled-Load
   Average Bitrate is 0G bits/sec, Maximum Burst is 1K bytes
   RRO:
    Empty
```

Example 2: Command is entered on the MP after a failure

```
Router# show ip rsvp request detail

RSVP Reservation. Tun Dest: 24.1.1.1 Tun Sender: 23.1.1.1,
   Tun ID: 1 LSP ID: 126
   Next Hop is 14.1.1.1 on POS0/1
   Label is 0
   Reservation Style is Shared-Explicit, QoS Service is Controlled-Load
   Average Bitrate is 0G bits/sec, Maximum Burst is 1K bytes
   RRO:
     Empty
   FRR is in progress (we are Merge Point)

 RSVP Reservation. Tun Dest: 24.1.1.1 Tun Sender: 23.1.1.1,
   Tun ID: 1 LSP ID: 126
   Next Hop is 15.1.1.1 on POS0/1
   Label is 0
   Reservation Style is Shared-Explicit, QoS Service is Controlled-Load
   Average Bitrate is 0G bits/sec, Maximum Burst is 1K bytes
   RRO:
     Empty
   FRR is in progress (we are Merge Point)
```

Notice that after the failure, there are two entries for the rerouted LSP. Information referenced in the following explanation is highlighted.

The first entry continues to show the pre-failure information (i.e., Resv messages are being sent to 14.1.1.1 on Ethernet1). This state is for the Resv being sent upstream before the failure, in response to Path messages sent before the failure. This state may time out quickly, or it may continue to be refreshed for a few minutes if, for example, an upstream node is unaware of the failure.

The second entry shows the post-failure information (i.e., Resv messages are being sent to 15.1.1.1 on Ethernet2). This state is for the Resv messages being sent upstream after the failure (to the PLR), and will remain and be refreshed as long as the LSP is rerouted.

In this example, the MP is also the tail of the LSP. There is no RRO information because there are no nodes downstream.

show ip rsvp reservation

To display downstream reservation state information (that is, information related to the Resv message arriving from downstream), use the **show ip rsvp reservation** command in EXEC mode.

show ip rsvp reservation [**detail**] {**destination** *ipaddress* | **source** *ipaddress* | **dst-port** *pnum* | **src-port** *pnum*}

Syntax Description

detail	(Optional) Displays additional reservation information.
destination *ipaddress*	(Optional) Displays information only for tunnels that go to a particular destination (tail) that is specified in the IP address.
source *ipaddress*	(Optional) Displays information only for tunnels originating at the headend specified in the IP address.
dst-port *pnum*	(Optional) Destination port. Displays information only for tunnels that have the tunnel ID specified in the *pnum* field.
src-port *pnum*	(Optional) Source port. Displays information only for tunnels that have the LSP ID (also called the tunnel instance) specified in the *pnum* field.

Defaults

This command has no default behavior or values.

Command Modes

EXEC

Command History

Release	Modification
11.2	This command was introduced.
12.2	The **detail** keyword was added to display additional reservation information.
12.0(22)S	The command displays useful Fast ReRoute information when an LSP is actively using a backup tunnel at this node (that is, when a node is the PLR). If desired, information for only a single tunnel or a subset of tunnels can be displayed. The command is supported on the Cisco 10000 series ESRs.

Usage Guidelines

When hundreds or thousands of tunnels exist and you are interested in only a few, it is useful to display output for only a single tunnel or a subset of tunnels. To request a limited display, enter the **show ip rsvp reservation** command with the appropriate keyword (which in this case is called an output filter): **destination**, **source**, **dst-port**, and **src-port**. You can enter any or all of the output filters, and you can enter them whether or not you specify the **detail** keyword.

Examples

Following is sample output from the **show ip rsvp reservation detail** command when the command is entered on the Point of Local Repair (PLR) before a failure and after a failure.

Note Refer to Figure 9.

Example 1: Command is entered on the PLR before a failure

```
Router# show ip rsvp reservation detail

RSVP Reservation. Tun Dest: 24.1.1.1 Tun Sender: 23.1.1.1,
   Tun ID: 1 LSP ID: 126
   Next Hop is 12.1.1.2 on POS1/2
   Label is 18
   Reservation Style is Shared-Explicit, QoS Service is Controlled-Load
   Average Bitrate is 0G bits/sec, Maximum Burst is 1K bytes
   RRO:
     14.1.1.1/32, Flags:0x0 (No Local Protection)
       Label record: Flags 0x1, ctype 1, incoming label 18
     14.1.1.2/32, Flags:0x0 (No Local Protection)
       Label record: Flags 0x1, ctype 1, incoming label 0
```

Example 2: Command is entered on the PLR after a failure

```
Router# show ip rsvp reservation detail

RSVP Reservation. Tun Dest: 24.1.1.1 Tun Sender: 23.1.1.1,
   Tun ID: 1 LSP ID: 126
   FRR is in progress: (we are PLR)
    Bkup Next Hop is  16.1.1.2 on POS1/1
        Label     is  0
    Orig Next Hop was 12.1.1.2 on POS1/2
        Label     was 18
   Reservation Style is Shared-Explicit, QoS Service is Controlled-Load
   Average Bitrate is 0G bits/sec, Maximum Burst is 1K bytes
  RRO:
    24.1.1.1/32, Flags:0x0 (No Local Protection)
      Label record: Flags 0x1, ctype 1, incoming label 0
```

Notice the following (see highlighted text) in Examples 1 and 2:

- At the PLR, you see “FRR is in progress (we are PLR)” when an LSP has been rerouted (that is, it is actively using a backup tunnel).
- Resv messages arrive on a different interface and from a different Next Hop after a failure. The pre-failure display shows the original NHOP and arriving interface; the post-failure display shows both the original and the new (Bkup) NHOP and arriving interface. The label is also shown.
- The RRO in arriving Resv messages changes after the failure, given that the Resv messages will avoid the failure (that is, it will traverse different links and/or hops).

Related Commands

Command	Description
ip rsvp reservation	Enables a router to simulate RSVP Resv message reception from the sender.
clear ip rsvp hello instance counters	Clears (refreshes) the values for Hello instance counters.

show ip rsvp sender

To display path state information (that is, information related to the Path messages arriving from upstream), and the state of Fast ReRoute for a given MPLS Traffic Engineering LSP, use the **show ip rsvp sender** command in EXEC mode.

show ip rsvp sender [**detail**] {**destination** *ipaddress* | **source** *ipaddress* | **dst-port** *pnum* | **src-port** *pnum*}

Syntax Description

detail	(Optional) Displays additional sender information.
destination *ipaddress*	(Optional) Displays information only for tunnels that go to a particular destination (tail) that is specified in the IP address.
source *ipaddress*	(Optional) Displays information only for tunnels originating at the headend specified in the IP address.
dst-port *pnum*	(Optional) Destination port. Displays information only for tunnels that have the tunnel ID specified in the *pnum* field.
src-port *pnum*	(Optional) Source port. Displays information only for tunnels that have the LSP ID (also called the tunnel instance) specified in the *pnum* field.

Defaults

This command has no default behavior or values.

Command Modes

EXEC

Command History

Release	Modification
11.2	This command was introduced.
12.0(22)S	The command output includes additional information that can be helpful when examining Fast ReRoute state or when troubleshooting Fast ReRoute. If desired, information for only a single tunnel or a subset of tunnels can be displayed. The command is supported on the Cisco 10000 series ESRs.

Usage Guidelines

This command is very useful for determining the state of RSVP signalling both before and after an LSP has been fast rerouted. The command is most useful when you enter it at the Point of Local Repair (PLR) or at the Merge Point (MP).

When hundreds or thousands of tunnels exist and you are interested in only a few, it is useful to display output for only a single tunnel or a subset of tunnels. To request a limited display, enter the **show ip rsvp sender** command with the appropriate keyword (which in this case is called an output filter): **destination**, **source**, **dst-port**, and **src-port**. You can enter any or all of the output filters, and you can enter them whether or not you specify the **detail** keyword.

Examples

Following is sample output from the **show ip rsvp sender detail** command under the following circumstances:

- Command is entered at the PLR before a failure (see Example 1)
- Command is entered at the PLR after a failure (see Example 2)
- Command is entered at the MP before a failure (see Example 3)
- Command is entered at the MP after a failure (see Example 4)
- Command output shows all senders (see Example 5)
- Command output only shows senders who have a specific destination (see Example 6)
- Show more detail about a sender who has a specific destination (see Example 7)

Note Refer to Figure 9.

Example 1: Command is entered at the PLR before a failure

The following is sample output from the **show ip rsvp sender detail** command when it is entered at the PLR before a failure:

```
Router# show ip rsvp sender detail

PATH:
   Tun Dest: 24.1.1.1 Tun ID: 1 Ext Tun ID: 23.1.1.1
   Tun Sender: 23.1.1.1, LSP ID: 126
   Path refreshes arriving on POS1/0 from PHOP 11.1.1.1
   Path refreshes being sent to NHOP 12.1.1.2 on POS1/1
   Session Attr::
     Setup Prio: 0, Holding Prio: 0
     Flags: Local Prot desired, Label Recording, SE Style
     Session Name:tagsw4500-23_t1
   ERO:
     12.1.1.2 (Strict IPv4 Prefix, 8 bytes, /32)
     14.1.1.1 (Strict IPv4 Prefix, 8 bytes, /32)
     14.1.1.2 (Strict IPv4 Prefix, 8 bytes, /32)
     24.1.1.1 (Strict IPv4 Prefix, 8 bytes, /32)
   Traffic params - Rate: 0G bits/sec, Max. burst: 1K bytes
   Fast-Reroute Backup info:
     Inbound  FRR: Not active
     Outbound FRR: Ready -- backup tunnel selected
       Backup Tunnel: Tu2        (label 0)
       Bkup Sender Template:
         Tun Sender: 15.1.1.1, LSP ID: 126
       Bkup FilerSpec:
         Tun Sender: 15.1.1.1, LSP ID 126
```

Table 7 describes the significant fields.

Note The flags field is important for Fast ReRoute. For information about flags that must be set, see the Flags field description in Table 7.

Table 7 show ip rsvp sender detail Field Descriptions —on PLR Before Failure

Field	Description
The first five fields provide information that uniquely identifies the LSP. **The first three fields identify the LSP's session (that is, the contents of the SESSION object in arriving Path messages).**	
Tun Dest	IP address of the destination of the tunnel.
Tun ID	Tunnel identification number.
Ext Tun ID	Extended tunnel identification number.
The next two fields identify the LSP's sender (SENDER_TEMPLATE object of arriving Path messages).	
Tun Sender	Tunnel sender.
LSP ID	LSP identification number.
The remaining fields indented under PATH provide additional information about this LSP.	
Session Attr—Session attributes. Refers to information included in the SESSION_ATTRIBUTE object of arriving Path messages, such as the Setup and Holding Priorities, Flags, and the Session Name.	
Setup Prio	Setup priority.
Holding Prio	Holding priority.
Flags	An LSP must have the "Local protection desired" Flag of the SESSION_ATTRIBUTE object set for the LSP to use a backup tunnel (that is, in order to receive local protection). If this flag is not set, you have not enabled Fast ReRoute for this tunnel at its headend (by entering the **tunnel mpls traffic-eng fast-reroute** command). NNHOP backup tunnels rely on label recording, so LSPs should have the "label recording desired" flag set too. This flag is set if the tunnel was configured for Fast ReRoute.
ERO—Refers to the EXPLICIT_ROUTE object of the Path messages. This field displays the contents of the ERO at this node. As a Path message travels from the sender (headend) to the receiver (tailend), each node removes its own IP address from the ERO. The displayed value reflects the remainder of hops between this node and the tail.	
Fast-Reroute Backup info—Information that is relevant to Fast ReRoute for this LSP.	
Inbound FRR	If this node is downstream from a rerouted LSP (for example, at a Merge Point for this LSP), the state is Active.

Table 7 show ip rsvp sender detail Field Descriptions —on PLR Before Failure (continued)

Field	Description
Outbound FRR	If this node is a PLR for an LSP, there are three possible states: • Active—This LSP is actively using its backup tunnel, presumably because there has been a downstream failure. • No Backup—This LSP does not have local (Fast ReRoute) protection. No backup tunnel has been selected for it to use in case of a failure. • Ready—This LSP is ready to use a backup tunnel in case of a downstream link or node failure. A backup tunnel has been selected for it to use.
Backup Tunnel	If the Outbound FRR state is Ready or Active, this field indicates the following: • Which backup tunnel has been selected for this LSP to use in case of a failure. • The inbound label that will be prepended to the LSP's data packets for acceptance at the backup tunnel tail (the Merge Point).

Table 7 show ip rsvp sender detail Field Descriptions —on PLR Before Failure (continued)

Field	Description
Bkup Sender Template	If the Outbound FRR state is Ready or Active, SENDER_TEMPLATE and FILTERSPEC objects are shown. These objects will be used in RSVP messages sent by the backup tunnel if/when the LSP starts actively using the backup tunnel. They differ from the original (pre-failure) objects only in that the node (the PLR) substitutes its own IP address for that of the original sender. For example, Path and PathTear messages will contain the new SENDER_TEMPLATE. Resv and ResvTear messages will contain the new FILTERSPEC object. If this LSP begins actively using the backup tunnel, the display changes as shown below.
Bkup FilerSpec	If the Outbound FRR state is Ready or Active, SENDER_TEMPLATE and FILTERSPEC objects are shown. These objects will be used in RSVP messages sent by the backup tunnel if/when the LSP starts actively using the backup tunnel. They differ from the original (pre-failure) objects only in that the node (the PLR) substitutes its own IP address for that of the original sender. For example, Path and PathTear messages will contain the new SENDER_TEMPLATE. Resv and ResvTear messages will contain the new FILTERSPEC object. If this LSP begins actively using the backup tunnel, the display changes as shown in Example 2.

Example 2: Command is entered at the PLR after a failure

If the LSP begins actively using the backup tunnel and the command is entered at the PLR after a failure, the display changes as shown below.

Note Highlighted fields are referenced in the explanation that follows the sample display.

```
Router# show ip rsvp sender detail

PATH:
   Tun Dest: 24.1.1.1 Tun ID: 1 Ext Tun ID: 23.1.1.1
   Tun Sender: 23.1.1.1, LSP ID: 126
   Path refreshes arriving on POS1/0 from PHOP 11.1.1.1
   Path refreshes being sent to NHOP 24.1.1.1 on Tunnel2
   Session Attr::
     Setup Prio: 0, Holding Prio: 0
     Flags: Local Prot desired, Label Recording, SE Style
     Session Name:tagsw4500-23_t1
   ERO:
     24.1.1.1 (Strict IPv4 Prefix, 8 bytes, /32)
     24.1.1.1 (Strict IPv4 Prefix, 8 bytes, /32)
   Traffic params - Rate: 0G bits/sec, Max. burst: 1K bytes
```

```
  Fast-Reroute Backup info:
    Inbound  FRR: Not active
    Outbound FRR: Active -- using backup tunnel
      Backup Tunnel: Tu2        (label 0)
      Bkup Sender Template:
        Tun Sender: 15.1.1.1, LSP ID: 126
      Bkup FilerSpec:
        Tun Sender: 15.1.1.1, LSP ID 126
      Orig Output I/F: Et2
      Orig Output ERO:
        12.1.1.2 (Strict IPv4 Prefix, 8 bytes, /32)
        14.1.1.1 (Strict IPv4 Prefix, 8 bytes, /32)
        14.1.1.2 (Strict IPv4 Prefix, 8 bytes, /32)
        24.1.1.1 (Strict IPv4 Prefix, 8 bytes, /32)
```

Once an LSP is actively using a backup tunnel, the following changes occur:

- Path refreshes are no longer sent to the original NHOP out the original interface. They are sent through the backup tunnel to the node that is the tail of the backup tunnel (NHOP or NNHOP).
- The ERO is modified so that it will be acceptable upon arrival at the NHOP or NNHOP.
- The display shows both the original ERO and the new one now being used.
- The display shows the original output interface (that is, the interface from which Path messages were sent for this LSP before the failure).

Example 3: Command is entered at the MP before a failure

If the same **show ip rsvp sender** command is entered at the Merge Point (the backup tunnel tail), the display changes from before to after the failure. Following is sample output before a failure:

```
Router# show ip rsvp sender detail

PATH:
   Tun Dest: 24.1.1.1 Tun ID: 1 Ext Tun ID: 23.1.1.1
   Tun Sender: 23.1.1.1, LSP ID: 126
   Path refreshes arriving on POS0/0 from PHOP 14.1.1.1
   Session Attr::
     Setup Prio: 0, Holding Prio: 0
     Flags: Local Prot desired, Label Recording, SE Style
     Session Name:tagsw4500-23_t1
   Traffic params - Rate: 0G bits/sec, Max. burst: 1K bytes
   Fast-Reroute Backup info:
     Inbound  FRR: Not active
     Outbound FRR: No backup tunnel selected
```

Example 4: Command is entered at the MP after a failure

After a failure, the following changes occur:

- The interface and previous hop (PHOP) from which Path messages are received will change.
- The inbound FRR becomes Active.
- The original PHOP and the original input interface are displayed as shown below.

Following is sample output after a failure:

```
Router# show ip rsvp sender detail

PATH:
   Tun Dest: 24.1.1.1 Tun ID: 1 Ext Tun ID: 23.1.1.1
   Tun Sender: 23.1.1.1, LSP ID: 126
   Path refreshes arriving on POS0/1 from PHOP 15.1.1.1 on Loopback0
   Session Attr::
     Setup Prio: 0, Holding Prio: 0
     Flags: Local Prot desired, Label Recording, SE Style
     Session Name:tagsw4500-23_t1
   Traffic params - Rate: 0G bits/sec, Max. burst: 1K bytes
   Fast-Reroute Backup info:
     Inbound  FRR: Active
       Orig Input I/F: POS0/0
       Orig PHOP: 14.1.1.1
       Now using Bkup Filterspec w/ sender: 15.1.1.1 LSP ID: 126
     Outbound FRR: No backup tunnel selected
```

Notice the following changes, which are highlighted in the sample command output:

- After a failure, Path refreshes arrive on a different interface and from a different PHOP.
- The original PHOP and input interface are shown under Fast-Reroute Backup information, along with the FILTERSPEC object that will now be used when sending messages (such as Resv and ResvTear).

Example 5: Command output shows all senders

In the following example, information about all senders is displayed.

```
Router# show ip rsvp sender

To              From        Pro DPort Sport Prev Hop        I/F  BPS  Bytes
24.1.1.1        23.1.1.1    0   1     59    11.1.1.1        Et1   0G     1K
24.1.1.1        25.1.1.1    0   2     9                           0G     1K
24.1.1.1        23.1.1.1    0   3     12    11.1.1.1        Et1   0G     1K
24.1.1.1        25.1.1.1    0   3     20                          0G     1K
26.1.1.1        25.1.1.1    0   0     23                          0G     1K
26.1.1.1        25.1.1.1    0   1     22                          0G     1K
26.1.1.1        25.1.1.1    0   1000  22                          0G     1K
```

Table 8 describes the fields displayed in this example.

Table 8 show ip rsvp sender Field Descriptions

Field	Description
To	IP address of the receiver.
From	IP address of the sender.
Pro	Protocol code. Code 1 indicates Internet Control Message Protocol (ICMP).
DPort	Destination port number.
Sport	Source port number.
Prev Hop	IP address of the previous hop.
I/F	Interface of the previous hop.

Table 8 show ip rsvp sender Field Descriptions (continued)

Field	Description
BPS	Reservation rate, in bits per second, the application is advertising it might achieve.
Bytes	Bytes of burst size the application is advertising it might achieve.

Example 6: Command output only shows senders who have a specific destination

To only show information about senders who have a specific destination, specify the destination filter as shown below. In this example, the destination is 26.6.6.6.

```
Router# show ip rsvp sender destination 26.6.6.6

To          From         Pro DPort Sport Prev Hop    I/F  BPS  Bytes
26.1.1.1    25.1.1.1     0   0     23                     0G     1K
26.1.1.1    25.1.1.1     0   1     22                     0G     1K
26.1.1.1    25.1.1.1     0   1000  22                     0G     1K
```

Example 7: Show more detail about a sender who has a specific destination

To show more detail about the sender whose destination is 1000 (as shown in Example 6), specify the command with the destination port filter.

```
Router# show ip rsvp sender detail dst-port 1000

PATH:
  Tun Dest 26.1.1.1 Tun ID 1000 Ext Tun ID 25.1.1.1
  Tun Sender: 25.1.1.1, LSP ID: 22
  Path refreshes being sent to NHOP 12.1.1.2 on Ethernet2
  Session Attr::
    Setup Prio: 7, Holding Prio: 7
    Flags: SE Style
    Session Name:tagsw4500-25_t1000
  ERO:
    12.1.1.2 (Strict IPv4 Prefix, 8 bytes, /32)
    26.1.1.1 (Strict IPv4 Prefix, 8 bytes, /32)
  Traffic params - Rate: 0G bits/sec, Max. burst: 1K bytes
  Fast-Reroute Backup info:
    Inbound  FRR: Not active
    Outbound FRR: No backup tunnel selected
```

Related Commands

Command	Description
ip rsvp sender	Enables a router to simulate RSVP Path message reception from the sender.

show mpls traffic tunnel backup

To display information about the backup tunnels that are currently configured, use the **show mpls traffic tunnel backup** command in EXEC mode.

show mpls traffic tunnel backup tunnel*tunnel-id*

Syntax Description

tunnel*tunnel-id*	Tunnel ID of the backup tunnel for which you want to display information.

Defaults

This command has no default behavior or values.

Command Modes

EXEC

Command History

Release	Modification
12.0(22)S	This command was introduced.

Examples

The following is sample output from the **show mpls traffic tunnel backup** command.

```
Router# show mpls traffic tunnel backup tunnel1000

Tunnel1000          Dest: 12.0.0.9        State: Up
any-pool cfg 100 inuse 0 num_lsps 0
     protects: ATM0.1
```

Table 9 describes the fields displayed in this example.

Table 9 show mpls traffic tunnel backup Field Descriptions

Field	Description
Tunnel	Tunnel ID of the backup tunnel for which this information is being displayed.
Dest	IP address of the destination of the backup tunnel.
State	State of the backup tunnel. Valid values are Up, Down, or Admin-down.
any-pool	Pool from which bandwidth is acquired. Valid values are any-pool, global-pool, and sub-pool.
cfg	Amount of bandwidth configured for that pool.
inuse	Amount of bandwidth currently being used.
num_lsps	Number of LSPs being protected.
protects	The protected interfaces that are using this backup tunnel.

Related Commands

Command	Description
tunnel mpls traffic-eng backup-bw	Specifies what types of LSPs can use a backup tunnel, whether the backup tunnel should provide bandwidth protection, and if so, how much.

show mpls traffic-eng fast-reroute database

To display the contents of the Fast ReRoute database, use the **show mpls traffic-eng fast-reroute database** command in EXEC mode.

show mpls traffic-eng fast-reroute database
{*network* [*mask* | *masklength*]
| labels *low label* [*-high label*] **|**
interface *ifname* |
backup-interface *ifname*}]
[**state** {**active** | **ready** | **partial** | **complete** }]
[**role** {**head** | **middle**}]
[**detail**]

Syntax Description

network	IP address of the destination network. This functions as the prefix of the Fast ReRoute rewrite.
mask	Bit combination indicating the portion of the IP address that is being used for the subnet address.
masklength	Number of bits in the mask of the destination.
labels	Shows only database entries that possess in-labels assigned by this router (local labels). Specify either a starting value or a range of values.
low label	Starting label value or lowest value in the range.
- high label	Highest label value in the range.
interface *ifname*	Shows only database entries related to the specified primary outgoing interface.
backup-interface *ifname*	Shows only database entries related to the specified backup outgoing interface.
state	Shows entries that match one of four possible states: active, ready, partial, or complete.
active	The FRR rewrite has been put into the forwarding database (where it can be placed onto appropriate incoming packets).
ready	The FRR rewrite has been created, but has not yet been moved into the forwarding database.
partial	State before the FRR rewrite has been fully created; its backup routing information is still incomplete.
complete	State after the FRR rewrite has been assembled: ready or active.
role	Shows entries associated either with the tunnel head or tunnel midpoint.
head	Entry associated with the tunnel head.
middle	Entry associated with the tunnel midpoint.
detail	Shows long-form information (LFIB-FRR total number of clusters, groups, and items) in addition to the short-form information (prefix, label, and state).

Defaults

This command has no default behavior or values.

Command Modes

EXEC

Command History

Release	Modification
12.0(10)ST	This command was introduced.
12.0(22)S	This command is used for Node Protection. The command is supported on the Cisco 10000 series ESRs.
12.0(23)S	Output display reflects reduction in rewrites-per-prefix when LDP is not enabled.

Usage Guidelines

A **tunnel head end item** is created and added to the FRR database for each TE tunnel that is protected by Fast ReRoute. The existence of the head end item indicates that the label rewrite stored in the LFIB is protected, and that any IP prefix that uses the tunnel as its next hop will be FRR protected. You can confirm this information for each individual prefix by using the **show mpls forwarding-table detail** command.

A **prefix item** is created and added to the FRR database for each prefix that has label information associated with it via TDP/LDP and is being routed over a protected TE tunnel. Those items correspond to unique label rewrites, which are FRR protected and used by LFIB.

LSP midpoint items are created and added to the FRR database for LSPs that use the node as a transit, if their LSP output interface is protected by Fast ReRoute.

Examples

The following is sample output from the **show mpls traffic-eng fast-reroute database** command at a tunnel head link. (Prefix item and LSP midpoint information categories are empty in this first example because LDP has not been enabled. In the second example, shown after Table 10, LDP has been enabled).

```
Router# show mpls traffic-eng fast-reroute database

Tunnel head end item frr information:
Protected Tunnel   In-label   Out intf/label   FRR intf/label   Status
Tunnel10           Tun hd     PO5/0:Untagged   Tu0:12304        ready

Prefix item frr information:
Prefix     Tunnel     In-label   Out intf/label   FRR intf/label   Status

LSP midpoint frr information:
LSP identifier     In-label   Out intf/label   FRR intf/label   Status
```

Table 10 describes the fields displayed in this example.

Table 10 show mpls traffic-eng fast-reroute database Field Descriptions

Field	Description
Protected Tunnel	The tunnel's identifying number.
In-label	Label advertised to other routers to signify a particular prefix. The value "Tun hd" indicates that no label has been advertised.
Out intf/label	Out interface—Short name of the physical interface through which traffic goes to the protected link. Out label: • At a tunnel head, this is the label that the tunnel destination device advertises. The value "Untagged" indicates that no such label has been advertised. • At a tunnel midpoint, this is the label selected by the next-hop device. The value "Pop Tag" indicates that the next hop is the tunnel's final hop.
FRR intf/label	Fast Reroute interface—The backup tunnel interface. Fast Reroute label: • At a tunnel head, this is the label that the tunnel tail selected to indicate the destination network. The value "Untagged" indicates that no label has been advertised. • At a tunnel midpoint, this has the same value as the Out label.
Status	State of the rewrite: partial, ready, or active.
Prefix	Address to which packets with this label are going.
LSP identifier	LSP's identifying number.

The following is sample output from the **show mpls traffic-eng fast-reroute database** command when LDP has been enabled.

```
Router# show mpls traffic-eng fast-reroute database

Tunnel head end item frr information:
Protected Tunnel   In-label   Out intf/label   FRR intf/label   Status
Tunnel10           Tun hd     PO5/0:Untagged   Tu0:14814        ready

Prefix item frr information:
Prefix         Tunnel     In-label   Out intf/label   FRR intf/label   Status
10.0.7.1/32    Tu10       12317      PO5/0:Untagged   Tu0:14814        ready
10.0.4.1/32    Tu10       12318      PO5/0:Untagged   Tu0:14814        ready
10.0.0.52/30   Tu10       12319      PO5/0:Untagged   Tu0:14814        ready
10.0.0.48/30   Tu10       12320      PO5/0:Untagged   Tu0:14814        ready
10.0.0.36/30   Tu10       12321      PO5/0:Untagged   Tu0:14814        ready
10.0.0.32/30   Tu10       12322      PO5/0:Untagged   Tu0:14814        ready
10.0.0.28/30   Tu10       12323      PO5/0:Untagged   Tu0:14814        ready

LSP midpoint frr information:
LSP identifier     In-label   Out intf/label   FRR intf/label   Status
```

The following example shows output, at a midpoint link, from the **show mpls traffic-eng fast-reroute database** command with the **labels** argument specified:

```
Router# show mpls traffic-eng fast-reroute database labels 250 - 255

Tunnel head end item frr information:
Protected Tunnel   In-label   Out intf/label   FRR intf/label   Status

LSP midpoint frr information:
LSP identifier          In-label  Out intf/label   FRR intf/label   Status
10.110.0.10 229 [7334]  255       PO0/0:694        Tu4000:694       active
10.110.0.10 228 [7332]  254       PO0/0:693        Tu4000:693       active
10.110.0.10 227 [7331]  253       PO0/0:692        Tu4000:692       active
10.110.0.10 226 [7334]  252       PO0/0:691        Tu4000:691       active
10.110.0.10 225 [7333]  251       PO0/0:690        Tu4000:690       active
10.110.0.10 224 [7329]  250       PO0/0:689        Tu4000:689       active
```

The following example shows output, at a tunnel head link, from the **show mpls traffic-eng fast-reroute database** command with the **detail** argument specified:

```
Router# show mpls traffic-eng fast-reroute database 12.0.0.0. detail

LFIB FRR Database Summary:
  Total Clusters:      2
  Total Groups:        2
  Total Items:         789
Link 10:PO5/0 (Down, 1 group)
  Group 51:PO5/0->Tu4000 (Up, 779 members)
    Prefix 12.0.0.0/16, Tu313, active
      Input label Tun hd, Output label PO0/0:773, FRR label Tu4000:773
    Prefix 12.0.0.0/16, Tu392, active
      Input label Tun hd, Output label PO0/0:775, FRR label Tu4000:775
    Prefix 12.0.0.0/16, Tu111, active
      Input label Tun hd, Output label PO0/0:16, FRR label Tu4000:16
    Prefix 12.0.0.0/16, Tu394, active
      Input label Tun hd, Output label PO0/0:774, FRR label Tu4000:774
```

Table 11 describes the fields displayed in this example.

Table 11 show mpls traffic-eng fast-reroute database with Detail Keyword Field Descriptions

Field	Description
Total Clusters	A cluster is the physical interface upon which Fast ReRoute Link Protection has been enabled.
Total Groups	A group is a database record that associates the link-protected physical interface with a backup tunnel. A cluster (physical interface) therefore can have one or more groups. For example, the cluster Ethernet4/0/1 is protected by backup Tunnel1 and backup Tunnel2, so it has two groups.
Total Items	An item is a database record that associates a rewrite with a group. A group therefore can have one or more items.

Table 11 show mpls traffic-eng fast-reroute database with Detail Keyword Field Descriptions

Field	Description
Link 10:PO5/0 (Down, 1 group)	Describes a cluster (physical interface): • "10" is the interface's unique IOS-assigned ID number. • ":" is followed by the interface's short name. • Parentheses contain the operating state of the interface (Up or Down) and the number of groups associated with it.
Group 51:PO5/0->Tu4000 (Up, 779 members)	Describes a group: • "51" is the ID number of the backup interface. • ":" is followed by the group's physical interface short name. • "->" is followed by the backup tunnel interface short name. • Parentheses contain the operating state of the tunnel interface (Up or Down) and the number of items—also called "members"—associated with it.

Related Commands

Command	Description
show mpls traffic-eng fast-reroute log reroutes	Displays contents of the Fast ReRoute event log.

show mpls traffic-eng tunnels

To show information about tunnels, use the **show mpls traffic-eng tunnels** command in EXEC mode.

show mpls traffic-eng tunnels
[**tunnel** *unit*]
[**destination** *address*]
[**source-id** {*num* | *ipaddress* | *ipaddress num*}]
[**role** {**all** | **head** | **middle** | **tail** | **remote**}]
[{**up** | **down**}]
[**name** *string*]
[**suboptimal constraints** {**none** | **current** | **max**}]
[{[**interface in** *phys_intf*] [**interface out** *phys_intf*] | [**interface** *phys_intf*]}]
[**property** {**backup** | **fast-reroute**}]
[**brief** | **backup** | **protection**]

Syntax Description

tunnel *unit*	(Optional) Displays information for the specified tunneling interface.
destination *address*	(Optional) Restricts the display to tunnels destined to the specified IP address.
source-id	(Optional) Restricts the display to tunnels with a matching source IP address and/or tunnel number.
num	(Optional) Tunnel number.
ipaddress	(Optional) Source IP address.
ipaddress num	(Optional) Source IP address and tunnel number.
role	(Optional) Restricts the display to tunnels with the indicated role (all, head, middle, tail, or remote).
all	(Optional) Displays all tunnels.
head	(Optional) Displays tunnels with their head at this router.
middle	(Optional) Displays tunnels with a midpoint at this router.
tail	(Optional) Displays tunnels with a tail at this router.
remote	(Optional) Displays tunnels with their head at some other router; this is a combination of **middle** and **tail**.
up	(Optional) Displays tunnels if the tunnel interface is up. Tunnel midpoints and tails are typically up or not present.
down	(Optional) Displays tunnels that are down.
name *string*	(Optional) Displays tunnel with the specified name. The tunnel name is derived from the interface description, if specified; otherwise, it is the interface name. The tunnel name is included in the signalling message so it is available at all hops.
suboptimal constraints none	(Optional) Displays tunnels whose path metric is greater than the shortest unconstrained path. Selected tunnels have a longer path than the IGP's shortest path.
suboptimal constraints current	(Optional) Displays tunnels whose path metric is greater than the current shortest path, constrained by the tunnel's configured options. Selected tunnels would have a shorter path if they were reoptimized immediately.

suboptimal constraints max	Displays tunnels whose path metric is greater than the current shortest path, constrained by the tunnel's configured options, and considering only the network's capacity. Selected tunnels would have a shorter path if no other tunnels were consuming network resources.
interface in *phys_intf*	Displays tunnels that use the specified input interface.
interface out *phys_intf*	Displays tunnels that use the specified output interface.
interface *phys_intf*	Displays tunnels that use the specified interface as an input or output interface.
property backup	Selects MPLS TE tunnels being used to protect physical interfaces on this router. A tunnel configured to protect a link against failure is a backup tunnel and has the backup tunnel property.
property fast-reroute	Selects Fast ReRoute-protected MPLS TE tunnels originating on, transmitting, or terminating on this router.
brief	Specifies a format with one line per tunnel.
backup	Displays information about the Fast ReRoute protection provided by each tunnel selected by other options specified with this command. The information includes the physical interface protected by the tunnel, the number of TE LSPs (tunnels) protected, and the bandwidth protected.
protection	Displays information about the protection provided to each tunnel selected by other options specified with this command. The information includes whether protection is configured for the tunnel, the protection (if any) provided to the tunnel by this router, and the tunnel bandwidth protected.

Defaults

If you specify this command without any arguments or keywords, the command displays general information about each MPLS TE tunnel known to the router.

Command Modes

EXEC

Command History

Release	Modification
12.0(5)S	This command was introduced.
12.0(10)ST	The new **brief** format includes input and output interface information. The **suboptimal** and **interface** keywords were added to the non-brief format. The non-brief, non-summary formats contain the history of LSP selection.
12.0(22)S	The **property**, **back-up**, **fast-reroute**, and **protection** keywords were added. The command is supported on the Cisco 10000 series ESRs.

Usage Guidelines

To select the tunnels for which information is displayed, use the **tunnel**, **destination**, **source-id**, **role**, **up**, **down**, **name**, **suboptimal**, **interface** and **property** keywords and options singly or combined.

To select the type of information displayed about the selected tunnels, use the **brief**, **accounting**, **backup**, or **protection** keywords.

Examples

The following is sample output from the **show mpls traffic-eng tunnel brief** command. It displays brief information about every MPLS TE tunnel known to the router.

```
Router# show mpls traffic-eng tunnels brief

Signalling Summary:
    LSP Tunnels Process:            running
    RSVP Process:                   running
    Forwarding:                     enabled
    Periodic reoptimization:        every 3600 seconds, next in 1706 seconds
TUNNEL NAME                      DESTINATION      UP IF     DOWN IF   STATE/PROT
Router_t1                        10.112.0.12      -         PO4/0/1   up/up
Router_t2                        10.112.0.12      -         unknown   up/down
Router_t3                        10.112.0.12      -         unknown   admin-down
Router_t1000                     10.110.0.10      -         unknown   up/down
Router_t2000                     10.110.0.10      -         PO4/0/1   up/up
Displayed 5 (of 5) heads, 0 (of 0) midpoints, 0 (of 0) tails
```

Table 12 describes the fields displayed in this example.

Table 12 show mpls traffic-eng tunnels Field Descriptions

Field	Description
LSP Tunnels Process	Status of the LSP tunnels process.
RSVP Process	RSVP has or has not been enabled. (This feature is enabled as a consequence of MPLS Traffic Engineering being enabled.)
Forwarding	Status of forwarding (enabled or disabled).
Periodic Reoptimization	Schedule for periodic reoptimization.
TUNNEL NAME	Name of the interface that is configured at the tunnel head.
DESTINATION	Identifier of the tail-end router.
Head	Summary information about tunnel heads at this device.
UP IF	Upstream interface that the tunnel used.
DOWN IF	Downstream interface that the tunnel used.
STATE/PROT	For tunnel heads, admin-down or up. For non-heads, signalled.

The following is sample output from the **show mpls traffic-eng tunnels property backup-tunnel brief** command. It displays brief information about all MPLS TE tunnels acting as Fast Reroute backup tunnels (**property backup**) for interfaces on the router.

```
Router# show mpls traffic-eng tunnels property backup brief

Signalling Summary:
    LSP Tunnels Process:            running
    RSVP Process:                   running
    Forwarding:                     enabled
    Periodic reoptimization:        every 3600 seconds, next in 2231 seconds
    Periodic FRR Promotion:         every 300 seconds, next in 131 seconds
    Periodic auto-bw collection:    disabled
TUNNEL NAME                      DESTINATION      UP IF     DOWN IF   STATE/PROT
Router_t578                      88.88.88.88      -         PO1/0     up/up
```

```
Router_t5710                      7.7.7.7          -          unknown   admin-down
Router_t5711                      7.7.7.7          -          PO1/1     up/up
Displayed 3 (of 9) heads, 0 (of 1) midpoints, 0 (of 0) tails
```

The following is sample output from the **show mpls traffic-eng tunnels backup** command. This command selects every MPLS TE tunnel known to the router and displays information about the Fast ReRoute protection each selected tunnels provides for interfaces on this router; the command does not generate output for tunnels that do not provide Fast ReRoute protection of interfaces on this router.

```
Router# show mpls traffic-eng tunnels backup

Router_t578
  LSP Head, Tunnel578, Admin: up, Oper: up
  Src 55.55.55.55, Dest 88.88.88.88, Instance 1
  Fast Reroute Backup Provided:
    Protected i/fs: PO1/0, PO1/1, PO3/3
    Protected lsps: 1
    Backup BW: any pool unlimited; inuse: 100 kbps
Router_t5710
  LSP Head, Tunnel5710, Admin: admin-down, Oper: down
  Src 55.55.55.55, Dest 7.7.7.7, Instance 0
  Fast Reroute Backup Provided:
    Protected i/fs: PO1/1
    Protected lsps: 0
    Backup BW: any pool unlimited; inuse: 0 kbps
Router_t5711
  LSP Head, Tunnel5711, Admin: up, Oper: up
  Src 55.55.55.55, Dest 7.7.7.7, Instance 1
  Fast Reroute Backup Provided:
    Protected i/fs: PO1/0
    Protected lsps: 2
    Backup BW: any pool unlimited; inuse: 6010 kbps
```

The following is sample output from the **show mpls traffic-eng tunnels property fast-reroute protection** command. This command selects every MPLS TE tunnel known to the router that was signaled as a Fast ReRoute protected LSP (**property fast-reroute**) and displays information about the protection this router provides each selected tunnel.

```
Router# show mpls traffic-eng tunnels property fast-reroute protection

Router_t1
  LSP Head, Tunnel1, Admin: up, Oper: up
  Src 55.55.55.55, Dest 88.88.88.88, Instance 25
  Fast Reroute Protection: Requested
    Outbound: FRR Ready
      Backup Tu5711 to LSP nhop
        Tu5711: out i/f: PO1/1, label: implicit-null
      LSP signalling info:
        Original: out i/f: PO1/0, label: 12304, nhop: 10.1.1.7
        With FRR: out i/f: Tu5711, label: 12304
      LSP bw: 6000 kbps, Backup level: any unlimited, type: any pool
Router_t2
  LSP Head, Tunnel2, Admin: up, Oper: up
  Src 55.55.55.55, Dest 88.88.88.88, Instance 2
  Fast Reroute Protection: Requested
    Outbound: FRR Ready
      Backup Tu578 to LSP nhop
        Tu578: out i/f: PO1/0, label: 12306
      LSP signalling info:
        Original: out i/f: PO3/3, label: implicit-null, nhop: 10.3.3.8
        With FRR: out i/f: Tu578, label: implicit-null
      LSP bw: 100 kbps, Backup level: any unlimited, type: any pool
```

```
r9_t1
  LSP Midpoint, signalled, connection up
  Src 9.9.9.9, Dest 88.88.88.88, Instance 2347
  Fast Reroute Protection: Requested
    Inbound:  FRR Inactive
      LSP signalling info:
        Original: in i/f: PO1/2, label: 12304, phop: 10.205.0.9
    Outbound: FRR Ready
      Backup Tu5711 to LSP nhop
        Tu5711: out i/f: PO1/1, label: implicit-null
      LSP signalling info:
        Original: out i/f: PO1/0, label: 12305, nhop: 10.1.1.7
        With FRR: out i/f: Tu5711, label: 12305
      LSP bw: 10 kbps, Backup level: any unlimited, type: any pool
```

show mpls traffic-eng tunnels summary

To show summary information about tunnels, use the **show mpls traffic-eng tunnels summary** command in EXEC mode.

show mpls traffic-eng tunnels summary

Syntax Description

This command has no arguments or keywords.

Defaults

This command has no default behavior or values.

Command Modes

EXEC

Command History

Release	Modification
12.0(5)S	This command was introduced.
12.0(22)S	The command output changed. The Periodic fastreroute line has been added. The command is supported on the Cisco 10000 series ESRs.

Examples

The following is sample output from the **show mpls traffic-eng tunnels summary** command.

Note The only change is that the Periodic fastreroute line has been added.

```
Router# show mpls traffic-eng tunnels summary

Signalling Summary:
  LSP Tunnels Process:         running
  RSVP Process:                running
  Forwarding:                  enabled
  Head: 4 interfaces, 3 active signalling attempts, 3 established
     5 activations, 2 deactivations
  Midpoints: 1, Tails: 0
  Periodic reoptimization:     every 3600 seconds, next in 2778 seconds
  Periodic fastreroute:        every 300 seconds, next in 168 seconds
  Periodic auto-bw collection: every 300 seconds, next in 78 seconds
```

Table 13 describes the fields displayed in this example.

Table 13 show mpls traffic-eng tunnels summary Field Descriptions

Field	Description
LSP Tunnels Process	MPLS Traffic Engineering has or has not been enabled.
RSVP Process	RSVP has or has not been enabled. (This feature is enabled as a consequence of MPLS Traffic Engineering being enabled.)

Table 13 show mpls traffic-eng tunnels summary Field Descriptions (continued)

Field	Description
Forwarding	Indicates whether appropriate forwarding is enabled. (Appropriate forwarding on a router is CEF switching.)
Head	Summary information about tunnel heads at this device.
Interfaces	Number of MPLS Traffic Engineering tunnel interfaces.
Active signalling attempts	LSPs currently successfully signaled or in the process of being signaled.
Established	LSPs currently signaled.
Activations	Signaling attempts initiated.
Deactivations	Signaling attempts terminated.
Periodic reoptimization	Frequency of periodic reoptimization and time until the next periodic reoptimization.
Periodic fastreroute	Frequency that scanning occurs to determine if LSPs should be promoted to better backup tunnels, and time until the next scanning.
Periodic auto-bw collection	Frequency of automatic bandwidth collection and time left until the next collection.

Related Commands

Command	Description
mpls traffic-eng fast-reroute timers	Specifies how often the router considers switching an LSP to a new (better) backup tunnel if additional backup-bandwidth becomes available.

tunnel mpls traffic-eng backup-bw

To specify what types of LSPs can use a backup tunnel, whether the backup tunnel should provide bandwidth protection, and if so, how much, use the **tunnel mpls traffic-eng backup-bw** command in interface configuration mode.

tunnel mpls traffic-eng backup-bw {*bandwidth* | [**sub-pool** {*bandwidth* | **Unlimited**}] [**global-pool** {*bandwidth* | **Unlimited**}]}

Syntax Description

global-pool	Only LSPs using bandwidth from the global-pool can use the backup tunnel.
sub-pool	Only LSPs using bandwidth from the sub-pool can use the backup tunnel.
bandwidth	The amount of bandwidth this backup tunnel can protect. The router limits the LSPs that can use this backup tunnel so that the sum of the bandwidth of the LSPs does not exceed the specified amount of bandwidth. If there are multiple backup tunnels, the router will use the best-fit algorithm (see "Backup Tunnel Selection Procedure") to determine which backup tunnel to use.
Unlimited	Backup tunnel does not provide bandwidth protection. Any number of LSPs can use the backup tunnel, regardless of their bandwidth.

Defaults

If neither **sub-pool** nor **global-pool** is entered, it is assumed that any LSP (those using bandwidth from the sub-pool or global-pool) can use this backup tunnel.

Command Modes

Interface configuration mode, on the backup tunnel interface

Command History

Release	Modification
12.0(22)S	This command was introduced.

Usage Guidelines

If both **sub-pool** and **global-pool** are specified, **sub-pool** must be specified first on the command line. For example, **tunnel mpls traffic-eng backup-bw sub-pool 100 global-pool Unlimited** is legal, but it is not legal to specify **tunnel mpls traffic-eng backup-bw global-pool Unlimited sub-pool 100**.

To limit both sub-pool and global pool LSPs, enter **tunnel mpls traffic-eng backup-bw sub-pool** *bandwidth* **global-pool** *bandwidth*.

If sub-pool is **Unlimited**, global-pool cannot also be **Unlimited**. Entering such a command (**tunnel mpls traffic-eng backup-bw sub-pool Unlimited global-pool Unlimited**) would be the same as entering nothing at all because it is the default behavior.

Examples

In the following example, backup tunnel 1 is to be used only by LSPs that take their bandwidth from the global pool. The backup tunnel does not provide bandwidth protection. Backup tunnel 2 is to be used only by LSPs that take their bandwidth from the sub-pool. Backup tunnel 2 provides bandwidth protection for up to 1000 units.

```
Router(config)# interface Tunnel1
Router(config-if)# tunnel mpls traffic-eng backup-bw global-pool Unlimited
Router(config-if)# end

Router(config)# interface Tunnel2
Router(config-if)# tunnel mpls traffic-eng backup-bw sub-pool 1000
Router(config-if)# end
```

Related Commands

Command	Description
mpls traffic-eng backup-path	Assigns one or more backup tunnels to a protected interface.

Bandwidth Protection Considerations

There are multiple methods to ensure bandwidth protection. Cisco implementation of FRR does not mandate a particular bandwidth protection method. However, it is important that the method you choose is consistent with the bandwidth protection strategy you choose.

The following sections describe some important issues to consider when choosing an appropriate configuration:

- Using Backup Tunnels with Explicitly Signaled Bandwidth, page 83
- Using Backup Tunnels Signaled with Zero Bandwidth, page 84

Using Backup Tunnels with Explicitly Signaled Bandwidth

When using the explicitly signaled bandwidth method, you must configure the following two bandwidth parameters for a backup tunnel:

- Signaled bandwidth
- Backup-bandwidth

The *signaled bandwidth* is used by the LSRs on the path of the backup tunnel to perform admission control and bandwidth accounting.

The *backup-bandwidth* is used by the PLR (the head-end of the backup tunnel) to decide how much primary traffic can use this backup tunnel if there is a failure.

You must configure both parameters, and the values of *signaled bandwidth* and *backup-bandwidth* must be the same.

To configure *signaled bandwidth*, use the **tunnel mpls traffic-eng bandwidth** command.

To configure *backup-bandwidth*, use the **tunnel mpls traffic-eng backup-bw** command.

Configuring *signaled bandwidth* allows you to specify both of the following:

- *Amount* of bandwidth a backup tunnel reserves
- *Pool (*global pool or sub-pool*)* from which the backup tunnel reserves its bandwidth

Note Only one pool can be selected (that is, the backup tunnel can explicitly reserve bandwidth from the global pool or the sub-pool, but not both).

Configuring *backup-bandwidth* allows you to specify the pool to which the traffic must belong in order to use this backup tunnel. Multiple pools are allowed.

You can configure *different* pools for signaled bandwidth and backup-bandwidth of the same backup tunnel.

Example—Assume the following:

- Bandwidth protection is desired only for sub-pool traffic. Best-effort traffic uses the global pool and does not require bandwidth protection.
- Scheduling is configured so that sub-pool traffic uses the priority queue, and global pool traffic is served at a lower priority.

Bandwidth protection for 10 Kbps of sub-pool traffic on a given link can be achieved by any of the following command combinations:

- **tunnel mpls traffic-eng bandwidth sub-pool 10**

  tunnel mpls traffic-eng backup-bw sub-pool 10
- **tunnel mpls traffic-eng bandwidth global-pool 10**

  tunnel mpls traffic-eng backup-bw sub-pool 10
- **tunnel mpls traffic-eng bandwidth global-pool 10**

  tunnel mpls traffic-eng backup-bw sub-pool 10 global-pool unlimited

In all of the above combinations, bandwidth is explicitly signaled on the backup tunnel.

In the first combination, the bandwidth to be used by sub-pool traffic in case of a failure has been explicitly reserved from the sub-pool. Primary sub-pool traffic cannot use this bandwidth (we have reduced the amount of sub-pool bandwidth available to primary sub-pool traffic). This bandwidth is reserved for use only by sub-pool traffic that has been rerouted onto a backup tunnel when a failure occurs.

The second combination is similar to the first, except that the bandwidth is reserved from the global-pool instead of the sub-pool. The bandwidth available to primary global-pool traffic (rather than primary sub-pool traffic) is reduced. The benefit is that the size of the sub-pool for primary traffic is not affected.

The third combination is similar to the second, except that global-pool traffic can use the backup tunnel if there is a failure. However, no bandwidth protection is given to global-pool traffic during a failure. Subpool traffic is served at a higher priority; therefore, it is protected even though an unlimited amount of global-pool traffic is allowed on the same backup tunnel.

Using Backup Tunnels Signaled with Zero Bandwidth

Frequently it is desirable to use backup tunnels with zero *signaled bandwidth*, even when bandwidth protection is required. It may seem that if no bandwidth is explicitly reserved, no bandwidth guarantees can be provided. However, that is not necessarily true.

Assume the following:

- Only link protection is desired.
- Bandwidth protection is desired only for sub-pool traffic.
- A protected link between nodes A and B has a max reservable sub-pool bandwidth of S.
- There is some path between A and B such that for each link along this path, the max reservable global-pool bandwidth of this link minus S is at least S.

If it is possible to find such a path, then it is possible to provide bandwidth protection for this link using a backup tunnel signaled with zero bandwidth.

For example, assume that the max reservable sub-pool bandwidth between A and B is 5. The objective is to protect 5 units of sub-pool bandwidth in case the link between A and B fails. So, an alternate path must be found between A and B which has at least 5 units of bandwidth available on each link along the path. On each link of this alternate path, protection can be provided by using 5 units of global-pool bandwidth without compromising primary sub-pool traffic. This is possible only if, for each link along the alternate path, there are 5 units of global-pool bandwidth available. For example, assume the following:

- The max reservable global-pool bandwidth on a link is 10.
- The max reservable sub-pool bandwidth is 5 or less.

Then there is at least 5 units of global-pool bandwidth available.

In general, if the max reservable global-pool bandwidth on a link is G, and the max reservable sub-pool bandwidth is S, then there is G minus S global-pool bandwidth available. If G minus S is at least S, then the global pool can be used to provide backup protection for 5 units of sub-pool traffic. In general, if we can find an alternate path which can be used in case of a link failure which has sufficient bandwidth to hold the sub-pool traffic using this link, then we can signal the backup tunnel (using zero signaled bandwidth) along this path and it will be able to provide bandwidth protection for the sub-pool traffic in case of a failure.

If it is possible to find such a path for all links in the network, then bandwidth protection can be provided for every link in the network using backup tunnels signaled with zero bandwidth.

The above approach allows sharing of the global pool bandwidth between backup tunnels protecting independent link failures. The backup tunnels are expected to be used for only a short period of time after a failure (until the head-ends of affected LSPs reroute those LSPs to other paths with available sub-pool bandwidth). The probability of multiple unrelated link failures is very small (in the absence of node or SRLG failures, which result in multiple link failures). Therefore, it is reasonable to assume that link failures are in practice independent with high probability. This “independent failure assumption” in combination with backup tunnels signaled without explicit bandwidth reservation enables efficient bandwidth sharing that yields substantial bandwidth savings.

Backup tunnels protecting the sub-pool traffic do not draw bandwidth from any pool. Primary traffic using the global pool can use the entire global pool, and primary traffic using the sub-pool can use the entire sub-pool. Yet, sub-pool traffic has a complete bandwidth guarantee if there is a single link failure.

A similar approach can be used for node and SRLG protection. However, the decision of where to put the backup tunnels is more complicated because both node and SRLG failures effectively result in the simultaneous failure of several links. Therefore, the backup tunnels protecting traffic traversing all affected links cannot be computed independently of each other. The backup tunnels protecting groups of links corresponding to different failures can still be computed independently of each other, which results in similar bandwidth savings.

Signaled Bandwidth versus Backup-Bandwidth

Backup-bandwidth is used locally (by the router that is the head-end of the backup tunnel) to determine which, and how many, primary LSPs can be rerouted on a particular backup tunnel. The router ensures that the combined bandwidth requirement of these LSPs does not exceed the *backup-bandwidth*.

Therefore, even when the backup tunnel uses zero *signaled bandwidth*, the *backup-bandwidth* must be configured with the value corresponding to the bandwidth requirement of the traffic protected by this backup tunnel. Unlike the explicitly signaled case, the value of the *signaled bandwidth* (which is zero) is not the same value as the *backup-bandwidth*.

Glossary

backup tunnel—An MPLS Traffic Engineering tunnel used to protect other (primary) tunnels' traffic when a link or node failure occurs.

bandwidth—The available traffic capacity of a link.

CEF—Cisco Express Forwarding. A means for accelerating the forwarding of packets within a router, by storing route lookup.

Cisco Express Forwarding—See CEF.

differentiated services code point—See DSCP.

DSCP—Differentiated Services Code Point. Six bits in the IP header, as defined by the IETF. These bits determine the class of service provided to the IP packet.

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

Fast ReRoute—Procedures that enable temporary routing around a failed link or node while a new LSP is being established at the head end.

flooding—Traffic passing techniques used by switches and bridges in which traffic received on an interface is sent out all the interfaces of that device except the interface on which the information was received originally.

global pool—The total bandwidth allocated to an MPLS Traffic Engineering link or node.

headend—The router that originates and maintains a given LSP. This is the first router in the LSP's path.

hop—Passage of a data packet between two network nodes (for example, between two routers).

instance—A Hello instance implements the RSVP Hello extensions for a given router interface address and remote IP address. Active Hello instances periodically send Hello Request messages, expecting Hello ACK messages in response. If the expected Ack message is not received, the active Hello instance declares that the neighbor (remote IP address) is unreachable (that is, it is lost). This can cause LSPs crossing this neighbor to be fast rerouted.

interface—A network connection.

intermediate nodes—Intermediate System-to-Intermediate System. Link-state hierarchical routing protocol that calls for intermediate system (IS) routers to exchange routing information based on a single metric to determine network topology.

link—A point-to-point connection between adjacent nodes. There can be more than one link between adjacent nodes. A network communications channel consisting of a circuit or transmission path and all related equipment between a sender and a receiver. Sometimes referred to as a line or a transmission link.

label-switched path—See LSP.

limited backup-bandwidth—Backup tunnels that provide bandwidth protection.

load balancing—A configuration technique that shifts traffic to an alternative link if a certain threshold is exceeded on the primary link. Load balancing is similar to redundancy in that if an event causes traffic to shift directions, alternative equipment must be present in the configuration. In load balancing, the alternative equipment is not necessarily redundant equipment that only operates in the event of a failure.

LSP—Label-switched path. A configured connection between two routers, in which label switching is used to carry the packets. The purpose of an LSP is to carry data packets.

merge point—The backup tunnel's tail.

MPLS—Multiprotocol Label Switching. Packet-forwarding technology, used in the network core, that applies data link layer labels to tell switching nodes how to forward data, resulting in faster and more scalable forwarding than network layer routing normally can do.

MPLS global label allocation—There is one label space for all interfaces in the router. For example, label 100 coming in one interface is treated the same as label 100 coming in a different interface.

next hop—The next downstream node along an LSP's path. Also called NHOP.

next-hop backup tunnel—See NHOP backup tunnel.

next-next hop—The node after the next downstream node along an LSP's path. Also called NNHOP.

next-next-hop backup tunnel—See NNHOP backup tunnel.

NHOP—Next hop. The next downstream node along an LSP's path.

NHOP backup tunnel—Next-hop backup tunnel. Backup tunnel terminating at the LSP's next hop beyond the point of failure, and originating at the hop immediately upstream of the point of failure. It bypasses a failed link, and is used to protect primary LSPs that were using this link before the failure.

NNHOP—Next-next hop. The node after the next downstream node along an LSP's path.

NNHOP backup tunnel—Next-next-hop backup tunnel. Backup tunnel terminating at the LSP's next-next hop beyond the point of failure, and originating at the hop immediately upstream of the point of failure. It bypasses a failed link and/or node, and is used to protect primary LSPs that were using this link or node before the failure.

node—Endpoint of a network connection or a junction common to two or more lines in a network. Nodes can be interconnected by links, and serve as control points in the network. Computers on a network, or any endpoint or a junction common to two or more lines in a network. Nodes can be processors, controllers, or workstations.

OSPF—Open Shortest Path First. A link-state hierarchical Interior Gateway Protocol routing algorithm, derived from the IS-IS protocol. OSPF features include least-cost routing, multipath routing, and load balancing.

primary LSP—The last LSP originally signaled over the protected interface before the failure. The LSP before the failure.

primary tunnel—Tunnel whose LSP may be fast rerouted if there is a failure. Backup tunnels cannot be primary tunnels.

promotion—Conditions, such as a new backup tunnel comes up, cause a reevaluation of a backup tunnel that was chosen for an LSP. If the reevaluation is successful, it is called a promotion.

protected interface—An interface that has one or more backup tunnels associated with it.

redundancy—The duplication of devices, services, or connections so that, in the event of a failure, the redundant devices, services, or connections can perform the work of those that failed.

RSVP—Resource Reservation Protocol. An IETF protocol used for signaling requests (setting up reservations) for Internet services by a customer before that customer is permitted to transmit data over that portion of the network.

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

state—Information that a router must maintain about each LSP. The information is used for rerouting tunnels.

sub-pool—The more restrictive bandwidth in an MPLS Traffic Engineering link or node. The sub-pool is a portion of the link or node's overall global pool bandwidth.

tailend—The router upon which an LSP is terminated. This is the last router in the LSP's path.

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

tunnel—Secure communications path between two peers, such as two routers.

unlimited backup-bandwidth—Backup tunnels that provide no bandwidth (best-effort) protection (that is, they provide best-effort protection).