



MPLS Traffic Engineering Inter-AS TE

The MPLS Traffic Engineering: Inter-AS TE feature provides Autonomous System Boundary Router (ASBR) node protection, loose path reoptimization, stateful switchover (SSO) recovery of label-switched paths (LSPs) that include loose hops, ASBR forced link flooding, Cisco IOS Resource Reservation Protocol (RSVP) local policy extensions for interautonomous system (Inter-AS), and per-neighbor keys:

- ASBR node protection--Protects interarea and Inter-AS TE label-switched paths (LSPs) from the failure of an Area Border Router (ABR) or ASBR.
 - Loose path reoptimization--Allows a Multiprotocol Label Switching (MPLS) traffic engineering (TE) tunnel's LSPs to traverse hops that are not in the tunnel headend router's topology database (that is, they are not in the same Open Shortest Path First (OSPF) area, Intermediate System-to-Intermediate System (IS-IS) level, or autonomous system as the tunnel's headend router).
 - Loose hop recovery--Supports SSO recovery of LSPs that include loose hops.
 - ASBR forced link flooding--Helps an LSP cross a boundary into another domain when information in the other domain is not available to the headend router.
 - Cisco IOS RSVP local policy extensions for Inter-AS--Allows network administrators to create controlled policies for TE tunnels that function across multiple autonomous systems.
 - Per-neighbor keys--Allows cryptographic authentication to be accomplished on a per-neighbor basis.
-
- [Finding Feature Information, page 2](#)
 - [Prerequisites for MPLS Traffic Engineering Inter-AS TE, page 2](#)
 - [Restrictions for MPLS Traffic Engineering Inter-AS TE, page 2](#)
 - [Information About MPLS Traffic Engineering Inter-AS TE, page 3](#)
 - [How to Configure MPLS Traffic Engineering Inter-AS TE, page 12](#)
 - [Configuration Examples for MPLS Traffic Engineering Inter-AS TE, page 22](#)
 - [Additional References, page 25](#)
 - [Feature Information for MPLS Traffic Engineering Inter-AS TE, page 26](#)
 - [Glossary, page 27](#)

Finding Feature Information

Your software release may not support all the features documented in this module. For the latest caveats and feature information, see [Bug Search Tool](#) and the release notes for your platform and software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the feature information table.

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

Prerequisites for MPLS Traffic Engineering Inter-AS TE

- Enable MPLS.
- Configure TE on routers.
- Ensure that your network supports the following Cisco features:
 - MPLS
 - Cisco Express Forwarding
 - IS-IS or OSPF
- For loose path reoptimization, know how to configure the following:
 - IP explicit paths for MPLS TE tunnels
 - Loose hops
 - Interarea and Inter-AS tunnels

Restrictions for MPLS Traffic Engineering Inter-AS TE

Loose Path Reoptimization

- Midpoint reoptimization is not supported.

ASBR Forced Link Flooding

- The TE metric and affinity attributes that are known at a headend router (and used as constraints when an LSP's path is computed) are not currently signaled. Consequently, explicit router (ERO) expansions do not consider these constraints.
- Each node in an autonomous system must have a unique router ID.
- The router ID configured on a link must not conflict with the router ID within the autonomous system.
- If a link is configured for forced link flooding, the link's neighbors are not learned by regular Interior Gateway Protocol (IGP) updates. If a link is already learned about neighbors by IGP on a link, you

cannot configure the link as passive. Therefore, to configure a link for forced flooding, be sure that the node does not already have a neighbor on that link.

Information About MPLS Traffic Engineering Inter-AS TE

MPLS Traffic Engineering Tunnels

MPLS TE lets you build LSPs across your network that you then forward traffic down.

MPLS TE LSPs, also called TE tunnels, let the headend of a TE tunnel control the path its traffic takes to a particular destination. This method is more flexible than forwarding traffic based only on a destination address.

Interarea tunnels allow you to do the following:

- Build TE tunnels between areas (interarea tunnels)
- Build TE tunnels that start and end in the same area, on multiple areas on a router (intra-area tunnels)

Some tunnels are more important than others. For example, you may have tunnels carrying Voice over IP (VoIP) traffic and tunnels carrying data traffic that are competing for the same resources. Or you may simply have some data tunnels that are more important than others. MPLS TE allows you to have some tunnels preempt others. Each tunnel has a priority, and more-important tunnels take precedence over less-important tunnels.

Multiarea Network Design

You can establish MPLS TE tunnels that span multiple IGP areas and levels. The tunnel headend routers and tailend routers do not have to be in the same area. The IGP can be either IS-IS or OSPF.

To configure an interarea tunnel, use the **next-address loose** command to specify on the headend router a loosely routed explicit path of the LSP that identifies each ABR the LSP should traverse. The headend router and the ABRs along the specified explicit path expand the loose hops, each computing the path segment to the next ABR or tunnel destination.

Fast Reroute

MPLS Fast Reroute (FRR) is a fast recovery local protection technique that protects TE LSPs from link, shared risk link group (SRLG), and node failure. One or more TE LSPs (called backup LSPs) are preestablished to protect against the failure of a link, node, or SRLG. If there is a failure, each protected TE LSP traversing the failed resource is rerouted onto the appropriate backup tunnels.

The backup tunnel must meet the following requirements:

- It should not pass through the element it protects.
- It should intersect with a primary tunnel at a minimum of two nodes: point of local repair (PLR) and merge point (MP). The PLR should be the headend LSR of the backup tunnel, and the MP should be the tailend LSR of the backup tunnel. The PLR is where FRR is triggered when a link, node, or SRLG failure occurs.

- FRR protection can be performed for an Inter-AS tunnel only if the backup tunnel's merge point can route packets to the PLR's backup tunnel's egress interface. You can configure a static route or you can configure Border Gateway Protocol (BGP) to export the backup tunnel's egress interface to other autonomous systems.
- If the preferred link is a passive link, you must assign an administrative-weight for it. To assign an administrative weight, use the **mpls traffic-eng administrative-weight** command in interface configuration mode.
- Each router must be configured with the **mpls traffic-eng reoptimize events link-up** command in global configuration mode.

ASBR Node Protection

A TE LSP that traverses an ASBR needs a special protection mechanism (ASBR node protection) because the MP and PLR will be in different autonomous systems that have different IGP.

A PLR ensures that the backup tunnel intersects with the primary tunnel at the MP by examining the Record Route Object (RRO) of the primary tunnel to see if any addresses specified in the RRO match the destination of the backup tunnel.


Addresses specified in RRO IPv4 and IPv6 subobjects can be node-IDs and interface addresses. The traffic engineering RFC 3209 specifies that you can use a router address or interface address, but recommends using the interface address of outgoing path messages. Therefore, in the figure below router R2 is more likely to specify interface addresses in the RRO objects carried in the resv messages of the primary tunnel (T1) and the backup tunnel.

Node IDs allow the PLR to select a suitable backup tunnel by comparing node IDs in the resv RRO to the backup tunnel's destination.

RSVP messages that must be routed and forwarded to the appropriate peer (for example, an resv message) require a route from the MP back to the PLR for the RSVP messages to be delivered. The MP needs a route to the PLR backup tunnel's outgoing interface for the resv message to be delivered. Therefore, you must configure a static route from the MP to the PLR. For the configuration procedure, see the [Configuring a Static Route from the MP to the PLR](#), on page 14.

The figure below illustrates ASBR node protection. Router R4 is node-protected with a backup tunnel from R2-R3-R5-R6.

Figure 1: ASBR Node Protection



Notes:

There are two autonomous systems.
 The numbers within the Ethernet serial connection indicate the OSPF area number.
 There is no IGP between R2 and R4, and R3 and R5.

In this configuration, IP addresses are as follows:

- R1--Loopback0 10.10.0.1
 - Ethernet 0--IP address of 10.10.1.1 is connected to R2 Ethernet 0
 - Ethernet 1--IP address of 10.10.2.1 is connected to R3 Ethernet 1
- R2--Loopback0 10.10.0.2
 - Ethernet 0--IP address of 10.10.1.2 is connected to R1 Ethernet 0
 - Ethernet 1--IP address of 10.10.3.1 is connected to R3 Ethernet 1
 - Serial 2--IP address of 10.10.4.1 is connected to R4 serial 2
- R3--Loopback0 10.10.0.3
 - Ethernet 0--IP address of 10.10.2.2 is connected to R1 Ethernet 1
 - Ethernet 1--IP address of 10.10.3.2 is connected to R2 Ethernet 1

- Serial 2--IP address of 10.10.5.1 is connected to R5 serial 2
- R4--Loopback0 10.10.0.4
 - Ethernet 0--IP address of 10.10.7.1 is connected to R6 Ethernet 0
 - Ethernet 1--IP address of 10.10.6.1 is connected to R5 Ethernet 1
 - Serial 2--IP address of 10.10.4.2 is connected to R2 serial 2
- R5--Loopback0 10.10.0.5
 - Ethernet 0--IP address of 10.10.8.1 is connected to R6 Ethernet 0
 - Ethernet 1--IP address of 10.10.6.2 is connected to R4 Ethernet 1
 - Serial 2--IP address of 10.10.5.2 is connected to R3 serial 2
- R6--Loopback0 10.10.0.6
 - Ethernet 0--IP address of 10.10.7.2 is connected to R4 Ethernet 0
 - Ethernet 1--IP address of 10.10.8.2 is connected to R5 Ethernet 1

In the figure above, the following situations exist:

- Routers R1, R2, and R3 are in AS 100. The R1-R2 and R1-R3 links are in OSPF area 1.
- Routers R4, R5, and R6 are in AS200. The R4-R6 and R5-R6 links are in OSPF area 2.
- The link R2-R3 is in AS100, and link R4-R5 is in AS200. The links R2-R3 and R4-R5 are in OSPF area 0.
- The links R2-R4 and R3-R5 are not running an IGP because they cross the Inter-AS boundary between AS100 and AS200. Because they are not running IGP, you must configure an administrative weight for each passive interface for FRR to work. Use the **mpls traffic-eng administrative-weight** command in interface configuration mode.
- There is a primary tunnel, tunnel 100, from R1-R2-R4-R6.
- There is a backup tunnel, tunnel 102, from R2-R3-R5-R6.
- There is a TE tunnel, tunnel 101, from R6-R5-R3-R1 for returning data traffic for tunnel 100.
- There is a TE tunnel, tunnel 103, from R6-R5-R3-R2 for returning data traffic for tunnel 102.
- The explicit paths of all the tunnels use loose hops.
- The R2-R4 link is configured to be link flooded in both R2's and R4's IGP. The R3-R5 link is configured to be link flooded in both R3's and R5's IGP.

Router R2 needs to ensure the following:

- Backup tunnel intersects with the primary tunnel at the MP, and therefore has a valid MP address. In the figure above, R2 needs to determine that tunnel 100 and backup tunnel 102 share MP node R6.
- Backup tunnel satisfies the request of the primary LSP for bandwidth protection. For example, the amount of bandwidth guaranteed for the primary tunnel during a failure, and the type of protection (preferably protecting against a node failure rather than a link failure).

Node-IDs Signaling in RROs

ASBR node protection includes a node-ID flag (0x20), which is also called a node-ID subobject. When it is set, the flag indicates that the address specified in the RRO object in the resv message is the node-ID address. The node-ID address refers to the traffic engineering router ID.

A node must always use the same address in the RRO (that is, it must use IPv4 or IPv6, but not both).

To display all the hops, enter the following command on the headend router. Sample command output is as follows:

```
Router(config)# show ip rsvp reservations detail
Reservation:
  Tun Dest: 10.10.0.6 Tun ID: 100 Ext Tun ID: 10.10.0.1
  Tun Sender: 10.10.0.1 LSP ID: 31
  Next Hop: 10.10.1.2 on Ethernet0/0
  Label: 17 (outgoing)
  Reservation Style is Shared-Explicit, QoS Service is Controlled-Load
  Average Bitrate is 10K bits/sec, Maximum Burst is 1K bytes
  Min Policed Unit: 0 bytes, Max Pkt Size: 0 bytes
RRO:
  10.10.0.2/32, Flags:0x29 (Local Prot Avail/to NNHOP, Is Node-id)
  10.10.4.1/32, Flags:0x9 (Local Prot Avail/to NNHOP)
    Label subobject: Flags 0x1, C-Type 1, Label 17
  10.10.0.4/32, Flags:0x20 (No Local Protection, Is Node-id)
  10.10.7.1/32, Flags:0x0 (No Local Protection)
    Label subobject: Flags 0x1, C-Type 1, Label 17
  10.10.0.6/32, Flags:0x20 (No Local Protection, Is Node-id)
  10.10.7.2/32, Flags:0x0 (No Local Protection)
    Label subobject: Flags 0x1, C-Type 1, Label 0
Resv ID handle: 0100040E.
Status:
Policy: Accepted. Policy source(s): MPLS/TE
```

For a description of the fields, see the Cisco IOS Quality of Service Solutions Command Reference.

Addition of the Node-ID Subobject

When a fast reroutable LSP is signaled, the following actions occur:

- An LSR adds a node-ID subobject and an incoming label subobject in the resv message.
- If there is an RRO object in the path message, an LSR adds a node-ID subobject, an RRO IPv4 subobject that records the interface address, and an incoming label subobject in the resv message.

If you enable record-route on the headend LSR, the interface addresses for the LSP are included in the RRO object of the resv message.

To enable record-route, enter the following command with the **record-route** keyword:

```
tunnel mpls traffic-eng record-route
```

Processing of an RRO with Node-ID Subobjects

The node-ID subobject is added to the RECORD_ROUTE object before the label route subobject. If RECORD_ROUTE is turned on, the RRO object consists of the following in this order: node-ID, interface address, and label.

Merge Point Location

The destination of the backup tunnel is the node-ID of the MP. A PLR can find the MP and appropriate backup tunnel by comparing the destination address of the backup tunnel with the node-ID subobjects included in the resv RRO for the primary tunnel.

When both the IPv4 node-ID and IPv6 node-ID subobjects are present, a PLR can use either or both of them to find the MP address.

Determination of Backward Compatibility

To remain compatible with nodes that do not support RRO IPv4 or IPv6 node-ID subobjects, a node can ignore those objects. Those nodes cannot be the MP in a network with interarea or Inter-AS traffic engineering.

Loose Path Reoptimization

Interarea and Inter-AS LSPs

If the LSP of an MPLS TE tunnel traverses hops that are not in the headend router's topology database (that is, the hops are in a different OSPF area or IS-IS level), the LSP is called an *interarea TE LSP*.

If the LSP of the tunnel traverses hops that are in a different autonomous system (AS) from the tunnel's headend router, the LSP is called an *Inter-AS TE LSP*.

Interarea LSPs and Inter-AS TE LSPs can be signaled using loose hop subobjects in their EROs. The headend does not have "strict" knowledge of hops beyond its area, so the LSP's path is "loosely" specified at the headend. Downstream routers processing these loose hop subobjects (which do have the knowledge) are relied upon to expand them into strict hops.

Loose Hop Configuration

Beyond the headend area, configure hops as loose hops. Typically you specify only the ABRs and the tailend router of a tunnel, but any other combination is allowed.

Loose Hop Expansion

Loose hop expansion is the conversion of a single ERO loose hop subobject into one or more strict hop subobjects.

Interarea and Inter-AS TE LSPs can be signaled using loose hop subobjects in their EROs. When a router receives a path message containing an ERO that has a loose hop as the next address, the router typically expands the ERO by converting the single loose hop subobject into one or more strict hop subobjects. The router typically has the knowledge, in its topology database, of the best way to reach the loose hop and computes this path by using constraint-based shortest path first (CSPF). So the router substitutes this more specific information for the loose hop subobject found in the ERO. This process is called loose hop expansion or ERO expansion.

Loose hop expansions can occur at one or more hops along an LSP's path. This process is referred to as loose path reoptimization.

Tunnel Reoptimization Procedure

Tunnel reoptimization is the signaling of an LSP that is more optimal than the LSP a TE tunnel is currently using (for example, it may be shorter or may have a lower cost), and the switching over of the tunnel's data to use this new LSP.

The new more optimal TE LSP is always established and the data moved onto it before the original LSP is torn down (so it is called the "make before break" procedure). This ensures that no data packets are lost during the transition to the new LSP.

For tunnel reoptimization to function:

- Each router must be configured with the **mpls traffic-eng reoptimize events link-up** command.
- Each passive link must have an assigned administrative weight. To configure an administrative weight, use the **mpls traffic-eng administrative-weight** command in interface configuration mode.

The TE LSPs reoptimization process is triggered under the following circumstances:

- Periodically (based on a timer)
- User entered a command (**mpls traffic-eng reoptimize**) requesting reoptimization
- Network event, such as a link-up

Regardless of how reoptimization is triggered, the headend router reoptimizes a tunnel only if it can find a better path than the one the tunnel currently uses. If there is not a better path in the local topology database, no new LSP is signaled and reoptimization does not occur.

Prior to the addition of loose path reoptimization, interarea TE LSPs were not reoptimized if a better path became available in any area beyond the headend area. This is because the headend router was not capable of finding a better path when the better path existed in an area beyond its view (that is, it was not in its local topology database).

With the addition of loose path reoptimization, a tunnel's headend can reoptimize LSPs even if they span multiple areas, levels, or autonomous systems. This is done via the implementation of a query and response protocol defined in *draft-vasseur-mpls-loose-path-reopt-02.txt*. This draft defines a protocol whereby a tunnel's headend may query downstream routers to perform ERO expansion for this tunnel's LSP. These downstream routers respond in the affirmative if they can find a more optimal path than the one in use. (This is done via a new ERO expansion.) Having received an affirmative answer to its query, a headend signals a new LSP for the tunnel, and the new LSP benefits from a new ERO expansion along the better path.

Loose path reoptimization is on by default, and cannot be disabled. Whenever an LSP reoptimization is attempted but the headend fails to find a better path, if the LSP contains loose ERO subobjects, a query is sent downstream to determine whether downstream routers can find a better path. If an affirmative answer comes back, the LSP is reoptimized. That is, a new LSP is signaled (which will follow the better path), the tunnel's data packets are switched over to use this new LSP, and the original LSP is torn down.

For details on this query and response protocol, see *draft-vasseur-mpls-loose-path-reopt-02.txt*.

ASBR Forced Link Flooding

When you configure forced link flooding on an interface, the MPLS TE link management module advertises the link to all nodes. As a result of this advertisement, the TE topology database on all the nodes within the Inter-AS is updated with this information.

ASBR forced link flooding allows the links to be advertised even if IGP adjacencies are not running over these links. TE LSPs can traverse these links at the edge of a network between two nodes running BGP (or static routes) even if the exit ASBR is not listed in the IP explicit path. Therefore, a headend LSR can consider that link when it computes its TE LSP path.

Configuration of ASBR Forced Link Flooding

To activate ASBR forced link flooding, configure a link as passive and provide neighbor information (that is, the neighbor IGP ID and the neighbor TE ID). The required configuration tasks are described in the [Configuring a Static Route from the MP to the PLR](#), on page 14.

Link Flooding

A passive link is configured on an interface of an ASBR. The link is flooded in the ASBR's IGP. All the links are flooded as point-to-point links.

Flooding notifications are also sent when there is a change to a link's property.

OSPF Flooding

OSPF floods opaque link-state advertisement (LSA) Type 10 link information.

If a multiaccess link has more than one neighbor, a Type 10 LSA is advertised for each neighbor. In the topology database, neighbors are represented by point-to-point neighbor relationships.

Link TLV

A link TLV describes a single link and contains multiple sub-TLVs.

An opaque LSA contains a single link TLV.

For each ASBR-to-ASBR link, an ASBR must flood an opaque LSA containing one link TLV that has the link's attributes.

A link TLV comprises the following sub-TLVs:

- Link type (1 octet)--(Required) Defines the type of the link. The link type of a passive interface always is 1 (point-to-point), even for a multiaccess subnetwork.
- Link ID (4 octets)--(Required) Identifies the other end of the link for a point-to-point link. Includes the system ID of the neighbor, requires static configuration for a multiaccess ASBR-to-ASBR link, and includes the system ID of the neighbor.
- Local interface IP address (4 octets)--Specifies the IP addresses of the neighbor's interface corresponding to this link.
- Remote interface IP address (4 octets)--Specifies the IP addresses of the neighbor's interface corresponding to this link. The remote interface IP address is set to the router ID of the next hop. There must be a static configuration for the ASBR-to-ASBR link.
- Traffic engineering metric (4 octets)
- Maximum bandwidth (4 octets)
- Maximum reservable bandwidth (4 octets)
- Unreserved bandwidth (32 octets)
- Administrative group (4 octets)

IS-IS TLV

In IS-IS, when autonomous system A1 floods its LSP, it includes the system ID and a pseudonode number.

If three autonomous systems are connected to a multiaccess network LAN, each link is considered to be a point-to-point link. The links are marked with the maximum metric value so that the inter-ASBR links are considered by CSPF and not by shortest path first (SPF).

TE uses the protocol TLV type 22, which has the following data structure:

- System ID and pseudonode number node (7 octets)

- Default metric (3 octets)
- Length of sub-TLVs (1 octet)
- Sub-TLVs (0 to 244 octets), where each sub-TLV consists of a sequence of the following: 1 octet for subtype, 1 octet for the length of the value field of the sub-TLV, and 0 to 242 octets for the value

The table below defines the sub-TLVs.

Table 1: Sub-TLVs

Sub-TLV	Length (Octets)	Name
3	4	Administrative group (color).
6	4	IPv4 address for the interface described by the main TLV.
8	4	IPv4 address for a neighboring router on this link. This will be set to the router ID of the next hop.
9	4	Maximum link bandwidth.
10	4	Reservable link bandwidth.
11	32	Unreserved bandwidth.
18	3	TE default metric.
250 to 254	--	Reserved for Cisco-specific extensions.
255	--	Reserved for future expansion.



Note

The TE router ID is TLV type 134.

Topology Database

When the topology database module receives a link-state advertisement (LSA), the module scans the LSA to find the neighbors of the links. The ASBR link is part of the same LSA and is installed in the TE topology database like any other link.

During the CSPF operation, the TE headend module uses the TE topology database to find a path to the destination. Because the Inter-AS links are part of the TE topology database, the CSPF operation uses these links to compute the LSP path.

Link Flooding

The IGP floods information about a link in the following situations:

- When a link goes down
- When a link's configuration is changed (for example, when the link cost is modified)
- When it is time to periodically reflood the router's IGP information
- When link bandwidth changes significantly

Flooding is a little different in IS-IS and OSPF. In OSPF, only information about the link that has changed is flooded, because a Type 10 LSA contains a single link advertisement. In IS-IS, information about all links on a node is flooded even if only one has changed, because the Type 22 TLV contains a list of all links on the router.

How to Configure MPLS Traffic Engineering Inter-AS TE



Note

There is no configuration procedure for loose path reoptimization.

Configuring Loose Hops

The section describes how to do the following so that there can be loose hops:

Configuring an Explicit Path on the Tunnel That Will Cross the Inter-AS Link

If you want a tunnel to span multiple networks, configure an explicit path on the tunnel that will cross the Inter-AS link by performing the following steps.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **ip explicit-path** {name *path-name* | identifier *number*} [**enable** | **disable**]
4. **next-address loose** *A.B.C.D*
5. **interface tunnel** *number*
6. **tunnel mpls traffic-eng fast-reroute**
7. **mpls traffic-eng reoptimize events link-up**

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: Router> enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: Router# configure terminal	Enters global configuration mode.
Step 3	ip explicit-path {name <i>path-name</i> identifier <i>number</i>} [enable disable] Example: Router(config)# ip explicit-path identifier 2 enable	Enters the subcommand mode for IP explicit paths and creates or modifies the explicit path. This command places the router in IP explicit path configuration mode.
Step 4	next-address loose <i>A.B.C.D</i> Example: Router(cfg-ip-expl-path)# next-address loose 10.10.0.2	Specifies the next loose IP address in the explicit path. Each area border router (ABR) the path must traverse should be specified in a next-address loose command. This command places the router in global configuration mode.
Step 5	interface tunnel <i>number</i> Example: Router(config)# interface tunnel 100	Configures a tunnel interface. This command places the router in interface configuration mode.
Step 6	tunnel mpls traffic-eng fast-reroute Example: Router(config-if)# tunnel mpls traffic-eng fast-reroute	Enables an MPLS traffic engineering tunnel to use an established backup tunnel in the event of a link failure.
Step 7	mpls traffic-eng reoptimize events link-up Example: Router(config)# mpls traffic-eng reoptimize events link-up	Enables automatic reoptimization of MPLS traffic engineering when an interface becomes operational.

Configuring a Route to Reach the Remote ASBR

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **ip route prefix mask {ip-address | interface-type interface-number}**

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: Router> enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: Router# configure terminal	Enters global configuration mode.
Step 3	ip route prefix mask {ip-address interface-type interface-number} Example: Router(config)# ip route 10.10.0.1 255.255.255.255 tunnel 101	Establishes static routes.

Configuring a Static Route from the MP to the PLR

To enable Fast Reroute protection that spans across different autonomous systems, configure a static route from the MP to the PLR by performing the following steps.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **ip route prefix mask ip-address outgoing-interface**

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: <pre>Router> enable</pre>	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: <pre>Router# configure terminal</pre>	Enters global configuration mode.
Step 3	ip route <i>prefix mask ip-address outgoing-interface</i> Example: <pre>Router(config)# ip route 10.10.3.1 255.255.255.255 10.0.0.0 FastEthernet0/0</pre>	Establishes static routes. Refer to the appropriate hardware manual for interface information. Note Enter this command on the MP. The destination is the PLR.

Configuring ASBR Forced Link Flooding

This section describes how to do the following so that you can configure ASBR forced link flooding:

Configuring the Inter-AS Link as a Passive Interface Between Two ASBRs

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **interface *type slot/port***
4. **ip address *ip-address mask* [secondary]**
5. **mpls traffic-eng passive-interface *nbr-te-id te-router-id* [*nbr-if-addr if-addr*] [*nbr-igp-id {isis sysid | ospf sysid}*]**
6. **mpls traffic-eng administrative-weight *weight***

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable	Enables privileged EXEC mode.

	Command or Action	Purpose
	Example: Router> enable	<ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: Router# configure terminal	Enters global configuration mode.
Step 3	interface <i>type slot/port</i> Example: Router(config)# interface serial 2/0	Specifies an interface and enters interface configuration mode. Refer to the appropriate hardware manual for interface information.
Step 4	ip address <i>ip-address mask [secondary]</i> Example: Router(config-if)# ip address 10.10.4.1 255.255.255.0	Sets a primary or secondary IP address for an interface.
Step 5	mpls traffic-eng passive-interface <i>nbr-te-id te-router-id [nbr-if-addr if-addr] [nbr-igp-id {isis sysid ospf sysid}]</i> Example: Router(config-if)# mpls traffic-eng passive-interface nbr-te-id 10.10.11.12 nbr-igp-id ospf 10.10.15.18	Configures a link as a passive interface between two ASBRs. Note For an RSVP Hello configuration on the Inter-AS link, all fields are required.
Step 6	mpls traffic-eng administrative-weight <i>weight</i> Example: Router(config-if)# mpls traffic-eng administrative-weight 20	Overrides the Interior Gateway Protocol (IGP) administrative weight (cost) of the link and assigns a specific weight for the link.

Creating LSPs Traversing the ASBRs

To create LSPs traversing the ASBRs, perform the following steps.



Note Perform Steps 3 through 7 for the primary LSP and then for the backup LSP.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **ip explicit path *name* enable**
4. **next-address loose *A.B.C.D***
5. **interface tunnel *number***
6. **tunnel mpls traffic-eng fast-reroute**
7. **tunnel mpls traffic-eng path-option $number$ {dynamic | explicit | {name *path-name* | *path-number*}}**
[lockdown]

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: Router> enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: Router# configure terminal	Enters global configuration mode.
Step 3	ip explicit path <i>name</i> enable Example: Router(config)# ip explicit path route1 enable	Specifies the name of the explicit path and enables the path.
Step 4	next-address loose <i>A.B.C.D</i> Example: Router(config)# next-address loose 10.10.10.2	Configures a loose hop.
Step 5	interface tunnel <i>number</i> Example: Router(config)# interface tunnel 100	Configures a tunnel interface and enters interface configuration mode.
Step 6	tunnel mpls traffic-eng fast-reroute Example: Router(config-if)# tunnel mpls traffic-eng fast-reroute	Enables an MPLS traffic engineering tunnel to use an established backup tunnel in the event of a link failure.

	Command or Action	Purpose
Step 7	tunnel mpls traffic-eng path-option <i>number</i> { dynamic explicit { name <i>path-name</i> <i>path-number</i> }} [lockdown] Example: <pre>Router(config-if)# tunnel mpls traffic-eng path-option 1 route1</pre>	Configures a path option for an MPLS traffic engineering tunnel.

Configuring Multiple Neighbors on a Link

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **interface** *type slot/port*
4. **mpls traffic-eng passive-interface** [**nbr-te-id**] [*router-id* | *te-id*] [**nbr-igp-id**] [*isis sysid* | *ospf sysid*]
5. **mpls traffic-eng administrative-weight** *weight*

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: <pre>Router> enable</pre>	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: <pre>Router# configure terminal</pre>	Enters global configuration mode.
Step 3	interface <i>type slot/port</i> Example: <pre>Router(config)# interface serial 2/0</pre>	Specifies an interface and enters interface configuration mode. Refer to the appropriate hardware manual for interface information.
Step 4	mpls traffic-eng passive-interface [nbr-te-id] [<i>router-id</i> <i>te-id</i>] [nbr-igp-id] [<i>isis sysid</i> <i>ospf sysid</i>] 	Configures a link as a passive link.

	Command or Action	Purpose
	<p>Example:</p> <pre>Router(config-if)# mpls traffic-eng passive-interface nbr-te-id 10.10.0.4 nbr-igp-id ospf 10.10.0.4</pre>	
Step 5	<p>mpls traffic-eng administrative-weight <i>weight</i></p> <p>Example:</p> <pre>Router(config-if)# mpls traffic-eng administrative-weight 20</pre>	Overrides the Interior Gateway Protocol (IGP) administrative weight (cost) of the link and assigns a specific weight for the link.

Troubleshooting Tips

The following debug commands are useful for troubleshooting issues with MPLS Traffic Engineering: Inter-AS TE.

Debugging Headend of TE LSPs

```
debug mpls traffic-eng path lookup
debug mpls traffic-eng path verify
debug mpls traffic-eng path spf
```

Debugging Head and Midpoint (Link-Related Debugs)

```
debug mpls traffic-eng link-management igp-neighbors
debug mpls traffic-eng link-management advertisements
debug mpls traffic-eng link-management bandwidth-allocation
debug mpls traffic-eng link-management routing
```

Verifying the Inter-AS TE Configuration

To verify the Inter-AS TE configuration, perform the following steps.



Note

Perform Step 1 for Fast Reroute ready, and Step 2 for Fast Reroute active.

SUMMARY STEPS

1. **show ip rsvp sender detail**
2. **show ip rsvp sender detail**
3. **show mpls traffic-eng link-management advertisements**

DETAILED STEPS

Step 1 **show ip rsvp sender detail**

Use this command to display the MP sender display for the primary tunnel when Fast Reroute is ready.

Example:

```
Router# show ip rsvp sender detail
PATH:
Tun Dest: 10.10.0.6 Tun ID: 100 Ext Tun ID: 10.10.0.1
Tun Sender: 10.10.0.1 LSP ID: 31
Path refreshes:
  arriving: from PHOP 10.10.7.1 on Et0/0 every 30000 msec
Session Attr:
  Setup Prio: 7, Holding Prio: 7
  Flags: (0x7) Local Prot desired, Label Recording, SE Style
  session Name: R1_t100
ERO: (incoming)
  10.10.7.2 (Strict IPv4 Prefix, 8 bytes, /32)
  10.10.0.6 (Strict IPv4 Prefix, 8 bytes, /32)
RRO:
  10.10.7.1/32, Flags:0x0 (No Local Protection)
  10.10.4.1/32, Flags:0x9 (Local Prot Avail/to NNHOP) !Available to NNHOP
  10.10.1.1/32, Flags:0x0 (No Local Protection)
Traffic params - Rate: 10K bits/sec, Max. burst: 1K bytes
  Min Policed Unit: 0 bytes, Max Pkt Size 4294967295 bytes
Fast-Reroute Backup info:
  Inbound FRR: Not active
  Outbound FRR: No backup tunnel selected
Path ID handle: 50000416.
Incoming policy: Accepted. Policy source(s): MPLS/TE
Status: Proxy-terminated
```

Step 2 **show ip rsvp sender detail**

Use this command to display the MP sender display when the primary tunnel is Fast Reroute active:

Example:

```
Router# show ip rsvp sender detail
PATH:
Tun Dest: 10.10.0.6 Tun ID: 100 Ext Tun ID: 10.10.0.1
Tun Sender: 10.10.0.1 LSP ID: 31
Path refreshes:
  arriving: from PHOP 10.10.3.1 on Et1/0 every 30000 msec
Session Attr:
  Setup Prio: 7, Holding Prio: 7
  Flags: (0x7) Local Prot desired, Label Recording, SE Style
  Session Name: R1_t100
ERO: (incoming)
  10.10.0.4 (Strict IPv4 Prefix, 8 bytes, /32)
  10.10.0.6 (Loose IPv4 Prefix, 8 bytes, /32)
RRO:
  10.10.3.1/32, Flags:0xB (Local Prot Avail/In Use/to NNHOP) !Ready
  10.10.1.1/32, Flags:0x0 (No Local Protection)
Traffic params - Rate: 10K bits/sec, Max. burst: 1K bytes
  Min Policed Unit: 0 bytes, Max Pkt Size 4294967295 bytes
Fast-Reroute Backup info:
  Inbound FRR: Active
  Orig Input I/F: Et0/0
  Orig PHOP: 10.10.7.1
  Now using Bkup Filterspec w/ sender: 10.10.3.1 LSP ID: 31
  Outbound FRR: No backup tunnel selected
Path ID handle: 50000416.
```

```
Incoming policy: Accepted. Policy source(s): MPLS/TE
Status: Proxy-terminated
```

Step 3 show mpls traffic-eng link-management advertisements

Use this command to display the influence of a passive link. On R2, the passive link to R4 is in the Link ID:: 1 section.

Example:

```
Router# show mpls traffic-eng link-management advertisements
```

```
Flooding Status: ready
Configured Areas: 2
IGP Area[1] ID:: ospf 1 area 0
System Information::
  Flooding Protocol: OSPF
Header Information::
  IGP System ID: 10.10.0.2
  MPLS TE Router ID: 10.10.0.2
  Flooded Links: 2
Link ID:: 1
  Link Subnet Type: Point-to-Point
  Link IP Address: 10.10.4.1
  IGP Neighbor: ID 0-0-0-0-0-0-0, IP 10.10.0.4
  Physical Bandwidth: 1544 kbits/sec
  Res. Global BW: 1158 kbits/sec
  Res. Sub BW: 0 kbits/sec
Downstream::
      Global Pool Sub Pool
      -----
Reservable Bandwidth[0]: 1158      0 kbits/sec
Reservable Bandwidth[1]: 1158      0 kbits/sec
Reservable Bandwidth[2]: 1158      0 kbits/sec
Reservable Bandwidth[3]: 1158      0 kbits/sec
Reservable Bandwidth[4]: 1158      0 kbits/sec
Reservable Bandwidth[5]: 1158      0 kbits/sec
Reservable Bandwidth[6]: 1158      0 kbits/sec
Reservable Bandwidth[7]: 1148      0 kbits/sec
Attribute Flags: 0x00000000
IGP Area[1] ID:: ospf 1 area 1
System Information::
  Flooding Protocol: OSPF
Header Information::
  IGP System ID: 10.10.0.2
  MPLS TE Router ID: 10.10.0.2
  Flooded Links: 2
Link ID:: 1
  Link Subnet Type: Point-to-Point
  Link IP Address: 10.10.4.1
  IGP Neighbor: ID 0-0-0-0-0-0-0, IP 10.10.0.4
  Physical Bandwidth: 1544 kbits/sec
  Res. Global BW: 1158 kbits/sec
  Res. Sub BW: 0 kbits/sec
Downstream::
      Global Pool Sub Pool
      -----
Reservable Bandwidth[0]: 1158      0 kbits/sec
Reservable Bandwidth[1]: 1158      0 kbits/sec
Reservable Bandwidth[2]: 1158      0 kbits/sec
Reservable Bandwidth[3]: 1158      0 kbits/sec
Reservable Bandwidth[4]: 1158      0 kbits/sec
Reservable Bandwidth[5]: 1158      0 kbits/sec
Reservable Bandwidth[6]: 1158      0 kbits/sec
Reservable Bandwidth[7]: 1148      0 kbits/sec
Attribute Flags: 0x00000000
```

Configuration Examples for MPLS Traffic Engineering Inter-AS TE

Configuring Loose Hops Examples

Configuring an Explicit Path on the Tunnel That Will Cross the Inter-AS Link Example

The following commands configure a loose IP explicit path named `route1` suitable for use as a path option with Inter-AS TE with the destination 10.10.10.6 that is to traverse ABRs 10.10.0.2 and 10.10.0.4. The tunnel headend and the specified ABRs will find a path from the source AS100 to the destination 10.10.0.6 in AS200. See the figure above.

```
Router(config)# ip explicit-path name route1 enable
Router(cfg-ip-expl-path)# next-address loose 10.10.0.2
Router(cfg-ip-expl-path)# next-address loose 10.10.0.4
Router(cfg-ip-expl-path)# next-address loose 10.10.0.6
```

Note that the explicit path for an interarea TE tunnel need not specify the destination router because the tunnel configuration specifies it in the tunnel destination command. The following commands configure an explicit path named `path-without-tailend` that would work equally well for the interarea tunnel created in the previous example:

```
Router(config)# ip explicit-path name path-without-tailend
Router(cfg-ip-expl-path)# next-address loose 10.10.0.2
Router(cfg-ip-expl-path)# next-address loose 10.10.0.4
```

Configuring a Route to Reach the Remote ASBR in the IP Routing Table Example

In the following example, packets for the ASBR whose router ID is 10.10.0.1 will be forwarded via tunnel 101:

```
Router> enable
Router# configure terminal
Router(config)# ip route 10.10.0.1 255.255.255.255 tunnel 101
```

Configuring a Static Route from the MP to the PLR Example

In the following example, a static route is configured from the MP to the PLR. The outgoing interface is tunnel 103.

```
Router> enable
Router# configure terminal
Router(config)# ip route 10.10.3.1 255.255.255.255 tunnel 103
```

Configuring ASBR Forced Link Flooding Examples

Configuring the Inter-AS Link as a Passive Interface Example

For this example, see the figure above.

Routers R2 and R4 have the following router IDs:

- Router R2--10.10.0.2
- Router R4--10.10.0.4

```
Router> enable
Router# configure terminal
Router(config)# interface serial 2/0
```

Configures OSPF on Router R2 When Its Neighbor Is Running OSPF Too

```
Router(config-if)# mpls traffic-eng passive-interface nbr-te-id 10.10.0.4
```



Note

Because both routers are running OSPF, the **nbr-igp-id** keyword is not specified.

Specifies That Both Router R2 and Its Neighbor Are Running OSPF (the nbr-igp-id Keyword Is Specified)

```
Router(config-if)# mpls traffic-eng passive-interface nbr-te-id 10.10.0.4 nbr-igp-id ospf
10.10.0.4
```

Configures IS-IS on Router R1

```
Router(config-if)# mpls traffic-eng passive-interface nbr-te-id 10.10.0.4 nbr-igp-id isis
40.0000.0002.0001.00
```

Configures the Neighbor IGP ID (nbr-igp-id) When There Is More than One Neighbor Specified on a Link

```
Router(config-if)# mpls traffic-eng passive-interface nbr-te-id 10.10.0.4 nbr-igp-id ospf
10.10.0.4
Router(config-if)# mpls traffic-eng passive-interface nbr-te-id 10.10.0.7 nbr-igp-id ospf
10.10.0.7
```

Overrides the Interior Gateway Protocol (IGP) Administrative Weight of the Link and Assigns a Specific Weight

```
Router(config-if)# mpls traffic-eng administrative-weight 20
```



Note

The ID is unique for each neighbor.

Configures a Link as a Passive Interface (Includes Global TE Commands)

```
interface serial 2/0
ip address 10.10.4.1.255.255.255.0
mpls traffic-eng tunnels
mpls traffic-eng administrative-weight 10
mpls traffic-eng passive-interface nbr-te-id 10.10.0.4 nbr-igp-id ospf 10.10.0.4
ip rsvp bandwidth 1000
mpls traffic-eng administrative-weight 20
```

Creating LSPs Traversing the ASBRs Example

In the following example, a primary LSP is created:

```
Router> enable
Router# configure terminal
Router(config)# ip explicit path routel enable
Router(config)# next-address loose 10.10.0.2
Router(config)# next-address loose 10.10.0.4
Router(config)# next-address loose 10.10.0.6
Router(config)# interface tunnel 100
Router(config-if)# tunnel mpls traffic-eng fast reroute
Router(config-if)# tunnel mpls traffic-eng path-option 1 routel
```

In the following example, a backup LSP is created:

```
Router> enable
Router# configure terminal
Router(config)# ip explicit path backpath1 enable
Router(config)# next-address loose 10.10.0.3
Router(config)# next-address loose 10.10.0.5
Router(config)# next-address loose 10.10.0.6
Router(config)# interface tunnel 102
Router(config)# mpls traffic-eng backup path tunnel 102
Router(config-if)# tunnel mpls traffic-eng path-option 1 backpath1
```

Configuring Multiple Neighbors on a Link Example

In the following example, there is more than one neighbor on a link:

```
Router> enable
Router# configure terminal
Router(config)# interface ethernet 2/0
Router(config-if)# mpls traffic-eng passive-interface nbr-te-id 10.10.0.4 nbr-igp-id ospf
10.10.0.4
Router(config-if)# mpls traffic-eng administrative-weight 20
```


Additional References

Related Documents

Related Topic	Document Title
MPLS traffic engineering commands: complete command syntax, command mode, command history, defaults, usage guidelines, and examples	<i>Cisco IOS Multiprotocol Label Switching Command Reference</i>
Fast Reroute	MPLS TE: Link and Node Protection, with RSVP Hellos Support (with Fast Tunnel Interface Down Detection)
Link flooding and node protection	MPLS Traffic Engineering: Interarea Tunnels
IS-IS configuration tasks	Configuring a Basic IS-IS Network
OSPF configuration tasks	Configuring OSPF
IS-IS and OSPF commands: complete command syntax, command mode, command history, defaults, usage guidelines, and examples	<i>Cisco IOS IP Routing Protocols Command Reference</i>
RSVP	RSVP Message Authentication

Standards

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

MIBs

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

RFCs

RFCs	Title
RFC 3209	Extensions to RSVP for LSP Tunnels
draft-ietf-mpls-rsvp-lsp-fastreroute-02.txt	<i>Fast Reroute Extensions to RSVP-TE for LSP Tunnels</i>
draft-vasseur-mpls-loose-path-reopt-02.txt	<i>Reoptimization of an Explicitly Loosely Routed MPLS TE Path</i>
draft-vasseur-mpls-inter-as-te-00.txt	<i>MPLS Inter-AS Traffic Engineering</i>
draft-ietf-mpls-soft-preemption-00.txt	<i>MPLS Traffic Engineering Soft Preemption</i>

Technical Assistance

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

Feature Information for MPLS Traffic Engineering Inter-AS TE

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

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

Table 2: Feature Information for MPLS Traffic Engineering: Inter-AS TE

Feature Name	Releases	Feature Information
MPLS Traffic Engineering: Inter-AS TE	12.0(29)S 12.2(33)SRA 12.2(33)SRB 12.2(33)SXH 12.4(20)T Cisco IOS XE Release 3.5S	<p>The MPLS Traffic Engineering: Inter-AS TE feature provides ASBR node protection, loose path reoptimization, SSO recovery of LSPs that include loose hops, ASBR forced link flooding, Cisco IOS RSVP local policy extensions for Inter-AS, and per-neighbor key capabilities.</p> <p>In 12.0(29)S, this feature was introduced.</p> <p>In 12.2(33)SRA, the nbr-if-addr keyword was added to the mpls traffic-eng passive-interface command.</p> <p>In 12.2(33)SRB, support was added for SSO recovery of LSPs that include loose hops.</p> <p>In 12.2(33)SXH, this feature was integrated into Cisco IOS Release 12.2(33)SXH.</p> <p>In 12.4(20)T, this feature was integrated into Cisco IOS Release 12.4(20)T.</p> <p>In Cisco IOS XE Release 3.5S, this feature was integrated into Cisco IOS XE Release 3.5S.</p>

Glossary

ABR --Area Border Router. A routers connecting two areas.

adjacency --The MPLS TE Forwarding Adjacency feature allows a network administrator to handle a traffic engineering, label-switched path (LSP) tunnel as a link in an Interior Gateway Protocol (IGP) network based on the Shortest Path First (SPF) algorithm. A forwarding adjacency can be created between routers regardless of their location in the network. The routers can be located multiple hops from each other.

area --A logical set of network segments (for example, one that is OSPF-based) and their attached devices. Areas usually are connected to other areas by routers, making up a single autonomous system. OSPF and IS-IS define their areas differently. OSPF area borders are marked by routers. Some interfaces are in one area, and other interfaces are in another area. With IS-IS, all the routers are completely within an area, and the area borders are on links, not on routers. The routers that connect the areas are level-2 routers, and routers that have no direct connectivity to another area are level-1 routers.

ASBR --Autonomous System Boundary Router. The router is located between an OSPF autonomous system and a non-OSPF network. ASBRs run both OSPF and another routing protocol, such as RIP. ASBRs must reside in a nonstub OSPF area.

autonomous system --A collection of networks under a common administration sharing a common routing strategy. Autonomous systems are subdivided by areas.

backup tunnel --An MPLS traffic engineering tunnel used to protect other (primary) tunnel's traffic when a link or node failure occurs.

BGP --Border Gateway Protocol. Interdomain routing protocol that replaces EGP. BGP exchanges reachability information with other BGP systems.

border router --A router at the edge of a provider network that interfaces to another provider's border router using extended BGP procedures.

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

Fast Reroute --A mechanism for protecting MPLS traffic engineering (TE) LSPs from link and node failure 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 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.

flooding --A traffic-passing technique 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.

forwarding adjacency --A traffic engineering link (or LSP) into an IS-IS or OSPF network.

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).

IGP --Interior Gateway Protocol. Internet protocol used to exchange routing information within an autonomous system. Examples of common IGPs include Interior Gateway Routing Protocol (IGRP), Open Shortest Path First (OSPF), and Routing Information Protocol (RIP).

Inter-AS LSP --An MPLS traffic engineering label-switched path (LSP) that traverses hops that are not in the headend's topology database (that is, it is not in the same OSPF area, IS-IS area, or autonomous system as the headend).

interface --A network connection.

IP explicit path --A list of IP addresses, each representing a node or link in the explicit path.

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

link --A point-to-point connection between adjacent nodes.

LSA --link-state advertisement. A broadcast packet used by link-state protocols that contains information about neighbors and path costs. LSAs are used by the receiving routers to maintain their routing tables.

LSP --label-switched path. A configured connection between two routers, in which MPLS is used to carry packets. An LSP is a path created by the concatenation of one or more label-switched hops, allowing a packet to be forwarded by swapping labels from an MPLS node to another MPLS node.

midpoint --A transit router for a given LSP.

midpoint reoptimization --Ability of a midpoint to trigger a headend reoptimization.

MP --merge point. The LSR where one or more backup tunnels rejoin the path of the protected LSP, downstream of the potential failure. An LSR can be both an MP and a PLR simultaneously.

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.

multicast --Single packets are copied by the network and sent to a specific subset of network addresses. These addresses are specified in the Destination address field. (Multicast is an efficient paradigm for transmitting the same data to multiple receivers, because of its concept of a Group address. This allows a group of receivers to listen to the single address.)

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.

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.

opaque LSA --If a router understands LSA Type 10 link information, the router continues flooding the link throughout the network.

passive link --When IGP is not running on the link between two ASBRs, traffic engineering informs the IGP to flood link information on behalf of that link (that is, it advertises that link).

PLR --point of local repair. The headend LSR of a backup tunnel.

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

RSVP --Resource Reservation Protocol. 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.

SPF --shortest path first. A routing algorithm used as the basis for OSPF operations. When an SPF router is powered up, it initializes its routing-protocol data structures and then waits for indications from lower-layer protocols that its interfaces are functional.

SRLG --Shared Risk Link Group. Sets of links that are likely to go down together (for example, because they have the same underlying fiber).

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

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

TLV --type, length, values. A block of information embedded in Cisco Discovery Protocol advertisements.

