CHAPTER 9

MPLS Traffic Engineering

MPLS Traffic Engineering (MPLS TE) is a growing implementation in today's service provider networks. MPLS adoption in service provider networks has increased manifold due to its inherent TE capabilities. MPLS TE allows the MPLS-enabled network to replicate and expand upon the TE capabilities of Layer 2 ATM and Frame Relay networks. MPLS uses the reachability information provided by Layer 3 routing protocols and operates like a Layer 2 ATM network. With MPLS, TE capabilities are integrated into Layer 3, which can be implemented for efficient bandwidth utilization between routers in the SP network.

This chapter provides you with information on the operation and configuration of MPLS TE on Cisco routers.

TE Basics

TE is the process of steering traffic across to the backbone to facilitate efficient use of available bandwidth between a pair of routers. Prior to MPLS TE, traffic engineering was performed either by IP or by ATM, depending on the protocol in use between two edge routers in a network. Though the term "traffic engineering" has attained popularity and is used more in the context of MPLS TE today, traditional TE in IP networks was performed either by IP or by ATM.

TE with IP was mostly implemented by manipulation of interface cost when multiple paths existed between two endpoints in the network. In addition, static routes enabled traffic steering along a specific path to a destination. Figure 9-1 outlines a basic IP network with two customers, A and B, connected to the same service provider.

As illustrated in Figure 9-1, two paths exist between customer routers CE1-A and CE2-A via the provider network. If all links between the routers in Figure 9-1 were of equal cost, the preferred path between customer routers CE1-A and CE2-A would be the one with the minimum cost (via routers PE1-AS1, P3-AS1, and PE2-AS1) or *PATH1*. The same would apply for the customer routers CE1-B and CE2-B belonging to Customer B. If all the links were T3 links, for example, in the event of CE1-A sending 45 Mbps of traffic and CE1-B simultaneously sending 10 Mbps of traffic, some packets will be dropped at PE1-AS1 because the preferred path for both customers is using *PATH1*. The path *PATH2* will not be utilized for traffic forwarding; therefore, TE can utilize this available bandwidth. To

implement TE using IP whereby the paths *PATH1* and *PATH2* are either load balanced or used equally, we will need to implement IGP features such as maximum paths with variance or change the cost associated with the suboptimal path, *PATH2*, to make it equal to the current optimal path, *PATH1*. In an SP environment, this is often cumbersome to implement as the number of routers is much larger.

Figure 9-1 *Traditional IP Networks*

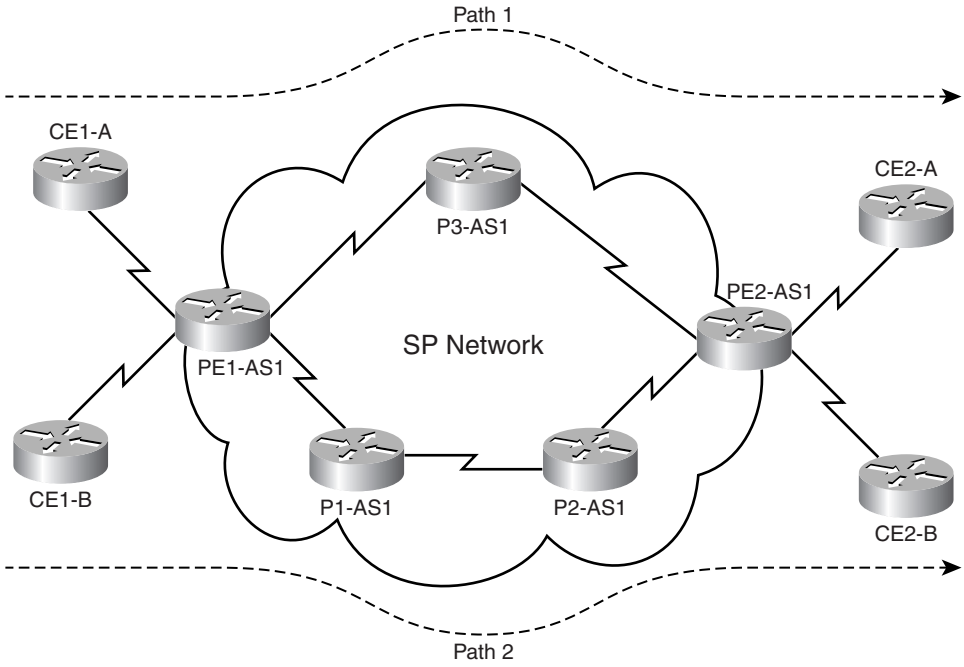

With ATM networks, the solution is a lot more feasible; PVCs can be configured between routers PE1-AS1 and PE2-AS1 with the same cost, but this would create a full mesh of PVCs between a group of routers. Implementing ATM for TE, however, has an inherent problem when a link or a node goes down. During link or node failure used in conjunction with ATM for TE, messages are flooded on the network. The Layer 3 topology must be predominantly fully meshed to take advantage of the Layer 2 TE implementation. Often, this might prove to be a scalability constraint for the IGP in use, due to issues with reconvergence at Layer 3.

The main advantage of implementing MPLS TE is that it provides a combination of ATM's TE capabilities along with the class of service (CoS) differentiation of IP. In MPLS TE, the headend router in the network controls the path taken by traffic to any particular destination in the network. The requirement to implement a full mesh of VCs, as in ATM, does not exist when implementing MPLS TE. Therefore, when MPLS TE is implemented, the IP network depicted in Figure 9-1 transforms into the label switched domain, as shown in Figure 9-2,

in which the TE label switched paths or TE tunnels (Tunnel1 and Tunnel2) define paths that can be used by traffic between PE1-AS1 and PE2-AS1.

Figure 9-2 *MPLS TE*

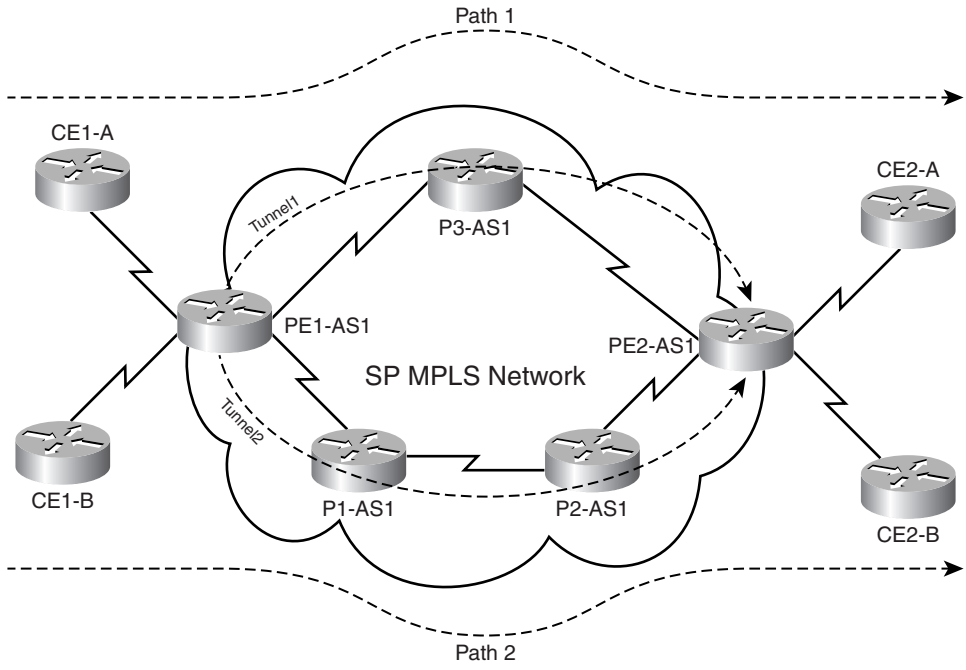

MPLS TE Theory

This section introduces you to the theoretical nuances in the implementation of MPLS TE. The primary topics covered will be the components of MPLS TE as well as RSVP and its function in the implementation of MPLS TE.

MPLS TE Overview

In a traditional IP forwarding paradigm, packets are forwarded on a per-hop basis where a route lookup is performed on each router from source to destination. As cited earlier, the destination-based forwarding paradigm leads to suboptimal use of available bandwidth between a pair of routers in the service provider network. Predominantly, the suboptimal paths are under-utilized in IP networks. To avoid packet drops due to inefficient use of available bandwidth and to provide better performance, TE is employed to steer some of the traffic destined to follow the optimal path to a suboptimal path to enable better bandwidth management and utilization between a pair of routers. TE, hence, relieves

temporary congestion in the core of the network on the primary or optimal cost links. TE maps flows between two routers appropriately to enable efficient use of already available bandwidth in the core of the network. The key to implementing a scalable and efficient TE methodology in the core of the network is to gather information on the traffic patterns as they traverse the core of the network so that bandwidth guarantees can be established. As illustrated in Figure 9-2, TE tunnels, *Tunnel1 and Tunnel2,* can be configured on PE1-AS1 that can map to separate paths *(PATH1, PATH2),* enabling efficient bandwidth utilization.

TE tunnels configured on routers are unidirectional. Therefore, to implement bidirectional TE deployment between routers PE1-AS1 and PE2-AS1 in Figure 9-2, a pair of tunnels must also be configured on PE2-AS1 in addition to *Tunnel1 and Tunnel2* configured on PE1-AS1. In an MPLS network, all pertinent tunnel configurations are always performed or provider edge (PE) routers. The TE tunnels or LSPs will be used to link the edge routers across the core of the service provider network.

MPLS TE can also map to certain classes of traffic versus destinations. If Customer A CE routers are connected into the SP network using OC3 links versus Customer B connecting into the SP network using a 64 K dialup link, preferential treatment can be configured on TE tunnels so that TE *Tunnel1* can carry Customer A traffic and Tunnel2 can carry Customer B traffic. This is shown in Figure 9-3. Also note that Figure 9-3 illustrates tunnels configured on both PE1-AS1 and PE2-AS1.

Figure 9-3 *TE Tunnels Based on Customer CoS*

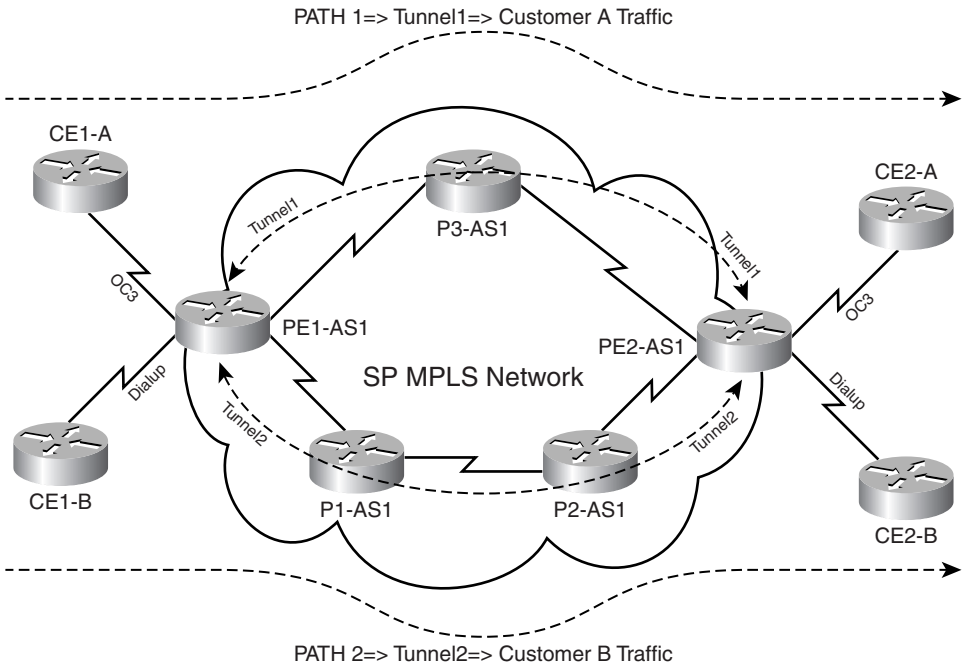

TE tunnels are, thus, data flows between a specific source and destination that might have properties or attributes associated with them. The attributes associated with a tunnel, in addition to the ingress (headend) and egress (tailend) points of the network, can include the bandwidth requirements and the CoS for data that will be forwarded utilizing this tunnel. Traffic is forwarded along the path defined as the TE tunnel by using MPLS label switching. Hence, TE tunnels are assigned specific label switched paths (LSPs) in the network from source to destination, which are usually PE routers. MPLS LSPs have a one-to-one mapping with TE tunnels, and TE tunnels are not bound to a specific path through the SP network to a destination PE router. Unless configured explicitly, TE tunnels can reroute packets via any path through the network associated with an MPLS LSP. This path might be defined by the IGP used in the core, which are discussed in the section on MPLS TE extensions.

The primary reason for the implementation of MPLS TE is to control paths along which traffic flows through a network. MPLS TE also lends itself to a resilient design in which a secondary path can be used when the primary path fails between two routers in a network. Data plane information is forwarded using label switching; a packet arriving on a PE from the CE router is applied labels and forwarded to the egress PE router. The labels are removed at the egress router and forwarded out to the appropriate destination as an IP packet.

OSPF or IS-IS with extensions for TE is used to carry information pertaining to the tunnel configured on a router. The extensions carry information on available resources for building a tunnel, like bandwidth on a link. As a result, a link that does not have the requested resources (like bandwidth) is not chosen to be a part of the LSP tunnel or TE tunnel. Signaling in an MPLS TE environment uses resource reservation protocol (RSVP) with extensions to support TE tunnel features.

The data plane ingress (headend) router in the MPLS domain requires information pertaining to the resource availability on all links capable of being a part of the MPLS TE tunnel. This information is provided by IGPs like OSPF and IS-IS due to the inherent operation of flooding information about links to all routers in the IGP domain. In IS-IS, a new TLV (type 22) has been developed to transmit information pertaining to resource availability and link status in the LS-PDUs. In OSPF, the type 10 LSA provides resource and links status information. When this information is flooded in IGP updates, the ingress (headend) router gathers information on all the available resources in the network along with the topology, which defines tunnels through the network between a set of MPLS-enabled routers.

The inspiration behind MPLS TE is *Constraint Based Routing (CBR)*, which takes into account the possibility of multiple paths between a specific source/destination pair in a network. With CBR, the operation of an IP network is enhanced so the least cost routing can be implemented as well as variables to find paths from a source to destination. CBR requires an IGP, like OSPF or IS-IS, for its operation. CBR is the backbone of the TE tunnel definition and is defined on the ingress routers to the MPLS domain when implementing MPLS TE. Resource availability and link status information are calculated using a *constrained SPF* calculation in which factors such as the bandwidth, policies, and topology are taken into consideration to define probable paths from a source to destination.

CSPF calculation results with an ordered set of IP addresses that map to next-hop IP addresses of routers forming an LSP, in turn mapping to the TE tunnel. This ordered set is defined by the headend router that is propagated to other routers in the LSP. The intermediate routers, thus, do not perform the function of path selection. RSVP with TE extensions is used to reserve resources in the LSP path as well as label association to the TE tunnel. The operation of RSVP for MPLS TE is introduced in the next section.

RSVP with TE Extensions: Signaling

RSVP reserves bandwidth along a path from a specific source to destination. RSVP messages are sent by the headend router in a network to identify resource availability along the path from a specific source to destination. The headend router is always the source of the MPLS TE tunnel, and the tailend router is the router that functions as the endpoint for the TE tunnel. After the RSVP messages are sent, the status of routers in the path (resource availability) information is stored in the path message as it traverses the network. RSVP, therefore, communicates the requirements of a specific traffic flow to the network and gathers information about whether the requirements can be fulfilled by the network.

The four main messages used in implementation of RSVP for TE are the *RSVP PATH message,* the *RSVP RESERVATION message, RSVP error messages,* and *RSVP tear messages.* In MPLS TE, RSVP is used to ensure and verify resource availability, as well as apply the MPLS labels to form the MPLS TE LSP through the routers in the network:

- **RSVP PATH message**—Generated by the headend router and is forwarded through the network along the path of a future TE LSP. At each hop, the PATH message checks the availability of requested resources and stores this information. In our network, shown in Figure 9-4, the PATH message is generated by Router PE1-AS1, the headend router, and is forwarded downstream where it checks resource availability at each hop (P1-AS1 and PE2-AS1). The RSVP PATH message functions as a label request in MPLS TE domain. Because all TE domains function with downstream-on-demand label allocation mode, the request to assign a label is generated at the headend router and propagated downstream.

Figure 9-4 *RSVP PATH and RESERVATION Messages*

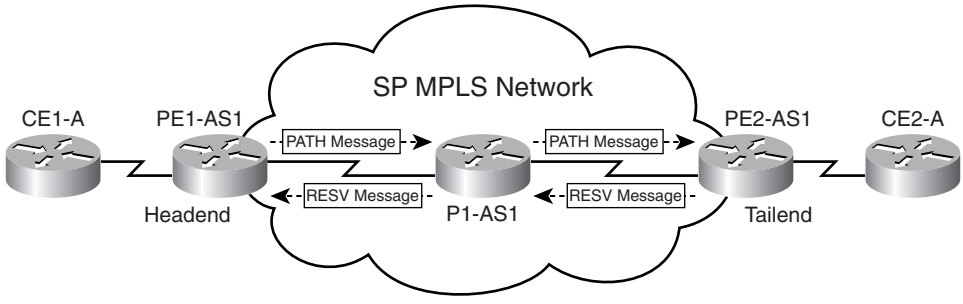

- **RSVP RESERVATION message**—Created by the tailend router in the MPLS TE domain and used to confirm the reservation request that was sent earlier with the PATH messages. In the network depicted in Figure 9-4, PE2-AS1 will generate the RSVP RESERVATION message in response to the PATH message. Therefore, PATH messages function as reservation requests and RESERVATION messages function as reservation confirmations for the availability of requested resources. The RSVP RESERVATION message performs the function of label assignment for a particular LSP mapping to the TE tunnel. As the MPLS domain label allocation and distribution is performed downstream-on-demand, the label mapping to a TE LSP is first generated by the tailend router or egress Edge LSR and then propagated upstream. This process is repeated at each hop upstream where local labels mapping to a TE tunnel are assigned and propagated upstream until the headend router is reached.
- **RSVP error messages**—In the event of unavailability of the requested resources, the router generates RSVP error messages and sends them to the router from which the request or reply was received. If Router P1-AS1 is unable to accommodate requested resources as defined in the PATH message generated by PE1-AS1 (headend router), the router generates a PATH ERROR (PATHERR) message and sends it to its upstream LSR PE1-AS1, as depicted in Figure 9-5.

Figure 9-5 *RSVP PATH Error and RESERVATION Error Messages*

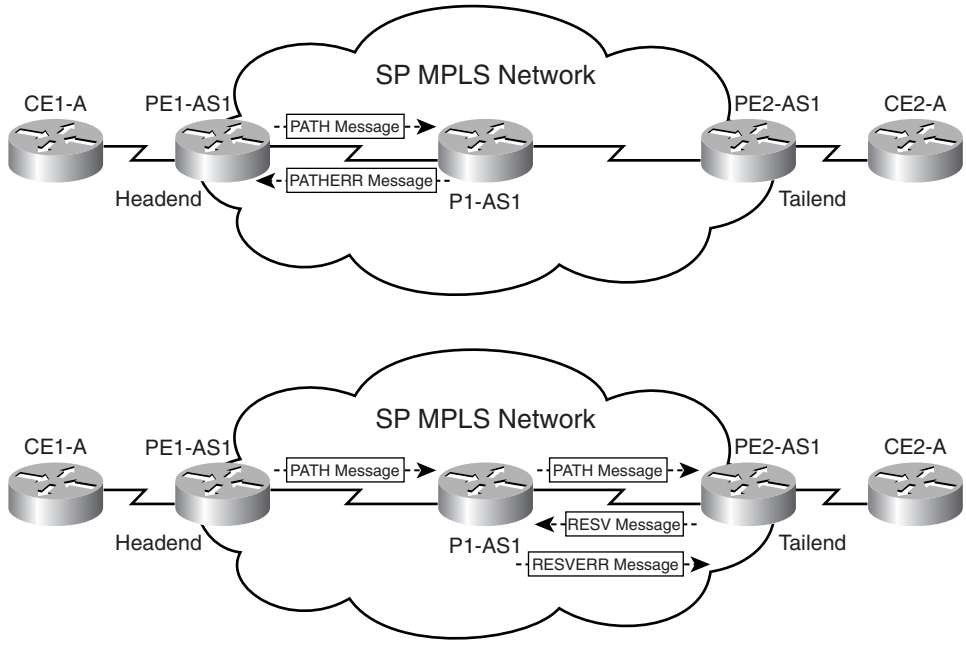

If the RSVP PATH message successfully reaches the tailend router, the tailend Router PE2-AS1 generates a RESERVATION message. If in the time lapsed between P1-AS1 receiving the PATH message from PE1-AS1 to receiving the RESERVATION message from PE2-AS1, P1-AS1 identifies a lack of resources to confirm the request, P1-AS1 will send a RESERVATION ERROR (RESVERR) message to its downstream LSR PE2-AS1 denying the reservation, as depicted in Figure 9-5.

- **RSVP tear messages**—RSVP creates two types of tear messages, namely, the PATH tear message and the RESERVATION tear message. These tear messages clear the PATH or RESERVATION states on the router instantaneously. The process of clearing a PATH or RESERVATION state on a router using tear messages enables the reuse of resources on the router for other requests. The PATH tear messages are usually generated in inter-area LSP creation where the inter-area LSP is not configured to be fast reroutable, and if a link failure occurs within an area, the LSR to which the failed link is directly attached will generate an RSVP PATH error and an RESV tear message to the headend. The headend will then generate an RSVP PATH tear message. The corresponding path option will be marked as invalid for a certain amount of time and the next path option will be immediately evaluated if it exists.

RSVP Operation in MPLS TE

As mentioned earlier, the result of a CSPF or CBR calculation on the headend router is an ordered list of IP addresses that identifies the next hops along the path of the TE tunnel or LSP. This list of routers is computed and is known only to the headend router that is the source of the TE tunnel. Other routers in the domain do not perform a CBR calculation. The headend router provides information to the routers in the TE tunnel path via RSVP signaling to request and confirm resource availability for the tunnel. RSVP with extensions for TE reserves appropriate resources on each LSR in the path defined by the headend router and assigns labels mapping to the TE tunnel LSP.

The RSVP extensions to enable RSVP use for signaling in an MPLS environment to implement TE are defined in Table 9-1. The functions of each of these extensions/objects in the messages are also outlined.

Table 9-1 *RSVP Objects*

Object	Message	Function
LABEL_REQUEST	PATH	Used to request a label mapping to the TE tunnel or LSP; generated by the headend router in the PATH message.
LABEL	RESERVATION	Used to allocate labels mapping to the TE tunnel or LSP; generated by the tailend router in the RESERVATION message and propagated upstream.

Table 9-1 *RSVP Objects (Continued)*

Object	Message	Function
EXPLICIT_ROUTE	PATH	Carried in PATH messages and is used to either request or confirm a specific path/route for the tunnel.
RECORD_ROUTE	PATH, RESERVATION	Similar to a record option with ICMP ping. It is added to the PATH or RESERVATION messages to notify the originating node about the actual route/path that the LSP TE tunnel traverses.
SESSION_ATTRIBUTE	PATH	Used to define specific session parameters local to the TE LSP tunnel.

During the path setup process for LSP TE tunnels, RSVP messages containing one or more of these extensions are used to identify the significance of each message type and its contents.

The path message contains the information outlined in Table 9-2.

Table 9-2 *RSVP Objects in Path Message*

Object	Message
SESSION	Defines the source and the destination of the LSP tunnel. Usually identified by IP addresses of corresponding loopback interfaces on headend and tailend routers.
SESSION_ATTRIBUTE	Defines the characteristics of the specific LSP tunnel, such as the bandwidth requirements and resources that would need to be allocated to the tunnel.
EXPLICIT_ROUTE	Populated by the list of next hops that are either manually specified or calculated using constraint-based SPF. The previous hop (PHOP) is set to the router's outgoing interface address. The Record_Route (RRO) is populated with the same address as well.
RECORD_ROUTE	Populated with the local router's outgoing interface address in the path of the LSP tunnel.
SENDER_TEMPLATE	In addition to the previously mentioned attributes, the sender template object in the path message depicts the interface address that will be used as the LSP-ID for the tunnel. This value is defined by the headend router.

The steps in the PATH and RESV message propagation in Figure 9-6 are depicted here:

Step 1 The appropriate values for the fields mentioned in Table 9-1 applied by the headend Router PE1-AS1 and the PATH message is sent to the next-hop router in the LSP tunnel path.

Step 2 When P1-AS1 receives this PATH message, the router checks the EXPLICIT_ROUTE object to see if the next hop is a directly connected network. This is checked in the *L-bit* of the RSVP path message. If the L-bit is set, the local router is not directly connected to the next hop in the LSP tunnel path. Therefore, the router would perform a constrained-SPF calculation to identify the next hop in the tunnel path.

If the L-bit is unset, the Router P1-AS1 knows that it is directly connected to the next hop in the LSP tunnel path. It then removes all entries in the EXPLICIT_ROUTE mapping to the local router (P1-AS1) and forwards the PATH message to the next hop as defined in the EXPLICIT_ROUTE object. In addition, P2-AS1 updates and appends the RECORD_ROUTE object to depict the local outgoing interface in the path of the LSP tunnel. Figure 9-6 depicts the PATH message values as the PATH message is forwarded from P1-AS1 to P2-AS1 after the appropriate values are updated. As previously mentioned, P1-AS1 removes references to its local interface in the EXPLICIT_ROUTE object and adds the outgoing interface in the RECORD_ROUTE object.

Figure 9-6 *RSVP PATH/RESERVATION Messages and Object Values*

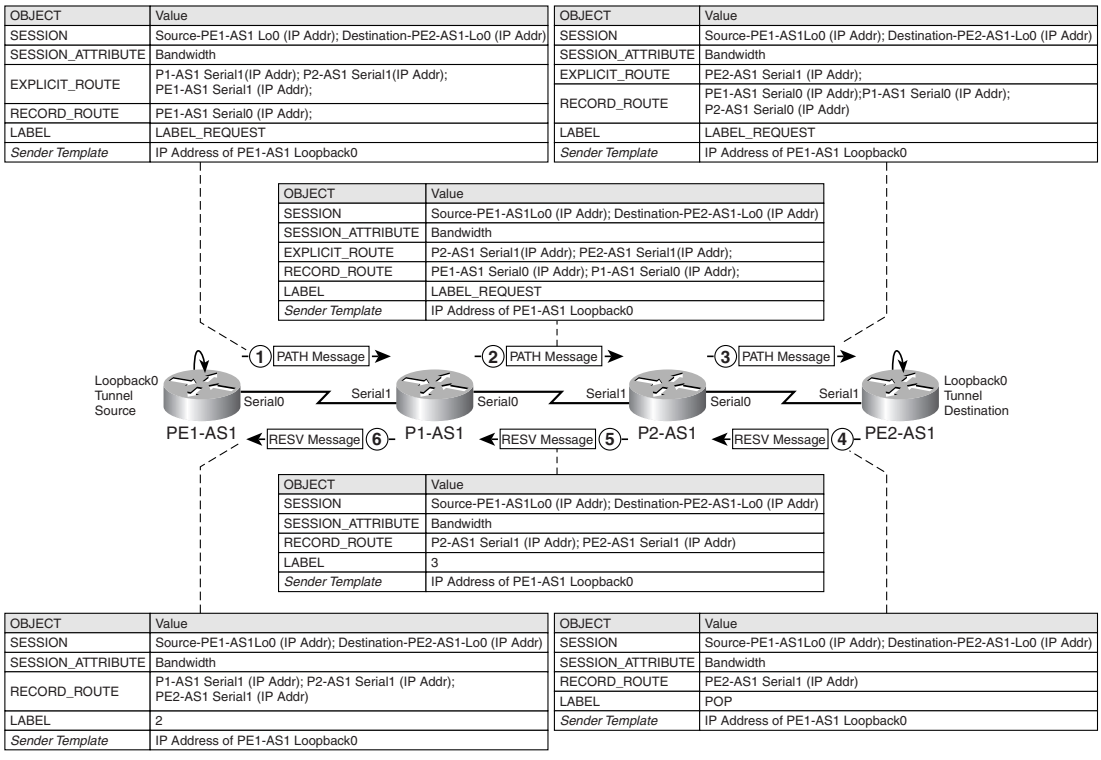

Step 3 The process is repeated at P2-AS1 in which references to its local interface in the EXPLICIT_ROUTE object are removed and appends the outgoing interface in the RECORD_ROUTE object.

Step 4 After the RSVP PATH message is received by the tailend Router PE2-AS1, it triggers the creation of a RESERVATION message. The key concept to note is that the *label allocation* process begins at the tailend router upon generation of the RESERVATION message upstream. Therefore, when PE2-AS1 generates a RESERVATION message, the router assigns a POP label to the LSP tunnel (penultimate hop popping). The RESERVATION message now has the RECORD_ROUTE object pointing to the outgoing interface on the tailend router toward the headend router. Therefore, the RECORD_ROUTE object is reinitiated in the RESERVATION message. The values are depicted in Figure 9-6.

Step 5 When this reservation message reaches P2-AS1, the RECORD_ROUTE is *prepended* with the outgoing interface and the local label mapping to the LSP is also generated and mapped in the LABEL object. An arbitrary value of 3 has been depicted for this LABEL value in Figure 9-6.

Step 6 This process is again repeated on P1-AS1 and the RESERVATION message is then received by PE1-AS1.

Step 7 When PE1-AS1 receives the RESERVATION message, the RECORD_ROUTE identifies the traffic engineered LSP associated to a specific bandwidth or resource requirement as defined in the SESSION object. The labels mapped to the LSP are thus used as in regular MPLS in which a local label is mapped to a next-hop label at each router that now maps to an RSVP-learned TE LSP versus a normal LSP.

In the implementation of RSVP for MPLS TE, RSVP with extensions for TE requests as well as confirms the LSP, reserves resources as requested on all LSP path routers, and applies MPLS labels to form the MPLS LSP through the network. Note that the routers store a copy of the PATH request as the request is forwarded to the next-hop LSR. This information identifies the interface as reservation messages are received on the same LSR to an egress interface to the headend router. In the next section, you will be introduced to the constraint-based SPF calculation process and the need for a link-state protocol to enable MPLS TE dynamically in a service provider core.

Constraint-Based Routing and Operation in MPLS TE

The most important requirement of TE is that the characteristics, as well as resource availability, on links on the network (in addition to bandwidth that would be used for cost computations) be propagated across the network to allow efficient choice of possible TE LSP paths. In link-state routing protocols, the preferred path still predominantly takes into consideration the bandwidth on the link between any two routers to compute the cost or

metric associated with that path, prior to preferred path allocation. Enabling the use of link-state routing protocols to efficiently propagate information pertaining to resource availability in their routing updates is performed by additional extensions to the actual operation of the link-state routing protocol. The mechanics of operation of a link-state routing protocol involves the flooding of updates in the network upon link-state or metric change or, in better terms, bandwidth availability from a TE perspective. The resource attributes are flooded by the routers in the network to make them available by the headend router in the TE tunnel during LSP path computation (dynamic tunnels). Link-state announcements carry information that lists that router's neighbors, attached networks, network resource information, and other relevant information pertaining to the actual resource availability that might be later required to perform a constraint-based SPF calculation. OSPF and IS-IS have been provided with extensions to enable their use in an MPLS TE environment to propagate information pertaining to resource availability and in dynamic LSP path selection.

Maximum Versus Available Bandwidth

Available bandwidth (AB) is a key value taken into consideration during the LSP path computation process to identify the preferred path for the TE tunnel. The available bandwidths on interfaces are configured on a priority basis. The number of priorities that can be configured in conjunction with the available bandwidth is 8: 0–7, where 0 represents the *highest* priority. *When the available bandwidth for a certain priority level on an interface is configured, it is subtracted from the available bandwidth on all priority levels below the one it is configured on.*

If Router PE1-AS1 has a serial interface (T1-1.544 Mbps), 1 Ethernet interface (10 Mbps) and one Fast Ethernet interface (100 Mbps), the actual bandwidths on the interfaces map to the *maximum bandwidth* (MB) values on the respective links. The *available bandwidth is usually the bandwidth of the required reservation subtracted from the maximum bandwidth.* However, this does not hold true if the available bandwidth value on the link is configured to be higher than the maximum bandwidth value on the link. Though the available bandwidth on the link can be configured to be higher than the max-bandwidth value, reservations exceeding the maximum bandwidth value are rejected.

When Router PE1-AS1 initially propagates information on the maximum bandwidth and the available bandwidth on all its links, the values for the available bandwidth at each priority level (P) for each link would be equal to their maximum bandwidth values (1.544 Mbps for serial, 10 Mbps for Ethernet, and 100 Mbps for Fast Ethernet).

When a tunnel request is accepted and the bandwidth deducted from the available bandwidth at a certain priority, it is also deducted from all the priorities lower than the priority at which the resource request was performed. If an LSP tunnel creation on PE1-AS1 consumes 40 Mbps of bandwidth on the Fast Ethernet interface at a priority level of 5, the available bandwidth values at the appropriate priorities on the Fast Ethernet interface would change for priorities 5 and above (100 – 40 = 60 Mbps).

Let us now consider the following sequence of requests:

1. Request for 10 Mbps of bandwidth on Ethernet interface at priority 1
2. Request for 20 Mbps of bandwidth on Fast Ethernet interface at priority 0
3. Request for 1 Mbps of bandwidth on serial interface at priority 0
4. Request for 2 Mbps of bandwidth on Ethernet interface at priority 3

This sequence will reduce the AB values, as depicted in Table 9-3.

Table 9-3 *PE1-AS1: Maximum Bandwidth and Available Bandwidth—All Interfaces*

Interface	AB P = 0 (Mbps)	AB P = 1 (Mbps)	AB P = 2 (Mbps)	AB P = 3 (Mbps)	AB P = 4 (Mbps)	AB P = 5 (Mbps)	AB P = 6 (Mbps)	AB P = 7 (Mbps)
Serial	1.544 – 1 = .544	1.544 – 1 = .544	1.544 – 1 = .544	1.544 – 1 = .544	1.544 – 1 = .544	1.544 – 1 = .544	1.544 – 1 = .544	1.544 – 1 = .544
Ethernet	10	10 – 10 = 0	10 – 10 = 0	10 – 10 = 0	10 – 10 = 0	10 – 10 = 0	10 – 10 = 0	10 – 10 = 0
Fast Ethernet	100 – 20 = 80	100 – 20 = 80	100 – 20 = 80	100 – 20 = 80	100 – 20 = 80	60 – 20 = 40	60 – 20 = 40	60 – 20 = 40

The outputs of Table 9-3 do not reflect the request for 2 Mbps of bandwidth on the Ethernet interface at priority 3. This request is rejected due to unavailable bandwidth at this priority level on the interface when the request is received. Link-state updates pertaining to resource availability are flooded when the status of the link changes, during manual reconfiguration of parameters mapping to the resource availability on the link, periodic updates on links and their status, and when the LSP path setup fails due to unavailability of requested resources for the LSP TE tunnel.

If the resources pertaining to the link change constantly, it will trigger update generation, which clearly must be avoided. During the instant when the resources pertaining to the links change constantly, the headend router might view the link as a probable link in the LSP path. Therefore, this probable nonupdated link might be used in path computation even though the link might not have the resources required for LSP path setup. However, after LSP path computation when the LSP path establishment is attempted, the router containing the link with the unavailable resources generates an update with information affirming a lack of resources.

Thresholds can be set up on a per interface or link basis on a router whereby updates are generated within a configured range of resource availabilities. Therefore, the upper limit, as well as the lower limit, when an update will be generated on the router containing the link, can be configured. For example, if the lower limit was configured to be 50% of link bandwidth with steps at 60, 70, 80, and 90 with the upper limit configured at 100%, updates with regards to link resource availability are generated and flooded in the network when 50%, 60%, 70%, 80%, 90%, and 100% of bandwidth are achieved.

Constraint-Based SPF

In the normal SPF calculation process, a router places itself at the head of the tree with shortest paths calculated to each of the destinations, only taking into account the least metric or cost route to the destination.

During regular SPF operation in the network, illustrated in Figure 9-7, only the cost is taken into consideration, and the least cost path from a loopback on PE1-AS1 to a loopback on PE2-AS1 is PE1-AS1->P1-AS1->PE2-AS1. In this calculation, a key concept to note is no consideration to the bandwidth of the links on the other paths from PE1-AS1 to PE2-AS1, namely via routers P3-AS1->P4-AS1 and P2-AS1. The bandwidth of the links is shown as an ordered pair in Figure 9-7 with the first value showing the cost of the link and the second showing the bandwidth across the link.

Figure 9-7 *SPF*

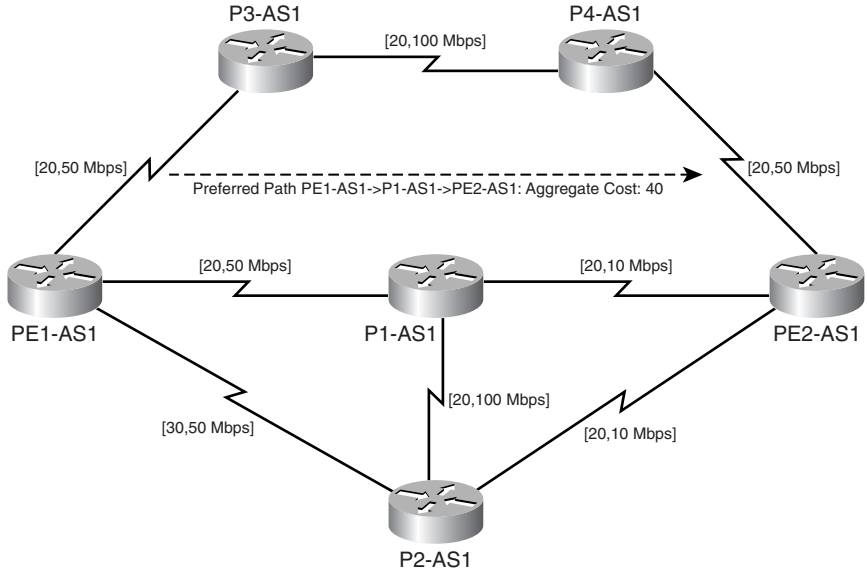

If the parameters chosen for the preferred path are not the least cost alone but also a requirement to support a bandwidth of 50 Mbps in Figure 9-7, we can eliminate the links that do not allow for the mentioned requirement. The network capable of supporting the requirement would look like what's shown in Figure 9-8.

With the just mentioned constraints, the only path capable of being used as an LSP for TE is the path from PE1-AS1 to PE2-AS1 via P3-AS1 and P4-AS1. If any of the links between P1-AS1, P2-AS1, and PE2-AS1 were to support a bandwidth more than the requirement, they would become a part of the CSPF tree structure with Router PE1-AS1 or the headend router as the root of the tree.

Figure 9-8 *CSPF*

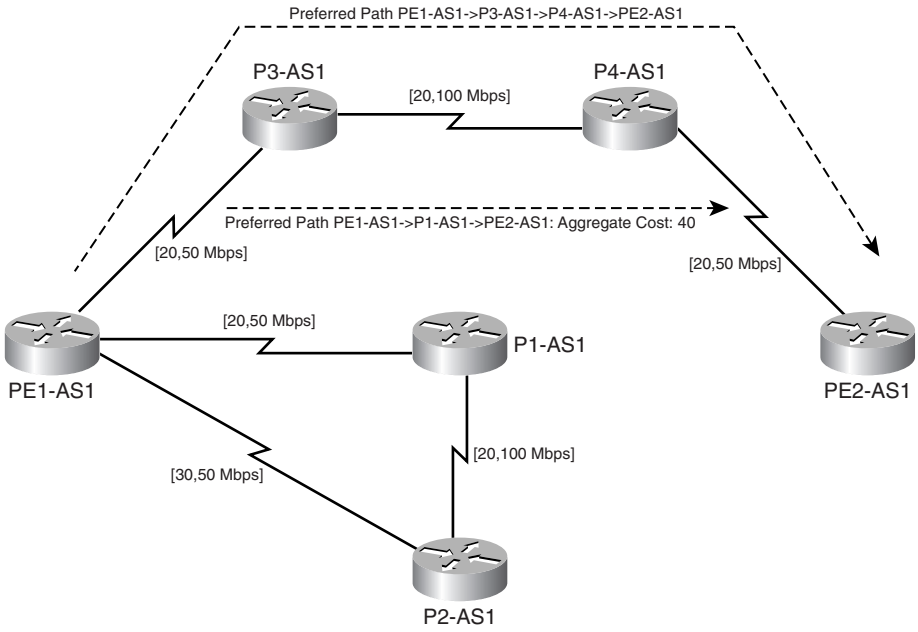

With CSPF, we use more than the link cost to identify the probable paths that can be used for TE LSP paths. The decision as to which path is chosen to set up a TE LSP path is performed at the headend router after ruling out all links that do not meet a certain criteria, such as bandwidth requirements in addition to the cost of the link. The result of the CSPF calculation at the headend router is an ordered set of IP addresses that maps to the next-hop addresses of routers that form the TE LSP. Therefore, multiple TE LSPs could be used by the use of CSPF to identify probable links in the network that meet the criteria. In addition, the user can configure a static TE tunnel or LSP on the headend router that outlines the next hops in the TE LSP path and, therefore, can use the statically defined LSP as the backup LSP path in the event of the primary TE LSP failing.

The result of the CSPF calculation is then passed over to the RSVP process to begin the RSVP request and reservation process, as mentioned in the earlier section. RSVP thus is used along with the result computed by CSPF or list of next hops configured by the user for LSP signaling and final establishment of the TE LSP. Note the TE LSP formed as a result of this process is unidirectional in nature.

Constraint-based SPF can use either administrative weights or IGP metric (also called TE metric) during the constraint-based computation. In the event of a tie, the path with the highest minimum bandwidth takes precedence, followed by the least number of hops along the path. If all else is equal, CSPF picks a path at random and chooses the same to be the TE LSP path of preference.

Therefore, the sequence of steps in the creation of an MPLS TE tunnel LSP in the network is as follows:

Step 1 CSPF calculation is performed from the headend router based on the constraints defined in the tunnel definition and requirements. This calculation is performed by the IGP in use, either OSPF or IS-IS.

Step 2 After the LSP path is calculated using the CSPF process, the output of the CSPF process, which is an ordered set of IP addresses mapping to next-hop addresses in the TE LSP, is passed to RSVP.

Step 3 RSVP now performs the resource reservation request and confirmation on the LSP, as defined by the CSPF process, to determine if the LSP meets the requirements of the specific resources requested by the tunnel definition.

Step 4 After the RSVP process receives a reservation message, it signals that the LSP is now established.

Step 5 At this juncture, the TE tunnel is available for the IGP to use. By default, the tunnel information is not added into the routing table; however, the router can be configured so that the tunnel interface is added to the routing table. You will be introduced to the configurations involved for TE on Cisco routers in the next section.

Link admission control performs a check at each hop in the desired LSP path to see if the resources requested are available prior to TE tunnel creation. The link admission control function is performed on a per hop basis with each router in the LSP path checking resource availability. If the requested resources are available, bandwidth is reserved and the router waits for the RESERVATION message to confirm this resource allocation. If, however, the resources requested are unavailable, the IGP in use sends messages stating resource unavailability. Link admission control then informs RSVP about lack of resources, and RSVP sends PATHERR messages to the headend requesting the resources and notifying a lack of resources.

When setting up TE LSP paths in link admission control, it is important that the priorities assigned to the available bandwidths are checked. Therefore, if the requested bandwidth is in use by a lower priority session (priorities 0–7, with 0 having highest priority), the lower priority session can be *preempted*. If preemption is supported, each preempted reservation leads to creation of PATHERR and RESVERR messages because the preempted session no longer fits the profile of the resource allocation.

OSPF Extension for MPLS TE

OSPF can be used as the link-state protocol of choice in MPLS TE for resource allocation information flooding through the network by implementing OSPF extensions or *Opaque LSAs*. The type of Opaque LSA in use is defined by the flooding scope of the LSA. OSPF

also now possesses TLV and sub-TLV attributes that can be configured to propagate resource availability information in link-state routing updates.

Opaque LSAs are of Type 9, 10, and 11 and differ in the flooding scope. Type 9 LSAs are not flooded beyond the local subnet and are of link-local scope. Type 10 LSAs are not flooded beyond the ABR and have an area-local scope. Type 11 LSAs are flooded throughout the autonomous system (AS). Cisco currently supports only Type 10 LSAs that have area-local scopes and are flooded within the area.

The Type 10 LSA, which is used in MPLS TE, has a number of TLV and sub-TLV values that map to specific resources in a TE domain. Figure 9-9 depicts the TLV and sub-TLV values and the appropriate values that they map to enable OSPF use for MPLS TE.

Figure 9-9 *OSPF TLV/Sub-TLV TE Extensions*

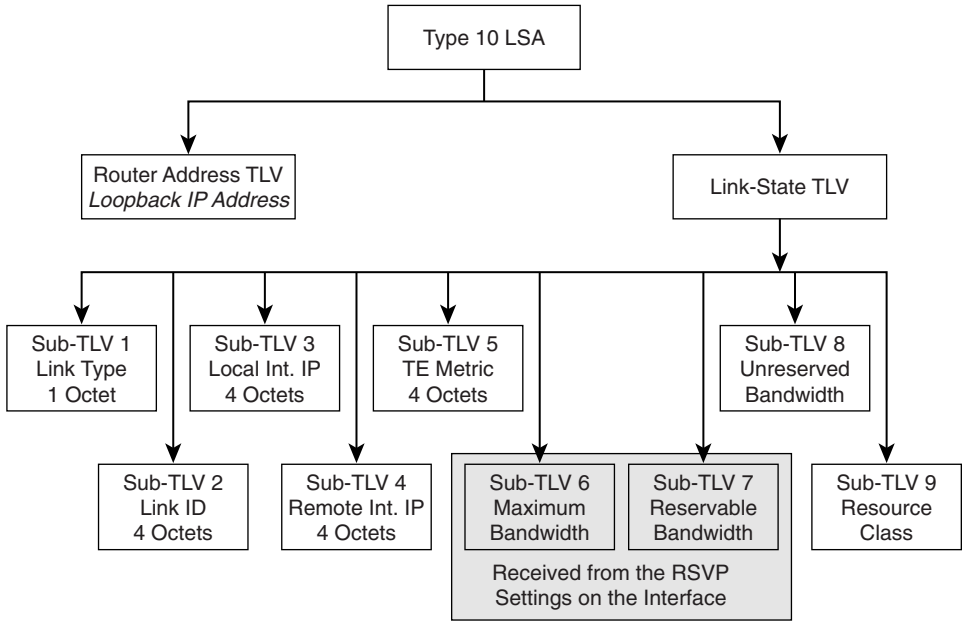

The most important sub-TLV values pertaining to TE are 6, 7, and 8. Values for sub-TLVs 6 and 7 are received from the RSVP configuration on the specific interface. Sub-TLV 8 defines the bandwidth available for reservation on each of the eight priorities. The value for sub-TLV 8 is received from the reservations active on the specific interface.

IS-IS Extensions for MPLS TE

Similar to OSPF, IS-IS can also be used as the link-state protocol of choice in the TE domain. IS-IS with extensions and newly defined TLVs can be used to propagate

information pertaining to resource allocation in an MPLS TE domain. The following TLVs have been defined for the use of IS-IS as the link-state IGP in a MPLS TE domain:

- **TLV22: Extended IS reachability**—This TLV propagates information about the state of links in the network and allows the use of "wide" metrics. In addition, this TLV provides information on resource availability, like link bandwidths.
- **TLV134: Router ID**—This TLV is used to identify the router with a distinct IP address, usually a loopback address. The source and destination IP addresses used to identify and define the tunnel endpoints must match the router ID.
- **TLV135: Extended IP reachability**—This TLV uses "wide" metrics and determines if a prefix is a level-1 or level-2 prefix. It also allows the flagging of routes when a prefix is leaked from level 2 into level 1.

In addition to the just mentioned TLVs, sub-TLVs have been defined that affix information pertaining to TE resource allocations to updates. Each sub-TLV consists of three octets except those explicitly mentioned in Figure 9-10. Most of the sub-TLVs are defined in draft-ietf-isis-traffic-xx.txt. Figure 9-10 depicts the TLVs and sub-TLVs in use by IS-IS to support MPLS TE functionality.

Figure 9-10 *IS-IS TLV/Sub-TLVs for MPLS TE*

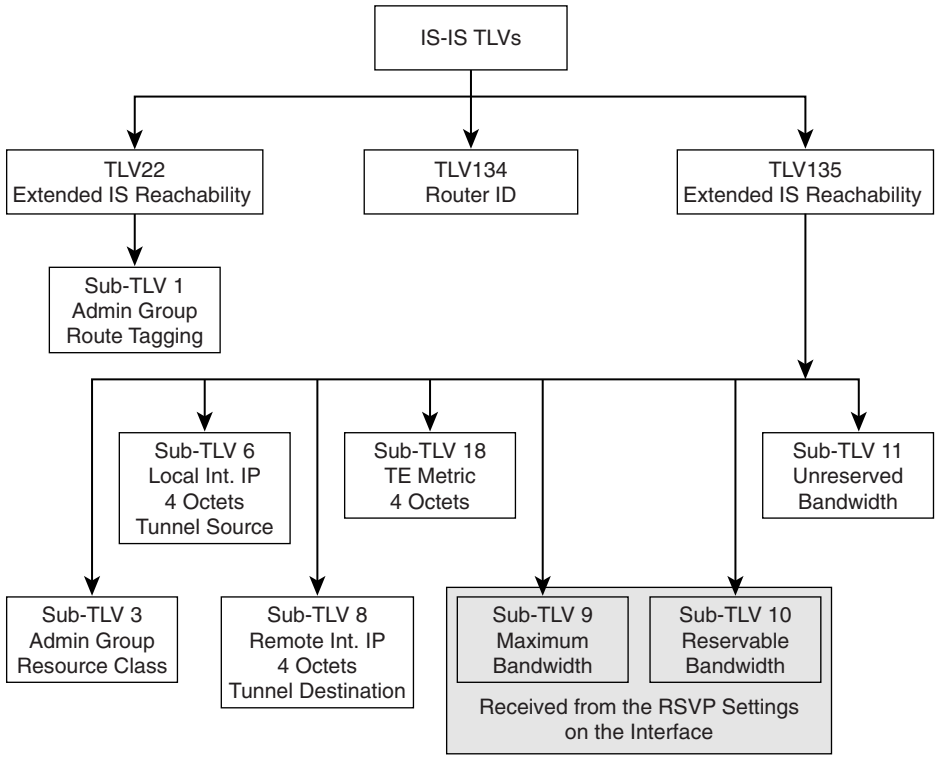

The key TLVs to note are Sub-TLV 6 and 8, which map to the tunnel endpoints or source and destination IP addresses that are usually loopback addresses; Sub-TLV 9 and 10, which map to the RSVP settings on a specific interface; and Sub-TLV 11, which maps to the unreserved bandwidth per priority on an interface after current resource allocations for active sessions have been established.

Configuring MPLS TE

This section introduces you to the steps involved in the configuration of Cisco routers to implement MPLS TE. The first subsection identifies the stepwise procedure involved in the configuration of Cisco routers for TE. It is then followed by a subsection depicting the actual configuration process on a topology consisting of six routers in which multiple paths can be used for TE purposes from a headend to tailend router.

MPLS TE Configuration Flowchart

The configuration of Cisco routers for MPLS TE support can be described in a systematic flowchart as depicted in the top row of Figure 9-11. It is assumed that the network is already configured with an IGP for NLRI exchange as well as MPLS forwarding on the appropriate interfaces prior to performing the following steps:

Step 1 Configure a loopback interface for tunnel association to the TE tunnel, as depicted in Figure 9-11.

Figure 9-11 *MPLS TE Configuration: Step 1*

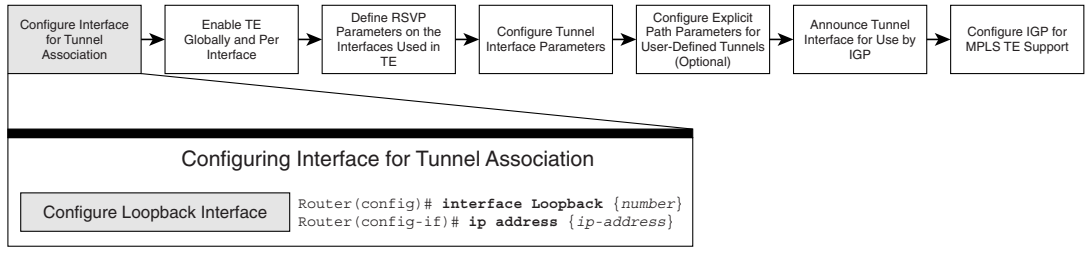

Step 2 The next step is the first configuration performed in relevance to enabling TE on the Cisco router. Figure 9-12 outlines the configurations performed on the Cisco router to enable TE functions globally on the router as well as interfaces that are possible candidates to be chosen for TE LSP paths.

Figure 9-12 *MPLS TE Configuration: Step 2*

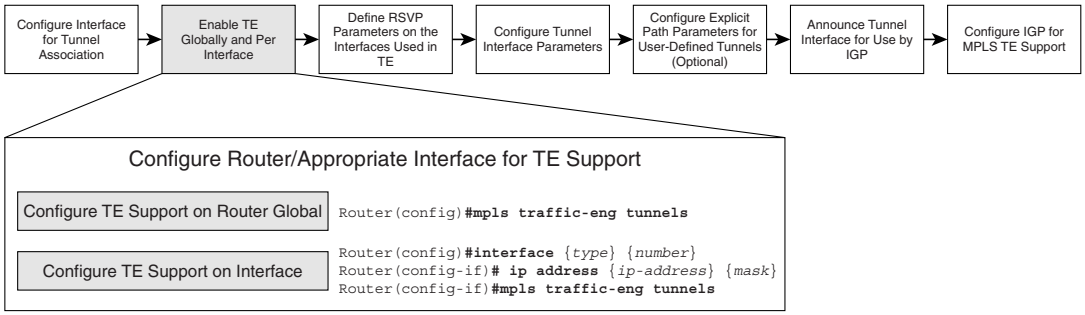

Step 3 Configure RSVP bandwidth parameters that will be used on the interface for signaling purposes and resource allocation for traffic engineered sessions. Figure 9-13 outlines the configurations that need to be performed on the interface.

Figure 9-13 *MPLS TE Configuration: Step 3*

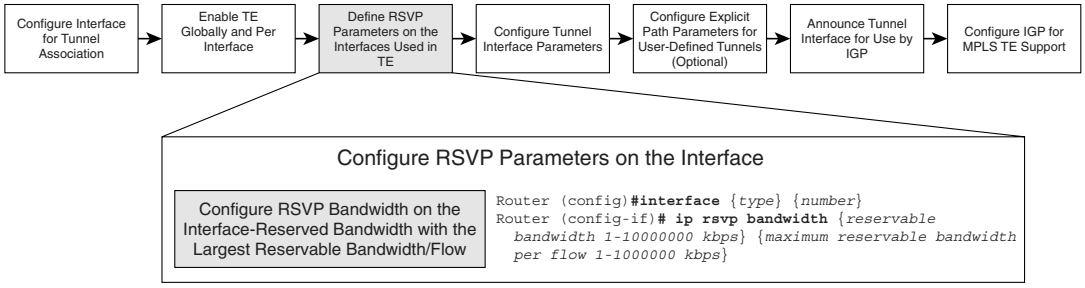

Step 4 After the interfaces that can be chosen to be a part of the LSP have been enabled for TE as well as RSVP parameters configured, the next step is to configure the tunnel interface. The main configurations of the tunnel interface would be association of the tunnel interface IP address to the loopback address configured earlier, the mode of the tunnel operation, and the destination address of the tunnel endpoint, which would map to the IP address of a loopback on the tailend router as well as the process by which the tunnel LSP path is chosen. In this step, if the path chosen for the LSP is done using the IGP and CSPF, the path option is chosen to be dynamic. Figure 9-14 depicts the configuration involved in setting up the tunnel interface.

Step 5 In addition to using the IGP for LSP path setup, the user can also define an **explicit-path** that will be used for the TE LSP. This optional step can be performed on the headend router so that the dynamic tunnel can be chosen to be the tunnel of choice for traffic forwarding and the explicit-path tunnel or user-defined static tunnel can be the backup path. In some cases, load balancing can also be achieved between the two types. Figure 9-15 depicts the configurations to set up an explicit-path LSP.

Figure 9-14 *MPLS TE Configuration: Step 4*

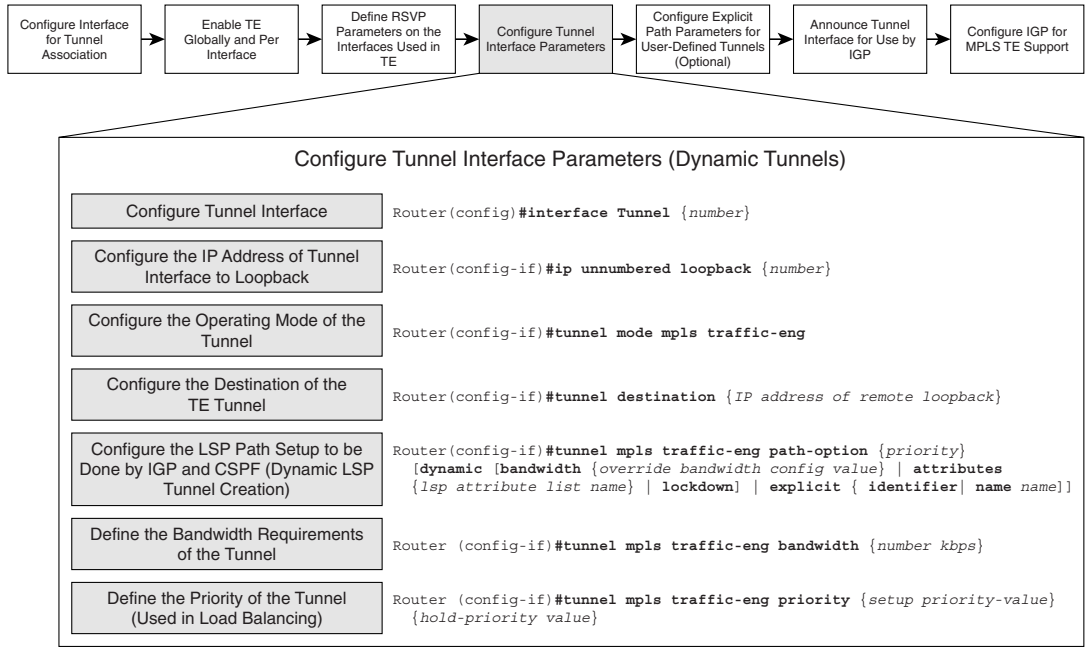

Figure 9-15 *MPLS TE Configuration: Step 5*

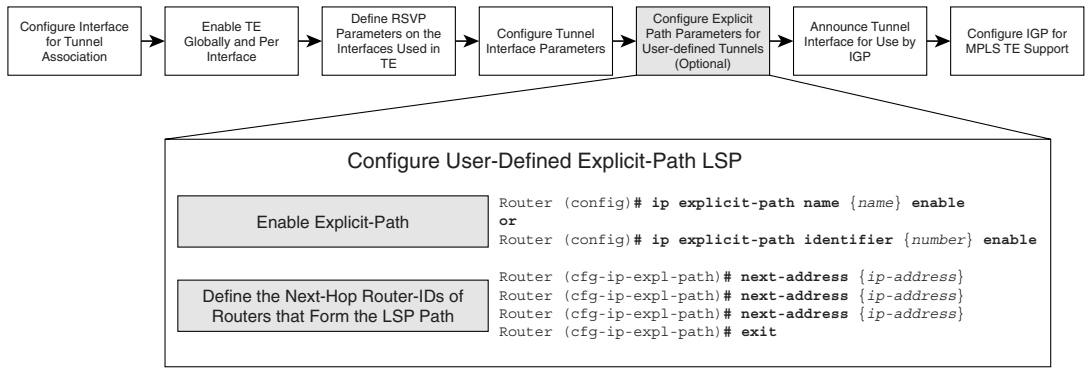

Step 6 By default, the tunnel interface is not announced into the IGP for use in the routing table. This will have to be configured explicitly for the tunnel interface to be used as the next hop in the routing table by the IGP. Figure 9-16 outlines the configurations that will have to be performed on the headend router to enable tunnel interface use as the next-hop address in the routing table for TE.

Figure 9-16 *MPLS TE Configuration: Step 6*

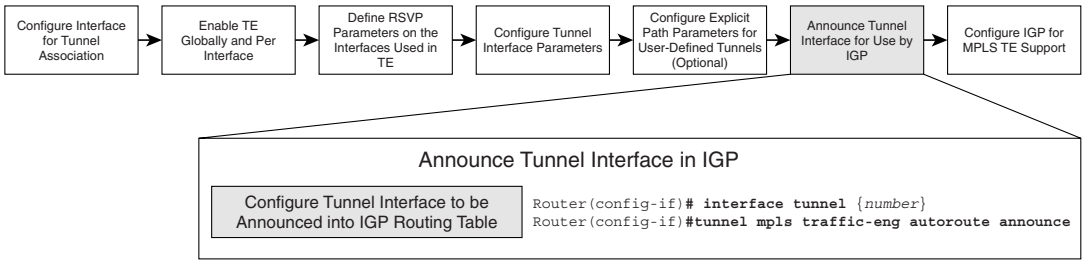

Step 7 The final step in the configuration of MPLS TE is the configuration of the IGP for TE support. The IGP in use can be either OSPF or IS-IS. The IGP process used for TE is the same as what's defined for NLRI reachability. The configurations involved for enabling TE extensions for both these protocols are outlined in Figure 9-17.

Figure 9-17 *MPLS TE Configuration: Step 7*

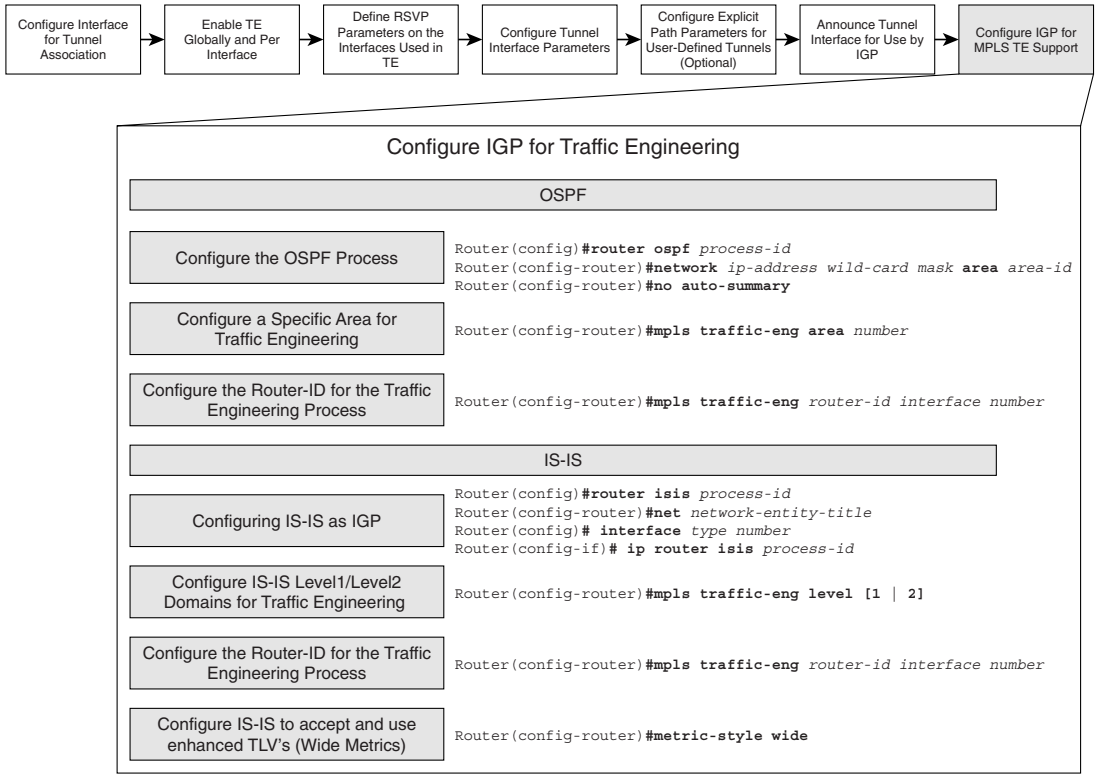

Configuring Dynamic Paths and Explicit Paths with MPLS TE

Figure 9-18 outlines the layout of the devices in the network that will be used to configure MPLS TE using dynamic and explicit paths. Prior to the following configurations, the devices shown in Figure 9-18 are configured with appropriate IP addresses on the interfaces as well as OSPF as the IGP. In addition, MPLS forwarding has been enabled on all interfaces in the network, as shown in Figure 9-18.

Figure 9-18 *MPLS TE Configuration Topology*

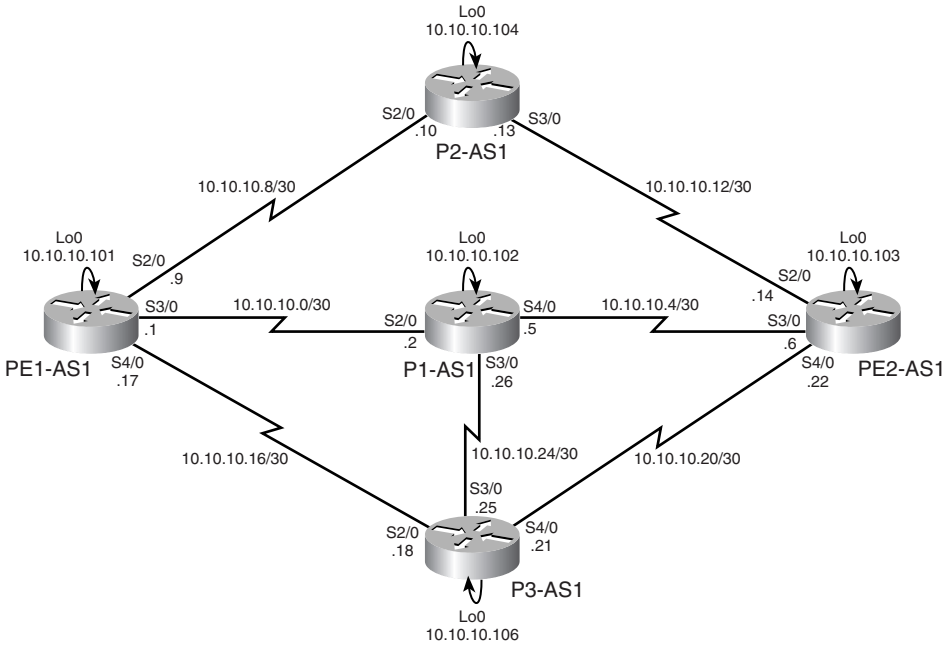

The following steps show how to configure dynamic paths and explicit paths with MPLS TE:

Step 1 Configure a loopback interface for tunnel association. This IP address can be used as the router ID in the various processes on the router (see Example 9-1).

Example 9-1 *Configure Loopback Interface for Tunnel Association*

```
PE1-AS1(config)#interface loopback 0
PE1-AS1(config-if)# ip address 10.10.10.101 255.255.255.255
```

Step 2 Enable TE globally on the router and per interface. Because we want the headend router to take all links in the network as possible links for LSP path setup, this interface-specific configuration is performed on all links

in the network topology shown in Figure 9-18. Only the configuration pertaining to PE1-AS1 is shown in Example 9-2.

Example 9-2 *Enable TE on the Router and per Interface*

```
PE1-AS1(config)#mpls traffic-eng tunnels
PE1-AS1(config)#interface serial 2/0
PE1-AS1(config-if)#mpls traffic-eng tunnels
PE1-AS1(config-if)#interface serial 3/0
PE1-AS1(config-if)#mpls traffic-eng tunnels
PE1-AS1(config-if)#interface serial 4/0
PE1-AS1(config-if)#mpls traffic-eng tunnels
```

Step 3 **Configuring RSVP bandwidth parameters**—Because we have chosen to include all interfaces in the network topology to be considered for LSP path setup, this configuration is performed on all interfaces. The chosen RSVP bandwidth configured on all interfaces is 256 K with the maximum that can be allotted to a single flow also 256 K. The configuration of headend Router PE1-AS1 is shown in Example 9-3.

Example 9-3 *Configure RSVP Parameters per Interface*

```
PE1-AS1(config)#interface serial 2/0
PE1-AS1(config-if)#ip rsvp bandwidth 256 256
PE1-AS1(config-if)#interface serial 3/0
PE1-AS1(config-if)#ip rsvp bandwidth 256 256
PE1-AS1(config-if)#interface serial 4/0
PE1-AS1(config-if)#ip rsvp bandwidth 256 256
```

Step 4 **Configuring tunnel interface parameters on the headend router**—On headend Router PE1-AS1, the tunnel destination is the loopback on Router PE2-AS1 (10.10.10.103). First, dynamic tunnels are configured in which the IGP performs a CSPF calculation and identifies the appropriate LSP path. Therefore, the path-option for this tunnel creation would be **dynamic**. The tunnel is defined with a priority of *1* and a bandwidth requirement of *100 K*. In addition, the tunnel is also provided a setup and hold priority of *1* to define that this is the most preferred tunnel LSP in the domain. See Example 9-4.

Example 9-4 *Configure Tunnel Interface Parameters on PE1-AS1*

```
PE1-AS1(config)#interface Tunnel0
PE1-AS1(config-if)# ip unnumbered Loopback0
PE1-AS1(config-if)# tunnel destination 10.10.10.103
PE1-AS1(config-if)# tunnel mode mpls traffic-eng
PE1-AS1(config-if)# tunnel mpls traffic-eng priority 1 1
PE1-AS1(config-if)# tunnel mpls traffic-eng bandwidth  100
PE1-AS1(config-if)# tunnel mpls traffic-eng path-option 1 dynamic
```

Step 5 **Configuring dynamic tunnel announcement into IGP**—In this step, the tunnel interface is configured to be added into the IGP routing table to enable traffic forwarding along the tunnel. See Example 9-5.

Example 9-5 *Announce Tunnel Interface into IGP*

```
PE1-AS1(config)#interface Tunnel0
PE1-AS1(config-if)#tunnel mpls traffic-eng autoroute announce
```

Step 6 **Configure explicit-path tunnel**—In this step, an explicit-path tunnel named *LSP1* is configured via P2-AS1 between PE1-AS1 and PE2-AS1. Configure the tunnel interface with appropriate parameters. The tunnel is configured with association to the same loopback address as used earlier with the same destination address on PE2-AS1. The resource requirements of the tunnel are also maintained. However, the tunnel priorities are configured to be 2 versus 1 in the prior dynamic tunnel configuration so that the dynamic tunnel is not chosen over the explicit tunnel for establishing primary LSP. Also, the path-option maps to the name of an explicit-path are configured on the headend router that maps to next-hop addresses in the LSP path. See Example 9-6.

Example 9-6 *Configure Tunnel Interface on PE1-AS1*

```
PE1-AS1(config)#interface Tunnel1
PE1-AS1(config-if)# ip unnumbered Loopback0
PE1-AS1(config-if)# tunnel destination 10.10.10.103
PE1-AS1(config-if)# tunnel mode mpls traffic-eng
PE1-AS1(config-if)# tunnel mpls traffic-eng priority 2 2
PE1-AS1(config-if)# tunnel mpls traffic-eng bandwidth  100
PE1-AS1(config-if)# tunnel mpls traffic-eng path-option 1 explicit name LSP1
```

Step 7 Configure the explicit-path with next-hop IP addresses of routers in the LSP path, as depicted in Example 9-7.

Example 9-7 *Configuration of Explicit LSP Path*

```
PE1-AS1(config)#ip explicit-path name LSP1
PE1-AS1(cfg-ip-expl-path)#next-address 10.10.10.10
Explicit Path name LSP1:
    1: next-address 10.10.10.10
PE1-AS1(cfg-ip-expl-path)#next-address 10.10.10.14
Explicit Path name LSP1:
    1: next-address 10.10.10.10
    2: next-address 10.10.10.14
PE1-AS1(cfg-ip-expl-path)#next-address 10.10.10.103
Explicit Path name LSP1:
    1: next-address 10.10.10.10
    2: next-address 10.10.10.14
    3: next-address 10.10.10.103
```

Step 8 Configure the tunnel interface to be announced into IGP to be the preferred path for traffic engineered traffic in the domain. See Example 9-8.

Example 9-8 *Announce Tunnel Interface into IGP*

```
PE1-AS1(config)#interface Tunnel1
PE1-AS1(config-if)# tunnel mpls traffic-eng autoroute announce
```

Step 9 **Enable IGP for MPLS TE**—The configurations on Router PE1-AS1 to enable OSPF for MPLS TE are shown in Example 9-9. The router ID configured under the MPLS TE module in OSPF and IS-IS is the loopback interface on the local router. This configuration needs to be performed on all routers in the TE domain.

Example 9-9 *Configure IGP for MPLS TE*

```
PE1-AS1(config)#router ospf 100
PE1-AS1(config-router)#mpls traffic-eng area 0
PE1-AS1(config-router)#mpls traffic-eng router-id loopback 0
```

Verification of MPLS TE Tunnel Creation

The following steps outline the various commands that can be entered on PE1-AS1 (after the just mentioned configuration) to determine if the TE tunnel has been created successfully on the router (headend):

Step 1 Perform a **show mpls traffic-eng tunnels brief** on the headend Routers PE1-AS1 and P1-AS1 in the LSP path, as well as the tailend Router PE2-AS1 to verify the tunnel state is up/up. The output of the command also gives us information on the LSP path in the tunnel setup. **UP IF** defines the upstream interface for the tunnel, and **DOWN IF** defines the downstream interface for the tunnel. See Example 9-10.

Example 9-10 **show mpls traffic-eng tunnels brief** *on Tunnel LSP Path Routers*

```
PE1-AS1#show mpls traffic-eng tunnels brief
Signalling Summary:
    LSP Tunnels Process:            running
    RSVP Process:                   running
    Forwarding:                     enabled
    Periodic reoptimization:        every 3600 seconds, next in 3206 seconds
    Periodic FRR Promotion:         Not Running
    Periodic auto-bw collection:    every 300 seconds, next in 206 seconds
TUNNEL NAME                     DESTINATION      UP IF     DOWN IF   STATE/PROT
PE1-AS1_t0                      10.10.10.103     -         Se3/0     up/up
PE1-AS1_t1                      10.10.10.103     -         Se2/0     up/up
Displayed 2 (of 2) heads, 0 (of 0) midpoints, 0 (of 0) tails
```

```
P1-AS1#show mpls traffic-eng tunnels brief
Signalling Summary:
    LSP Tunnels Process:            running
    RSVP Process:                   running
    Forwarding:                     enabled
    Periodic reoptimization:        every 3600 seconds, next in 2951 seconds
    Periodic FRR Promotion:         Not Running
    Periodic auto-bw collection:    every 300 seconds, next in 251 seconds
```

Example 9-10 **show mpls traffic-eng tunnels brief** *on Tunnel LSP Path Routers (Continued)*

```
TUNNEL NAME                       DESTINATION      UP IF     DOWN IF   STATE/PROT
PE1-AS1_t0                        10.10.10.103     Se2/0     Se4/0     up/up
Displayed 1 (of 1) heads, 1 (of 1) midpoints, 0 (of 0) tails
```

```
PE2-AS1#show mpls traffic-eng tunnels brief
Signalling Summary:
    LSP Tunnels Process:            running
    RSVP Process:                   running
    Forwarding:                     enabled
    Periodic reoptimization:        every 3600 seconds, next in 2857 seconds
    Periodic FRR Promotion:         Not Running
    Periodic auto-bw collection:    every 300 seconds, next in 157 seconds
TUNNEL NAME                       DESTINATION      UP IF     DOWN IF   STATE/PROT
PE1-AS1_t0                        10.10.10.103     Se3/0     -         up/up
PE1-AS1_t1                        10.10.10.103     Se2/0     -         up/up
```

Step 2 A view of the actual parameters of the tunnel can be retrieved by performing a **show mpls traffic-eng tunnels destination** *ip-address* (only Tunnel 0 depicted in Example 9-8) or a **show mpls traffic-eng tunnels tunnel** *interface-number*. As illustrated in Example 9-11, the output shows the status of the tunnel and the information about the parameters associated with the tunnel. In addition, it shows the preferred path chosen by the CSPF process under the explicit-path field in the output of the command, as shaded in Example 9-11.

Example 9-11 *MPLS TE Verification: Tunnel Parameters*

```
PE1-AS1#show mpls traffic-eng tunnels destination 10.10.10.103

Name: PE1-AS1_t0                                (Tunnel0) Destination: 10.10.10.103
  Status:
    Admin: up         Oper: up     Path: valid       Signalling: connected

    path option 1, type dynamic (Basis for Setup, path weight 20)

  Config Parameters:
    Bandwidth: 100      kbps (Global)  Priority: 1  1   Affinity: 0x0/0xFFFF
    Metric Type: TE (default)
    AutoRoute:  enabled   LockDown: disabled  Loadshare: 100      bw-based
    auto-bw: disabled
  Active Path Option Parameters:
    State: dynamic path option 1 is active
    BandwidthOverride: disabled  LockDown: disabled  Verbatim: disabled

  InLabel  :  -
  OutLabel : Serial3/0, 26
  RSVP Signalling Info:
       Src 10.10.10.101, Dst 10.10.10.103, Tun_Id 0, Tun_Instance 71
    RSVP Path Info:
      My Address: 10.10.10.101
      Explicit Route: 10.10.10.2 10.10.10.6 10.10.10.103
```

continues

Example 9-11 *MPLS TE Verification: Tunnel Parameters (Continued)*

```
      Record   Route:   NONE
      Tspec: ave rate=100 kbits, burst=1000 bytes, peak rate=100 kbits
    RSVP Resv Info:
      Record   Route:   NONE
      Fspec: ave rate=100 kbits, burst=1000 bytes, peak rate=100 kbits
  History:
    Tunnel:
      Time since created: 3 hours, 42 minutes
      Time since path change: 33 minutes, 26 seconds
    Current LSP:
      Uptime: 33 minutes, 26 seconds
```

```
PE1-AS1#show mpls traffic-eng tunnels tunnel 0

Name: PE1-AS1_t0                                (Tunnel0) Destination: 10.10.10.103
  Status:
    Admin: up         Oper: up     Path: valid       Signalling: connected

    path option 1, type dynamic (Basis for Setup, path weight 20)

  Config Parameters:
    Bandwidth: 100      kbps (Global)  Priority: 1  1   Affinity: 0x0/0xFFFF
    Metric Type: TE (default)
    AutoRoute:  enabled   LockDown: disabled  Loadshare: 100      bw-based
    auto-bw: disabled
  Active Path Option Parameters:
    State: dynamic path option 1 is active
    BandwidthOverride: disabled  LockDown: disabled  Verbatim: disabled

  InLabel  :  -
  OutLabel : Serial3/0, 26
  RSVP Signalling Info:
       Src 10.10.10.101, Dst 10.10.10.103, Tun_Id 0, Tun_Instance 71
    RSVP Path Info:
      My Address: 10.10.10.101
      Explicit Route: 10.10.10.2 10.10.10.6 10.10.10.103
      Record   Route:   NONE
      Tspec: ave rate=100 kbits, burst=1000 bytes, peak rate=100 kbits
    RSVP Resv Info:
      Record   Route:   NONE
      Fspec: ave rate=100 kbits, burst=1000 bytes, peak rate=100 kbits
  Shortest Unconstrained Path Info:
    Path Weight: 20 (TE)
    Explicit Route: 10.10.10.2 10.10.10.6 10.10.10.103
  History:
    Tunnel:
      Time since created: 3 hours, 42 minutes
      Time since path change: 33 minutes, 47 seconds
    Current LSP:
      Uptime: 33 minutes, 47 seconds
```

Step 3 Verify that the next hop to the destination IP address points to the tunnel interfaces in the IGP routing table. Only routes to network 10.l0.10.103 (destination) pointing to the tunnel interface as the next hop are shown for brevity. See Example 9-12. Because we have two tunnels configured on Router PE1-AS1 (dynamic and explicit) with the same parameters, the traffic to destination 10.10.10.103 is *load balanced equally* among the two paths, as shown in Example 9-12, because the bandwidths configured on the TE tunnels are the same. Traffic from PE1-AS1 to PE2-AS1 is *equally* load balanced across the two tunnels.

Example 9-12 *Verify Next-Hop Mapping to Tunnel Interface (Truncated)*

```
PE1-AS1#show ip route 10.10.10.103
Routing entry for 10.10.10.103/32
  Known via "ospf 100", distance 110, metric 97, type intra area
  Routing Descriptor Blocks:
  * directly connected, via Tunnel0
      Route metric is 97, traffic share count is 1
    directly connected, via Tunnel1
      Route metric is 97, traffic share count is 1
```

Step 4 By performing an extended ping to the destination loopback address on PE2-AS1, we see that the packets are load balanced across the two tunnel paths. See Example 9-13.

Example 9-13 *Extended Ping Verification for MPLS TE Tunnel Path*

```
PE2-AS1#ping
Protocol [ip]:
Target IP address: 10.10.10.103
Repeat count [5]: 2
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]: y
Source address or interface: 10.10.10.101
Type of service [0]:
Set DF bit in IP header? [no]:
Validate reply data? [no]:
Data pattern [0xABCD]:
Loose, Strict, Record, Timestamp, Verbose[none]: r
Number of hops [ 9 ]: 4
Loose, Strict, Record, Timestamp, Verbose[RV]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 2, 100-byte ICMP Echos to 10.10.10.103, timeout is 2 seconds:
Reply to request 0 (28 ms).  Received packet has options
 Total option bytes= 40, padded length=40
 Record route:
   (10.10.10.103)
   (10.10.10.6)
```

continues

Example 9-13 *Extended Ping Verification for MPLS TE Tunnel Path (Continued)*

```
    (10.10.10.2)
    (10.10.10.101) <*>
End of list

Reply to request 1 (28 ms).  Received packet has options
 Total option bytes= 40, padded length=40
 Record route:
    (10.10.10.103)
    (10.10.10.14)
    (10.10.10.10)
    (10.10.10.101) <*>
End of list
```

Final Configurations for Dynamic and Explicit Tunnels with MPLS TE

Example 9-14 and Example 9-15 outline the final configurations for all devices in Figure 9-18 for implementation of dynamic and explicit tunnels from PE1-AS1 to PE2-AS1.

Example 9-14 *Final Configurations for PE1-AS1 and PE2-AS1 to Implement Dynamic and Explicit Tunnels*

```
hostname PE1-AS1
!
ip cef
!
mpls traffic-eng tunnels
!
interface Loopback0
 ip address 10.10.10.101 255.255.255.255
!
interface Tunnel0
 ip unnumbered Loopback0
 tunnel destination 10.10.10.103
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng autoroute announce
 tunnel mpls traffic-eng priority 1 1
 tunnel mpls traffic-eng path-option 1 dynamic
 tunnel MPLS traffic-eng bandwidth 100
!
interface Tunnel1
 ip unnumbered Loopback0
 tunnel destination 10.10.10.103
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng autoroute announce
 tunnel mpls traffic-eng priority 2 2
 tunnel mpls traffic-eng path-option 1 explicit name LSP1
 tunnel MPLS traffic-end bandwidth 100
!
interface Serial2/0
 ip address 10.10.10.9 255.255.255.252
 mpls traffic-eng tunnels
```

Example 9-14 *Final Configurations for PE1-AS1 and PE2-AS1 to Implement Dynamic and Explicit Tunnels (Continued)*

```
 tag-switching ip
 fair-queue 64 256 48
 ip rsvp bandwidth 1000
!
interface Serial3/0
 ip address 10.10.10.1 255.255.255.252
mpls traffic-eng tunnels
mpls ip
ip rsvp bandwidth 1000
!
interface Serial4/0
 ip address 10.10.10.17 255.255.255.252
 mpls traffic-eng tunnels
 MPLS ip
 ip rsvp bandwidth 1000
!
router ospf 100
 mpls traffic-eng router-id Loopback0
 mpls traffic-eng area 0
 network 10.0.0.0 0.255.255.255 area 0
!
ip explicit-path name LSP1 enable
 next-address 10.10.10.10
 next-address 10.10.10.14
 next-address 10.10.10.103
!
end
```

```
hostname PE2-AS1
!
ip cef
!
mpls traffic-eng tunnels
!
interface Loopback0
 ip address 10.10.10.103 255.255.255.255
!
interface Serial2/0
 ip address 10.10.10.14 255.255.255.252
mpls traffic-eng tunnels
mpls ip
ip rsvp bandwidth 1000
!
interface Serial3/0
 ip address 10.10.10.6 255.255.255.252
 mpls traffic-eng tunnels
 mpls ip
 ip rsvp bandwidth 1000
!
interface Serial4/0
 ip address 10.10.10.22 255.255.255.252
```

continues

Example 9-14 *Final Configurations for PE1-AS1 and PE2-AS1 to Implement Dynamic and Explicit Tunnels (Continued)*

```
 mpls traffic-eng tunnels
 mpls ip
ip rsvp bandwidth 1000
!
router ospf 100
mpls traffic-eng router-id Loopback0
 mpls traffic-eng area 0
 network 10.0.0.0 0.255.255.255 area 0
!
end
```

Example 9-15 *Final Configurations for P1-AS1, P2-AS1, and P3-AS1 to Implement Dynamic and Explicit Tunnels*

```
hostname P1-AS1
!
ip cef
!
mpls traffic-eng tunnels
!
interface Loopback0
 ip address 10.10.10.102 255.255.255.255
!
interface Serial2/0
 ip address 10.10.10.2 255.255.255.252
mpls traffic-eng tunnels
 mpls ip
 ip rsvp bandwidth 1000
!
interface Serial3/0
 ip address 10.10.10.26 255.255.255.252
 mpls traffic-eng tunnels
 mpls ip
ip rsvp bandwidth 1000
!
interface Serial4/0
 ip address 10.10.10.5 255.255.255.252
 mpls traffic-eng tunnels
 mpls ip
ip rsvp bandwidth 1000
!
router ospf 100
 mpls traffic-eng router-id Loopback0
 mpls traffic-eng area 0
 network 10.0.0.0 0.255.255.255 area 0
!
end
```

```
hostname P2-AS1
!
ip cef
```

Example 9-15 *Final Configurations for P1-AS1, P2-AS1, and P3-AS1 to Implement Dynamic and Explicit Tunnels (Continued)*

```
!
mpls traffic-eng tunnels
!
interface Loopback0
 ip address 10.10.10.104 255.255.255.255
!
interface Serial2/0
 ip address 10.10.10.10 255.255.255.252
 mpls traffic-eng tunnels
 MPLS ip
 ip rsvp bandwidth 1000
!
interface Serial3/0
 ip address 10.10.10.13 255.255.255.252
 mpls traffic-eng tunnels
 mpls ip
 ip rsvp bandwidth 1000
!
router ospf 100
 mpls traffic-eng router-id Loopback0
 mpls traffic-eng area 0
 network 10.0.0.0 0.255.255.255 area 0
!
end
```

```
hostname P3-AS1
!
ip cef
!
mpls traffic-eng tunnels
!
interface Loopback0
 ip address 10.10.10.105 255.255.255.255
!
interface Serial2/0
 ip address 10.10.10.18 255.255.255.252
 mpls traffic-eng tunnels
 mpls ip
 ip rsvp bandwidth 1000
!
interface Serial3/0
 ip address 10.10.10.25 255.255.255.252
 no ip directed-broadcast
 mpls traffic-eng tunnels
 mpls ip
ip rsvp bandwidth 1000
!
interface Serial4/0
 ip address 10.10.10.21 255.255.255.252
```

continues

Example 9-15 *Final Configurations for P1-AS1, P2-AS1, and P3-AS1 to Implement Dynamic and Explicit Tunnels (Continued)*

```
 mpls traffic-eng tunnels
 mpls ip
 ip rsvp bandwidth 1000
!
router ospf 100
 mpls traffic-eng router-id Loopback0
 mpls traffic-eng area 0
 network 10.0.0.0 0.255.255.255 area 0
!
end
```

Unequal Cost Load Balancing Across Multiple TE Tunnels

In this section, we will configure another tunnel via the path PE1-AS1, P3-AS1, and PE2-AS1 with bandwidth requirements of 50 kbps versus 100 kbps. In every five packets, load balancing is performed so that two packets are sent on Tunnel 0, two on Tunnel 1, and one packet on Tunnel 2. In this case, if the source and destination of the tunnel interfaces are the same, the traffic between the sites performs unequal cost load balancing among the various tunnels between Routers PE1-AS1 and PE2-AS1. The configuration on PE1-AS1 (headend router) for another explicit LSP path setup via the path PE1-AS1, P3-AS1, and PE2-AS1 is shown in Example 9-16.

Example 9-16 *Unequal Cost Load Balancing Configuration on PE1-AS1*

```
PE1-AS1(config)#interface Tunnel2
PE1-AS1(config-if)# ip unnumbered Loopback0
PE1-AS1(config-if)# tunnel destination 10.10.10.103
PE1-AS1(config-if)# tunnel mode mpls traffic-eng
PE1-AS1(config-if)# tunnel mpls traffic-eng autoroute announce
PE1-AS1(config-if)# tunnel mpls traffic-eng priority 3 3
PE1-AS1(config-if)# tunnel mpls traffic-eng bandwidth  50
PE1-AS1(config-if)# tunnel mpls traffic-eng path-option 1 explicit name LSP2

PE1-AS1(config)#ip explicit-path name LSP2 enable
PE1-AS1(cfg-ip-expl-path)# next-address 10.10.10.18
Explicit Path name LSP2:
    1: next-address 10.10.10.18
PE1-AS1(cfg-ip-expl-path)# next-address 10.10.10.22
Explicit Path name LSP2:
    1: next-address 10.10.10.18
    2: next-address 10.10.10.22
PE1-AS1(cfg-ip-expl-path)# next-address 10.10.10.103
Explicit Path name LSP2:
    1: next-address 10.10.10.18
    2: next-address 10.10.10.22
    3: next-address 10.10.10.103
PE1-AS1(cfg-ip-expl-path)#end
```

After the configuration is performed, the output of the routing table entry for 10.10.10.103/32 shows the unequal cost load balancing in effect (see Example 9-17).

Example 9-17 *Verification of Unequal Cost Load Balancing*

```
PE1-AS1#show ip route 10.10.10.103
Routing entry for 10.10.10.103/32
  Known via "ospf 100", distance 110, metric 97, type intra area
  Routing Descriptor Blocks:
  * directly connected, via Tunnel0
      Route metric is 97, traffic share count is 2
    directly connected, via Tunnel1
      Route metric is 97, traffic share count is 2
    directly connected, via Tunnel2
      Route metric is 97, traffic share count is 1
```

Therefore, the final configuration for PE1-AS1 includes, in addition to Example 9-14, the configuration shown in Example 9-18.

Example 9-18 *Additional Configuration on PE1-AS1 for Unequal Cost Load Balancing*

```
interface Tunnel2
 ip unnumbered Loopback0
 no ip directed-broadcast
 tunnel destination 10.10.10.103
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng autoroute announce
 tunnel mpls traffic-eng priority 3 3
 tunnel mpls traffic-eng bandwidth  50
 tunnel mpls traffic-eng path-option 1 explicit name LSP2
```

MPLS TE Fast ReRoute Link Protection

Fast ReRoute (FRR) is a procedure used in conjunction with MPLS TE to reroute around a link in the case of link failure. Protection in networks can be provided by SONET, optical protection, or, more recently, MPLS FRR. With MPLS FRR, we can implement both link and node protection. In addition, different protection policies can be applied to different classes of traffic traversing the MPLS backbone. In FRR operation, a *backup tunnel* is configured to be used if the primary tunnel LSP fails. The backup tunnel must be configured so that the LSP can get to the next-hop LSR downstream without attempting to use the failed link.

The configuration for implementing FRR for link protection is simple to implement. If you use a subset of the network shown in Figure 9-18 to implement link protection, as illustrated in Figure 9-19, you can configure a backup tunnel on the LSR P1-AS1. If the primary tunnel from PE1-AS1 via P1-AS1 to PE2-AS1 fails due to link failure between P1-AS1 and PE2-AS1, the backup tunnel is used to forward traffic.

Figure 9-19 *MPLS FRR Network Topology, Configuration, and Verification*

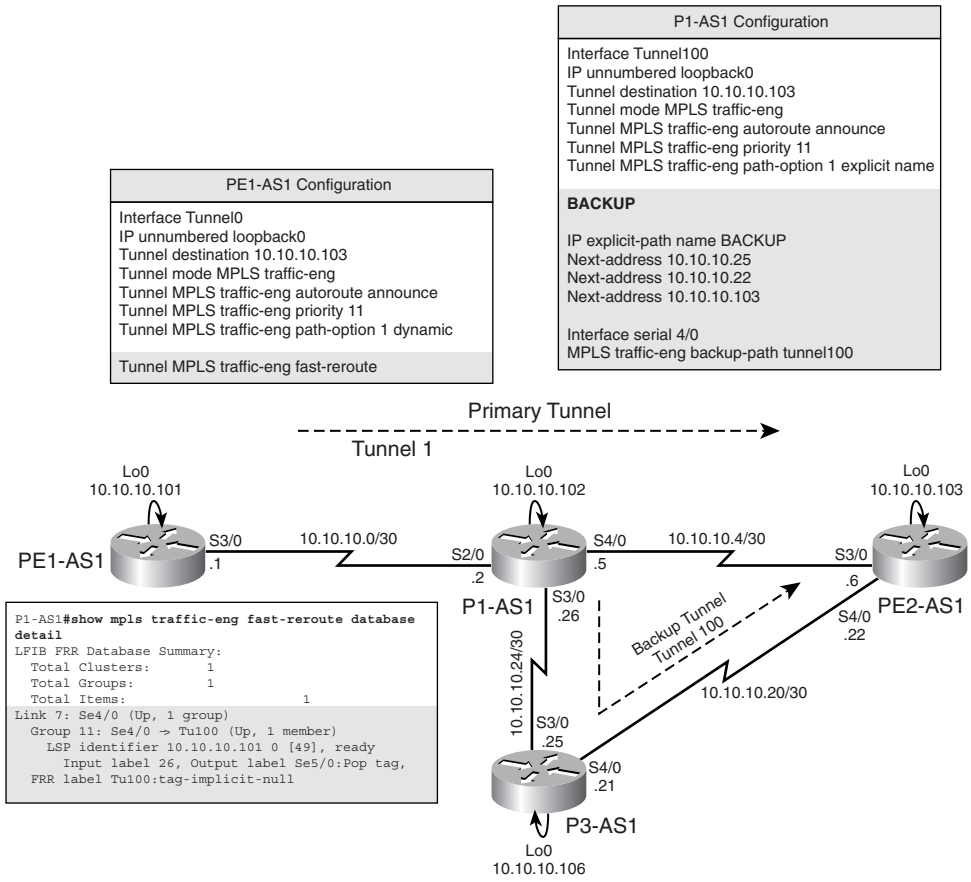

Configuration of the tunnel (Tunnel0 on PE1-AS1) to be protected from a link failure includes the **tunnel mpls traffic-eng fast-reroute** command under the tunnel interface configuration on the headend router (PE1-AS1) to enable FRR protection on the tunnel. In addition, a backup tunnel, Tunnel100, is configured on the downstream LSR (in our case, P1-AS1) to reroute traffic if the link between P1-AS1 and PE2-AS1 fails. Configuration is performed following the procedure shown in the earlier sections with an explicit path from P1-AS1 to PE2-AS1 via P3-AS1. Finally, this tunnel (Tunnel100) on P1-AS1 is associated to the link to be protected by using the command **mpls traffic-eng backup-path tunnel tunnel100** under the interface to be protected (Serial 4/0 on P1-AS1).

Verification of FRR capabilities can be performed by issuing the **show mpls traffic-eng fast-reroute database detail** command on the downstream LSR configured with a backup tunnel, as shown in Figure 9-19.

Implementing MPLS VPNs over MPLS TE

MPLS was initially adopted due to its inherent properties to deliver VPNs. However, in recent years, MPLS TE has gained popularity due to the robust TE capabilities it provides. In this section, we will discuss the configurations involved in the implementation of MPLS VPN over TE tunnels. TE tunnels can be configured between PE to PE routers as well as from PE to provider core or P routers. The configurations involved in both of these implementations of MPLS TE in the provider core are introduced. The network used to implement MPLS VPN over TE tunnels is shown in Figure 9-20.

Figure 9-20 *MPLS VPN Over TE Network Topology/Configuration*

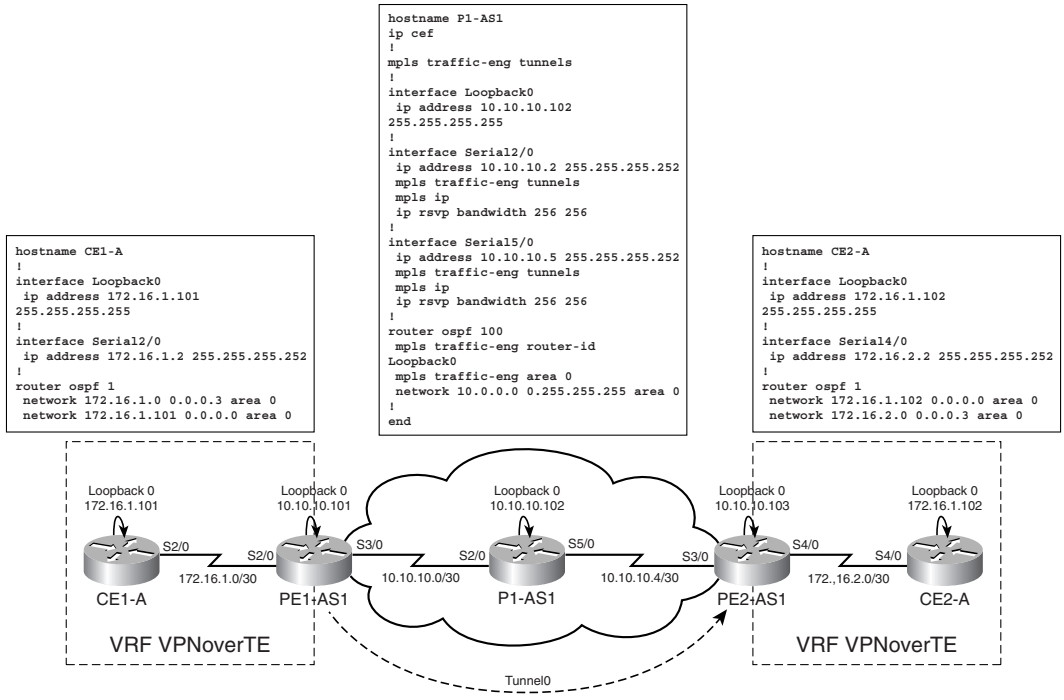

For simplicity, the OSPF PE-CE connectivity implementation is used on both PE Routers PE1-AS1 and PE2-AS1 in Figure 9-20. For this section, the IGP used in the core is OSPF with process-id 100. The process-id for the PE to CE connections is configured under OSPF 1. All networks are in area 0.

The configurations on Routers P1-AS1, CE1-A, and CE2-A are illustrated in Figure 9-20. Configurations for PE1-AS1 and PE2-AS1 are illustrated in Example 9-19. A tunnel is already configured with a dynamic path-option between PE1-AS1 and PE2-AS1.

Example 9-19 *PE1-AS1 and PE2-AS1 Configuration: MPLS VPN Over TE with PE to PE Tunnels*

```
hostname PE1-AS1
!
ip cef
!
ip vrf VPNoverTE
 rd 1:100
 route-target export 1:100
 route-target import 1:100
!
mpls traffic-eng tunnels
!
interface Loopback0
 ip address 10.10.10.101 255.255.255.255
!
interface Tunnel0
 ip unnumbered Loopback0
 tunnel destination 10.10.10.103
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng autoroute announce
 tunnel mpls traffic-eng priority 1 1
 tunnel mpls traffic-eng bandwidth  100
 tunnel mpls traffic-eng path-option 1 dynamic
!
interface Serial2/0
 ip vrf forwarding VPNoverTE
 ip address 172.16.1.1 255.255.255.252
!
interface Serial3/0
 ip address 10.10.10.1 255.255.255.252
 mpls traffic-eng tunnels
 mpls ip
 ip rsvp bandwidth 256 256
!
router ospf 1 vrf VPNoverTE
 redistribute bgp 100 metric 10 subnets
 network 172.16.1.0 0.0.0.3 area 0
!
router ospf 100
 mpls traffic-eng router-id Loopback0
 mpls traffic-eng area 0
 network 10.10.10.0 0.0.0.3 area 0
 network 10.10.10.101 0.0.0.0 area 0
!
router bgp 100
 bgp router-id 10.10.10.101
 neighbor 10.10.10.103 remote-as 100
 neighbor 10.10.10.103 update-source Loopback0
 no auto-summary
 !
 address-family vpnv4
 neighbor 10.10.10.103 activate
 neighbor 10.10.10.103 send-community extended
 exit-address-family
 !
```

Example 9-19 *PE1-AS1 and PE2-AS1 Configuration: MPLS VPN Over TE with PE to PE Tunnels (Continued)*

```
 address-family ipv4 vrf VPNoverTE
 redistribute ospf 1 vrf VPNoverTE metric 2
exit-address-family
!
end
```

```
hostname PE2-AS1
!
ip cef
!
ip vrf VPNoverTE
 rd 1:100
 route-target export 1:100
 route-target import 1:100
!
mpls traffic-eng tunnels
!
interface Loopback0
 ip address 10.10.10.103 255.255.255.255
!
interface Serial3/0
 ip address 10.10.10.6 255.255.255.252
 mpls traffic-eng tunnels
 mpls ip
 ip rsvp bandwidth 256 256
!
interface Serial4/0
 ip vrf forwarding VPNoverTE
 ip address 172.16.2.1 255.255.255.252
!
router ospf 1 vrf VPNoverTE
 redistribute bgp 100 metric 2 subnets
 network 172.16.2.0 0.0.0.3 area 0
!
router ospf 100
 mpls traffic-eng router-id Loopback0
 mpls traffic-eng area 0
 network 10.10.10.4 0.0.0.3 area 0
 network 10.10.10.103 0.0.0.0 area 0
!
router bgp 100
 bgp router-id 10.10.10.103
 neighbor 10.10.10.101 remote-as 100
 neighbor 10.10.10.101 update-source Loopback0
 !
 address-family vpnv4
 neighbor 10.10.10.101 activate
 neighbor 10.10.10.101 send-community extended
 exit-address-family
 !
 address-family ipv4 vrf VPNoverTE
 redistribute ospf 1 vrf VPNoverTE metric 2
 exit-address-family
!
end
```

Verification of MPLS VPN over TE with PE to PE Tunnels

Figure 9-21 outlines the various verification steps for identifying the operation of MPLS VPNs over TE with PE to PE tunnels.

Figure 9-21 *MPLS VPN over TE Verification—PE to PE Tunnels*

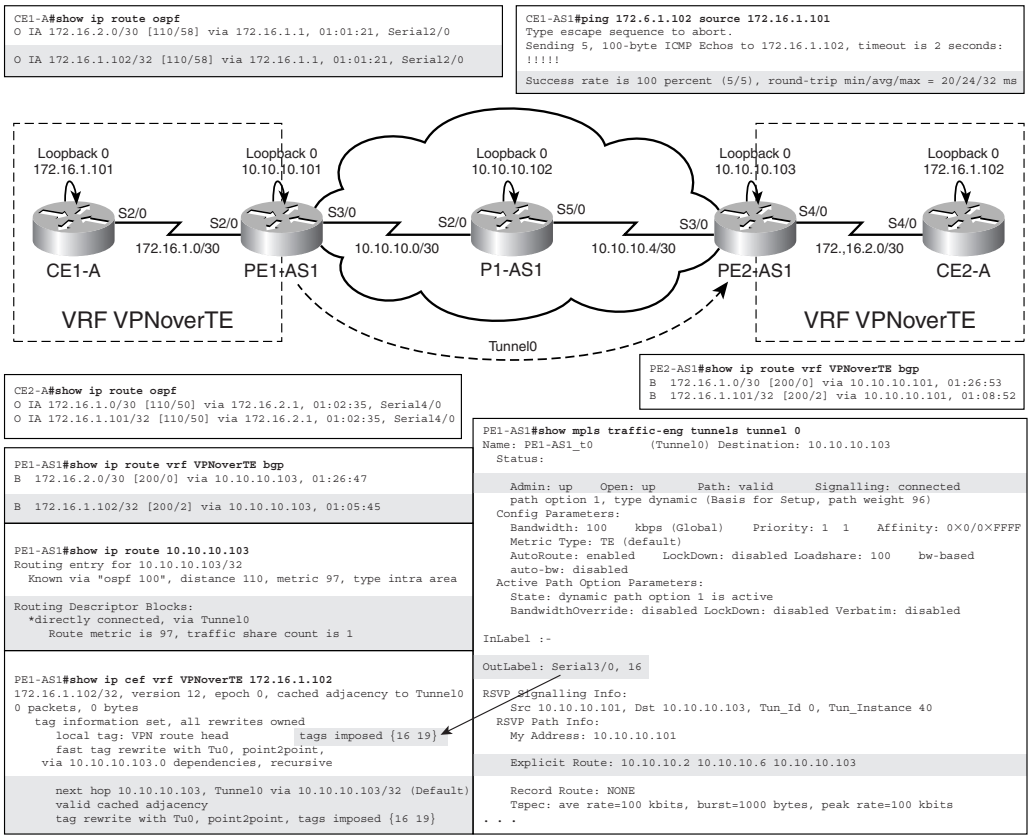

Figure 9-21 illustrates the routing tables on CE routers in which the CE routers learn the routes from the remote CEs via the MPLS backbone and place them in their local routing tables as OSPF IA routes, though all CE routes are in area 0 because sham-links are not configured.

Figure 9-21 also shows the routing table on the respective PE routers for the VRF VPNoverTE to check for route propagation in the MPLS VPN domain. As can be derived from the output, the appropriate VPN routes obtained from the remote CEs are learned from the next hop that maps to the remote PE loopback.

A closer look at the prefix 172.16.1.102 (loopback0 on CE2-A), learned across the MPLS domain one PE1-AS1, indicates a next-hop address of the remote PE loopback 10.10.10.103 (lo0 on PE2-AS1). In the global routing table, if this VPN forwards traffic

over the MPLS TE tunnel configured on PE1-AS1, the next-hop address of 10.10.10.103 must point to the tunnel interface (Tunnel0) as shown in Figure 9-21 by the output of **show ip route 10.10.10.103** on PE1-AS1. In addition, note that in the label-stack imposed on the packets in the MPLS domain when implementing MPLS VPN over TE (one label for MPLS VPN and the top label for TE), the top label maps to the label assigned by RSVP for the traffic engineered LSP path. Therefore, the out-label value in the output of **show MPLS traffic-eng tunnels tunnel0** (16) maps to the top label in the label stack, as highlighted in the output of **show ip cef vrf VPNoverTE 172.16.1.102** in Figure 9-21.

For final verification of connectivity, an extended ping is performed between loopback interfaces on CE routers, as shown in Figure 9-21.

Configuration of MPLS VPN over TE with PE to P Tunnels

In the preceding section, MPLS VPN was configured over TE tunnels in which the TE tunnel was configured between the two PE routers in the MPLS domain. Another possibility that might arise while deploying MPLS VPN over a TE enabled domain is a tunnel existing between a PE router and a provider core router. In our existing setup, the tunnel interface, Tunnel 0, configured on the PE Router PE1, is changed so that the destination of the tunnel is the loopback address on P1 or 10.10.10.102/32 (see Example 9-20). This configuration might be used in conjunction with FRR to enable link protection in the SP backbone for MPLS forwarded traffic belonging to a customer.

Example 9-20 *Configuration on PE1-AS1: Tunnel Destination Changed to 10.10.10.102/32*

```
PE1-AS1(config)#interface tunnel 0
PE1-AS1(config-if)# tunnel destination 10.10.10.102
```

If no other changes in configuration are made on any router, the CE routers no longer have connectivity to one another because the LSP is broken, as shown in Example 9-21.

Example 9-21 *CE1-AS1 Cannot Reach CE2 Because LSP Is Broken*

```
CE1-AS1#ping 172.16.1.102 source 172.16.1.101
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.1.102, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

To enable a complete LSP, MPLS is enabled on the tunnel interface on PE1-AS1. Also, P1-AS1 is configured to accept directed hellos, as shown in Example 9-22.

Example 9-22 *Enabling MPLS on the Tunnel Interface and Configuring Directed-Hello Accept on P1-AS1*

```
PE1-AS1(config)#interface tunnel 0
PE1-AS1(config-if)#mpls ip
```

```
P1-AS1(config)#mpls ldp discovery targeted-hello accept
```

Because the P1-AS1 router can accept directed hellos from neighbors who are not directly connected, the LSP is now established using the tunnel. This is shown in Figure 9-22 where the next hop for the remote CE loopback interfaces point to the interface tunnel 0 on PE1-AS1.

Figure 9-22 *MPLS VPN Over TE Verification—PE to P Tunnels*

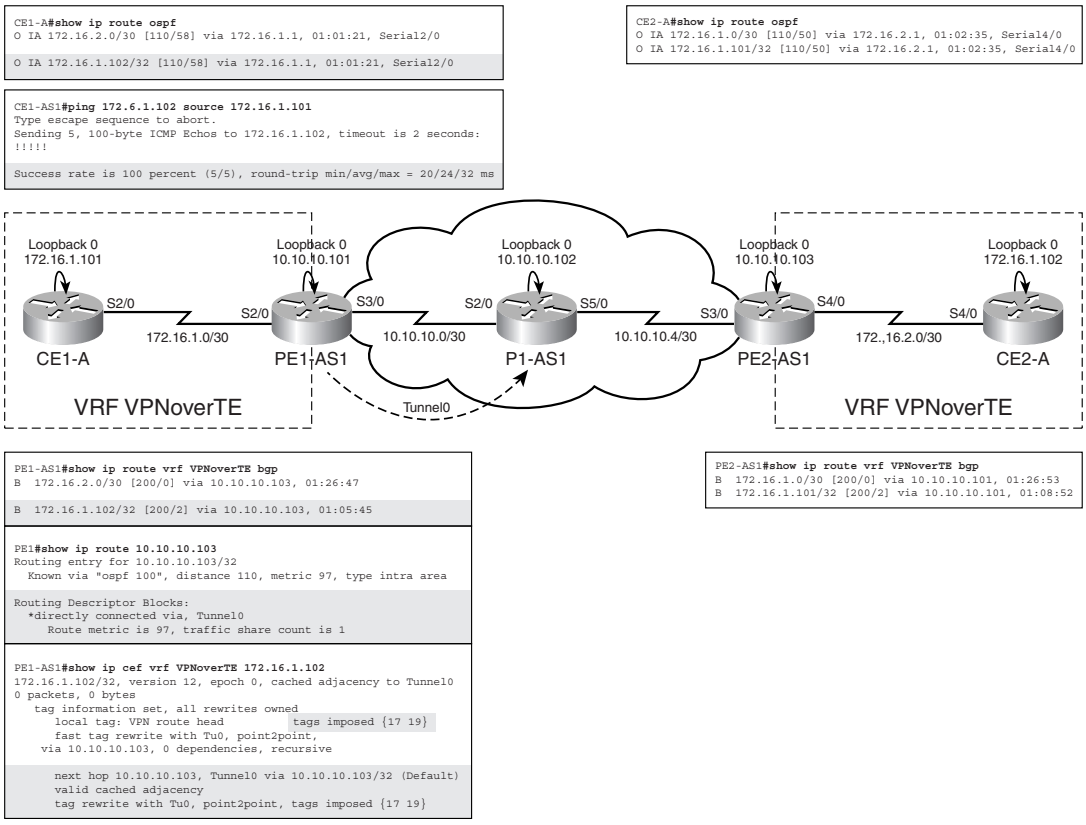

Connectivity between CE routers is verified using extended pings between loopback interfaces on CE routers, as shown in Figure 9-22.

Command Reference

Command	Description
Router(config)#**mpls traffic-eng tunnels**	Configures TE support on router in the global configuration mode.
Router(config-if)#**mpls traffic-eng tunnels**	Configures MPLS TE support per interface.
Router(config-if)# **ip rsvp bandwidth** {*reservable bandwidth 1-10000000 kbps*} {*maximum reservable bandwidth per flow 1-1000000 kbps*}	Configures RSVP bandwidth on the interface-reserved bandwidth with the largest reservable bandwidth/flow.
Router(config)#**interface tunnel** {*number*}	Configures tunnel interface.

(Continued)

Command	Description
Router(config-if)#**ip unnumbered loopback** {*number*}	Configures the loopback interface IP address to be associated with the tunnel interface under tunnel interface configuration.
Router(config-if)#**tunnel mode mpls traffic-eng**	Configures the tunnel mode to be an MPLS traffic-engineered tunnel.
Router(config-if)#**tunnel destination** {*IP address of remote loopback*}	Configures the MPLS traffic-engineered tunnel's destination or end-point.
Router(config-if)#**tunnel mpls traffic-eng path-option** {*priority*} **dynamic** [**bandwidth** {*override bandwidth config value*} | **attributes** {*lsp attribute list name*} | **lockdown**]	Configures the LSP path setup to be done by IGP and CSPF (dynamic LSP tunnel creation). The tunnels can be configured with the associated priority and attributes.
Router(config)# **ip explicit-path name** {*name*} **enable** or Router(config)# **ip explicit-path identifier** {*number*} **enable**	Configures an explicit path to be associated with a TE tunnel.
Router(cfg-ip-expl-path)#**next-address** {*ip-address*} Router(cfg-ip-expl-path)#**exit**	Configures the IP next-hop addresses for the explicit MPLS traffic engineered tunnel.
Router(config-if)#**tunnel mpls traffic-eng priority** {*setup priority-value*} {*hold-priority value*}	Defines the priority of the tunnel (used in load balancing).
Router(config-if)#**tunnel mpls traffic-eng autoroute announce**	Configures tunnel interface to be announced into IGP routing table (configured under tunnel interface configuration).
Router(config-router)#**mpls traffic-eng area** *number*	Enables OSPF for TE (under router OSPF configuration).
Router(config-router)#**mpls traffic-eng** *router-id interface number*	Configures the router ID for the TE process under OSPF or IS-IS.
Router(config-router)#**mpls traffic-eng level** [**1** | **2**]	Configures IS-IS Level1/Level2 domains for TE.
Router(config-router)#**metric-style wide**	Configures IS-IS to accept and use enhanced TLVs (wide metrics).
Router(config-if)# **tunnel mpls traffic-eng fast-reroute**	Enables the MPLS tunnel for FRR protection.
Router(config-if)# **mpls traffic-eng backup-path tunnel** {*interface-number*}	Configures the backup tunnel to be used during interface failure.