Chapter 4

# RSVP Reservation Styles in an MPLS Network

A Resource Reservation Protocol (RSVP) reservation style is a request for a bandwidth reservation that includes a set of options. The options are described in the three reservation styles: fixed filter (FF), shared explicit (SE), and wildcard filter (WF) that determine how senders, receivers, and sessions are treated. The reservation style also determines how RSVP signaling reroutes an existing label-switched path (LSP). The egress router must establish each LSP with one of the three reservation styles.

This chapter describes explicitly routed LSPs that are established using the FF or SE styles. (See Table 7.) The RSVP WF reservation style is not used for explicitly routed LSPs because of its lack of applicability for traffic engineering.

The terms *node* and *router* are used interchangeably throughout this book.

**Table 7: Checklist for RSVP Reservation Styles**

| RSVP Reservation Styles Tasks | Command or Action |
|---|---|
| **RSVP Reservation Styles Overview on page 54** | |
| **Fixed Filter Style Overview on page 55** | show configuration protocols mpls<br>show rsvp session detail<br>show mpls lsp extensive ) |
| **Shared Explicit Style Overview on page 58** | |
| **Configuring and Verifying an Adaptive LSP on page 59** | [edit]<br>edit protocol mpls<br>[edit protocols mpls]<br>set label-switched-path *lsp-path-name* adaptive<br>show<br>commit |
| **Rerouting the LSP Tunnel for the SE Reservation Style on page 62** | |
| Establish the Initial LSP Tunnel on page 62 | Not applicable. |
| Reroute an LSP Tunnel on page 63 | Not applicable. |

## RSVP Reservation Styles Overview

RSVP was originally a protocol for resource reservation. In the context of resource reservation, different reservation styles were developed to determine the degree to which resources are shared; the FF and SE reservation styles.

The FF reservation style dedicates a particular reservation to an individual sender (ingress router). This reservation style is useful for concurrent and independent traffic from different senders. When used with Multiprotocol Label Switching (MPLS), the FF reservation style allows the establishment of multiple parallel unicast point-to-point LSPs to support load balancing. It can also be used with primary and secondary paths to achieve minimal disruption to traffic. Examples of applications that use FF-style reservations are video applications and unicast applications, which both require flows that have a separate reservation for each sender.

The SE reservation style allows an explicit list of senders to share the largest bandwidth request across shared links. In an MPLS environment, this style is important for rerouting LSPs with no disruption to the flow of subscriber traffic. An example application for shared explicit reservations is an audio application in which each sender transmits a distinct data stream. Typically, only a few senders are transmitting at any one time. Such a flow does not require a separate reservation for each sender; a single reservation is sufficient.

In RSVP with traffic engineering, the ingress router can request the SE style by setting the appropriate bit in the Session Attribute object. If the Session Attribute object is present but the particular bit is *not* present, the egress router can use either style (FF or SE). All values in the Session Attribute object are advisory, so an egress router can ignore the bits when it selects a style; however, to date, this behavior has not been implemented. Selection of a style can be determined by non-support of a particular style, an explicit policy, or available resources.

In the context of traffic engineering, FF is the default reservation style. The SE style allows an LSP to share reservations which is useful when the ingress router is trying to set up an alternate path before tearing down the existing path. Clearly traffic is sent on the active path only, but from the point of view of reservations, sharing resources avoids double counting the resources.

## Fixed Filter Style Overview

The FF reservation style specifies an explicit list of senders and a distinct bandwidth reservation for each sender. The distinct bandwidth reservation is not shared with other senders, and is identified by an IP address and a local identification number (LSP_ID). Because each sender has its own particular reservation, a unique label and a separate LSP are constructed for each sender-receiver pair.

In RSVP with traffic engineering, each sender and receiver represent a different sender or receiver on a router, not necessarily different end systems. (See Figure 8).

**Figure 8: Fixed Filter Reservation Style**

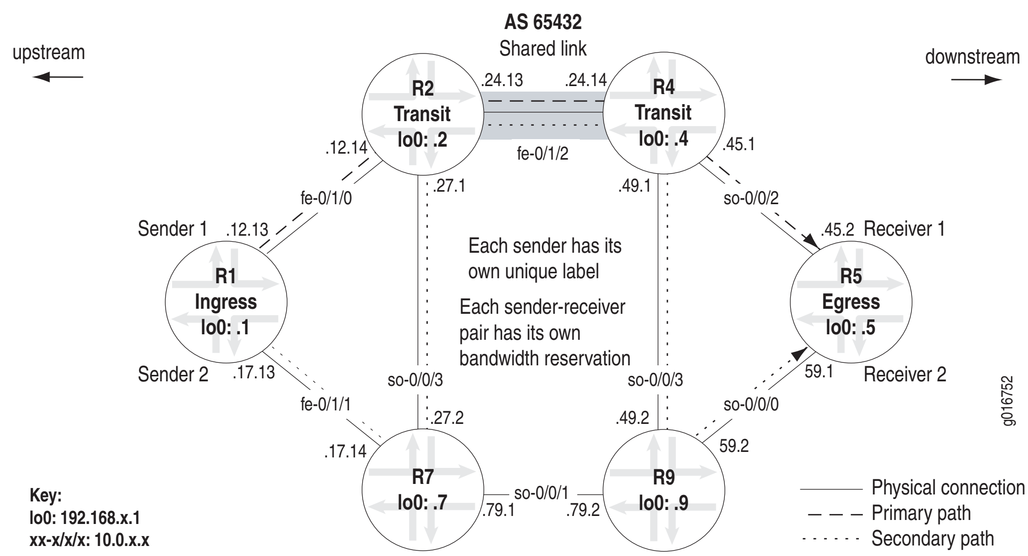


Figure 8 shows a primary and secondary path that share the Fast Ethernet link fe-0/1/2 between R2 and R4. Each path has a separate RSVP session in the FF reservation style. When sessions share a link, the total amount of reserved bandwidth on the shared link is the sum of the reservations for each individual session. If the sum of reservations is larger than the available bandwidth, the LSP cannot be established, as illustrated in the example network in Figure 8.

In the example network in Figure 8, R1 requests a 75-Mbps bandwidth reservation for all configured primary and secondary paths. Therefore, to establish a primary and standby secondary path, a 50-Mbps bandwidth reservation is required. Because the Fast Ethernet link has a total of 100 Mbps of bandwidth available, 75 Mbps of which is reserved for the primary path, leaving 25 Mbps for the standby secondary path, the standby secondary path cannot be established.

**Action** For an illustration of this situation, see the output for the following commands:

**show configuration protocols mpls** (See Sample Output on page 56)
**show rsvp session detail** (See Sample Output on page 56)
**show mpls lsp extensive** (See Sample Output on page 57)

**Sample Output**

```
user@R1> show configuration protocols mpls
bandwidth 75m;
label-switched-path lsp1 {
    to 192.168.5.1;
    primary via-r2;
    secondary via-r7 {
        standby;
    }
}
path via-r7 {
    10.0.17.14 strict;
    10.0.27.1 strict;
    10.0.24.14 strict;
    10.0.49.2 strict;
}
path via-r2 {
    10.0.12.14 strict;
    10.0.24.14 strict;
}
interface fe-0/1/0.0;
interface fe-0/1/1.0;
interface so-0/0/3.0;
```

**What It Means**

Sample output from R1 for the show configuration protocols mpls command shows the MPLS configuration that includes a bandwidth of 75 Mbps for all paths, LSP lsp1, a primary path, and a standby secondary path. Both named paths, path via-r7 and path via-r2, specify all transit routers up to the egress. The egress router is not specified. Both paths are strict, indicating that the route taken from one router to the next router is a direct path and cannot include any other routers. All specified addresses are interface addresses, ensuring that the incoming interface is the one specified and enforcing routing on a per-link basis.

From the network topology shown in Figure 8 on page 55, the link shared by both paths is from R2 to R4, fe-0/1/2, or address 10.0.24.14.

**Sample Output**

```
user@R1> show rsvp session detail
Ingress RSVP: 1 sessions

192.168.5.1
  From: 192.168.1.1, LSPstate: Up, ActiveRoute: 0
  LSPname: lsp1, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 102720
  Resv style: 1 FF, Label in: -, Label out: 102720
  Time left:    -, Since: Fri Jul 21 11:08:12 2006
  Tspec: rate 75Mbps size 75Mbps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 60165 protocol 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.12.14 (fe-0/1/0.0) 6 pkts
  RESV rcvfrom: 10.0.12.14 (fe-0/1/0.0) 6 pkts
  Explct route: 10.0.12.14 10.0.24.14 10.0.45.2
  Record route: <self> 10.0.12.14 10.0.24.14 10.0.45.2
Total 1 displayed, Up 1, Down 0
[...Output truncated...]
```

**What It Means** The sample output from R1 for the show rsvp session detail command shows that R1 has one ingress RSVP session established in the FF style and associated with the primary path, indicating that the standby secondary path is not established. If the secondary standby path was established, we would expect to see two ingress sessions, one for the primary path and another for the secondary standby path.

**Sample Output**

```
user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

192.168.5.1
  From: 192.168.1.1, State: Up, ActiveRoute: 0, LSPname: lsp1
  ActivePath: via-r2 (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   via-r2          State: Up
    Bandwidth: 75Mbps
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
 10.0.12.14 S 10.0.24.14 S 10.0.45.2 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node
10=SoftPreempt):
          10.0.12.14 10.0.24.14 10.0.45.2
    5 Jul 21 11:08:12 Selected as active path
    4 Jul 21 11:08:12 Record Route:  10.0.12.14 10.0.24.14 10.0.45.2
    3 Jul 21 11:08:12 Up
    2 Jul 21 11:08:12 Originate Call
    1 Jul 21 11:08:12 CSPF: computation result accepted
  Standby   via-r7          State: Dn
    Bandwidth: 75Mbps
    SmartOptimizeTimer: 180
        No computed ERO.
  Created: Fri Jul 21 11:08:11 2006
Total 1 displayed, Up 1, Down 0
[...Output truncated...]
```

**What It Means** Sample output from R1 for the show mpls lsp extensive command shows that 75 Mbps of bandwidth is allocated for each path. The secondary standby path is down (State: Dn) because there is not enough available bandwidth.

## Shared Explicit Style Overview

The SE RSVP reservation style creates shared reservations among explicit senders. For a single RSVP bandwidth reservation, the egress router (receiver) lists the senders sharing the reservation in a Resv message, resulting in the following:

- A multipoint-to-point LSP if the Path message does not contain an ERO or if the ERO is identical across senders. A common label is assigned.
- A separate LSP for each sender if the Path message contains a different ERO for each sender. A different label is assigned to different senders.
- Each LSP shares the bandwidth of the largest request across the shared link.

While any LSP can be established with an SE style reservation, the SE reservation is most useful during LSP reroute, for example, when a standby secondary path or link protection is configured. In general, the secondary LSP inherits the reservation style of the primary LSP, which is FF by default, or SE if link protection is used, unless the secondary LSP is configured with the adaptive statement at the secondary path level. See Figure 9.

**Figure 9: Shared Explicit Style**

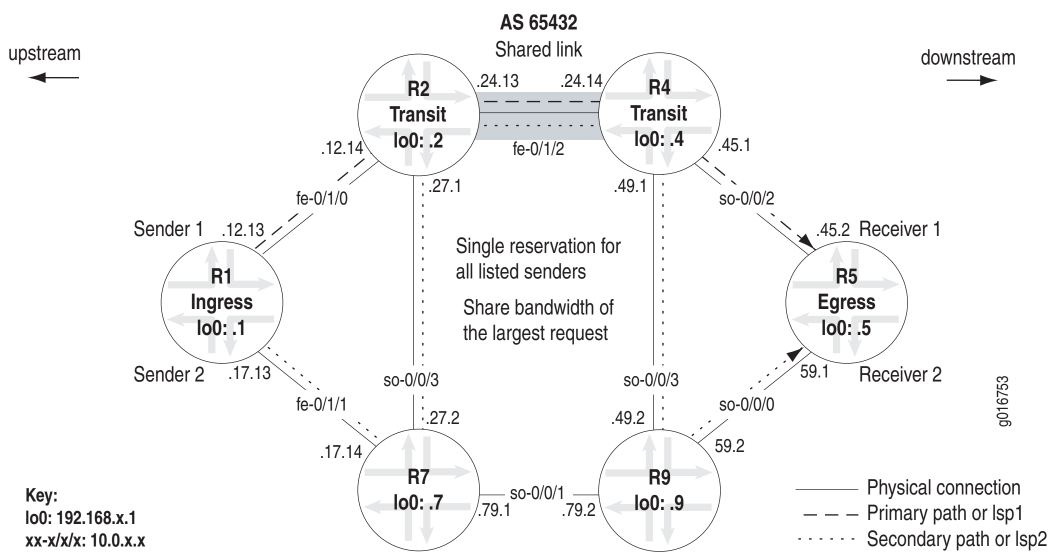


The network shown in Figure 9 shows two paths, R1-R2-R4-R5 and R1-R7-R9-R5. The paths are configured as either a strict primary or strict standby secondary path for an adaptive LSP, or as lsp1 and lsp2. Both configurations originate from R1 and share the link between R2 and R4. For more information about adaptive LSPs, see “Configuring and Verifying an Adaptive LSP” on page 59.

If a network problem results in an LSP reroute, the SE reservation style allows a smooth transition from either a primary path to a standby secondary path, or from on old LSP to a new LSP with the make-before-break operation. This style also permits the old and new LSPs to share a single reservation over links they have in common, preventing double counting of resources.

## Configuring and Verifying an Adaptive LSP

When you include the adaptive statement in the configuration, the LSP becomes adaptive and is established with the SE reservation style. The adaptive statement can be configured at two hierarchy levels:

- The [edit protocols mpls label-switched-path *lsp-path-name*] hierarchy level, which keeps the RSVP session information the same for all primary and secondary paths.
- The [edit protocols mpls label-switched-path *lsp-path-name* secondary *secondary-name*] hierarchy level, resulting in different Tunnel ID values for each path and causes the paths to be viewed as separate RSVP sessions, that may not share the same bandwidth reservation and possibly double-count resources.

Using an adaptive LSP at the [edit protocols mpls label-switched-path *lsp-path-name*] hierarchy level provides two advantages. The first advantage is the prevention of double-counting of bandwidth for links that share old and new paths. Double-counting occurs when an intermediate router does not recognize that the new and old paths belong to the same LSP and counts them as two separate LSPs, requiring separate bandwidth allocations. If some links are close to saturation, double-counting might cause the setup of the new path to fail. When the adaptive statement is included at the [edit protocols mpls label-switched-path *lsp-path-name*] hierarchy level, a standby secondary path is established, sharing physical links in common with the LSP's primary path.

The second advantage is the prevention of disruption to subscriber traffic by performing a make-before-break operation. When an established path attempts to reroute onto a new path, the ingress router maintains existing paths and allocated bandwidths, ensuring that the existing path is not prematurely torn down and allowing the current traffic to continue flowing while the new path is set up.

The following steps describe the process of configuring an adaptive LSP that keeps the RSVP session information the same for all primary and secondary paths. Before you can configure an adaptive LSP, you must have an LSP already configured with the primary and secondary paths you want to use, and any other options. For information on configuring a LSP with a primary path and secondary path, see "Path Protection in an MPLS Network" on page 9.

**Action** To configure an adaptive LSP, follow these steps:

1. In configuration mode, go to the following hierarchy level:

   [edit]
   user@R1# **edit protocol mpls**

2. Configure adaptive mode for the LSP:

   [edit protocols mpls]
   user@R1# **set label-switched-path** *lsp-path-name* **adaptive**

For example:

[edit protocols mpls]
user@R1# **set label-switched-path lsp1 adaptive**

3. Verify and commit the configuration:

[edit protocols mpls]
user@R1# **show**
user@R1# **commit**

**Sample Output**

```
[edit protocols mpls]
user@R1# show
bandwidth 75m;
label-switched-path lsp1 {
    to 192.168.5.1;
    adaptive;
    primary via-r2;
    secondary via-r7 {
        standby;
    }
}
path via-r7 {
    10.0.17.14 strict;
    10.0.27.1 strict;
    10.0.24.14 strict;
    10.0.49.2 strict;
}
path via-r2 {
    10.0.12.14 strict;
    10.0.24.14 strict;
}
interface fe-0/1/0.0;
interface fe-0/1/1.0;
interface so-0/0/3.0;

[edit protocols mpls]
user@R1# commit
commit complete
```

**What It Means**

Sample output from R1 for the show command shows bandwidth of 75 Mbps, the adaptive statement, and strict primary and secondary paths. 75 Mbps of bandwidth for each path is more combined bandwidth than the Fast Ethernet link fe-0/1/2 can accommodate. Because lsp1 is adaptive, both paths are up, indicating that the bandwidth is not double-counted, as shown in the following output for the show mpls lsp extensive command.

**Sample Output**

```
[edit protocols mpls]
user@R1# run show mpls lsp extensive
Ingress LSP: 1 sessions
192.168.5.1
  From: 192.168.1.1, State: Up, ActiveRoute: 0, LSPname: lsp1
  ActivePath: via-r2 (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   via-r2           State: Up
    Bandwidth: 75Mbps
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
 10.0.12.14 S 10.0.24.14 S 10.0.45.2 S
```

```
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 
10=SoftPreempt):
          10.0.12.14 10.0.24.14 10.0.45.2
    5 Jul 21 14:34:16 Selected as active path
    4 Jul 21 14:34:16 Record Route:  10.0.12.14 10.0.24.14 10.0.45.2
    3 Jul 21 14:34:16 Up
    2 Jul 21 14:34:16 Originate Call
    1 Jul 21 14:34:16 CSPF: computation result accepted
  Standby   via-r7           State: Up
    Bandwidth: 75Mbps
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 5)
 10.0.17.14 S 10.0.27.1 S 10.0.24.14 S 10.0.49.2 S 10.0.59.1 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node ...
          10.0.17.14 10.0.27.1 10.0.24.14 10.0.49.2 10.0.59.1
    4 Jul 21 14:34:45 Record Route:  10.0.17.14 10.0.27.1 10.0.24.14 10.0.49.2 
10.0.59.1
    3 Jul 21 14:34:45 Up
    2 Jul 21 14:34:45 Originate Call
    1 Jul 21 14:34:45 CSPF: computation result accepted
  Created: Fri Jul 21 14:34:15 2006
Total 1 displayed, Up 1, Down 0
```

**What It Means**

The sample output from R1 for the show mpls lsp extensive command shows that lsp1 is up with an active primary path that is up (*Primary via-r2 State: Up), and a standby secondary path that is also up (Standby via-r7 State: Up). Both paths have 75 Mbps of bandwidth, which is not double-counted because the adaptive statement ensures that new and old paths are recognized as belonging to the same LSP lsp1, as shown in the following sample output for the show rsvp session detail command. You can also use the show rsvp interface command to show the reserved and available bandwidth.

**Sample Output**

```
user@R1> show rsvp session detail
Ingress RSVP: 2 sessions

192.168.5.1
  From: 192.168.1.1, LSPstate: Up, ActiveRoute: 0
  LSPname: lsp1, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 102736
  Resv style: 1 SE, Label in: -, Label out: 102736
  Time left:    -, Since: Fri Jul 21 14:34:16 2006
  Tspec: rate 75Mbps size 75Mbps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 60167 protocol 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.12.14 (fe-0/1/0.0) 7 pkts
  RESV rcvfrom: 10.0.12.14 (fe-0/1/0.0) 7 pkts
  Explct route: 10.0.12.14 10.0.24.14 10.0.45.2
  Record route: <self> 10.0.12.14 10.0.24.14 10.0.45.2

192.168.5.1
  From: 192.168.1.1, LSPstate: Up, ActiveRoute: 0
  LSPname: lsp1, LSPpath: Secondary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 102608
  Resv style: 1 SE, Label in: -, Label out: 102608
  Time left:    -, Since: Fri Jul 21 14:34:45 2006
  Tspec: rate 75Mbps size 75Mbps peak Infbps m 20 M 1500
  Port number: sender 2 receiver 60167 protocol 0
```

```
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.17.14 (fe-0/1/1.0) 5 pkts
  RESV rcvfrom: 10.0.17.14 (fe-0/1/1.0) 5 pkts
  Explct route: 10.0.17.14 10.0.27.1 10.0.24.14 10.0.49.2 10.0.59.1
  Record route: <self> 10.0.17.14 10.0.27.1 10.0.24.14 10.0.49.2 10.0.59.1
Total 2 displayed, Up 2, Down 0
[...Output truncated...]
```

**What It Means** The sample output from R1 for the show rsvp session detail command shows two RSVP sessions for lsp1. Both sessions originate on R1 (192.168.1.1) and end in R5 (192.168.5.1). The first session is for the primary path and the second session is for the secondary path. Both paths are in the SE reservation style. The port number is the protocol ID and sender/receiver port used in this RSVP session. In the port number field, the primary session shows sender 1, while the secondary session shows sender 2, indicating that two senders are using the LSP tunnel.

## Rerouting the LSP Tunnel for the SE Reservation Style

An LSP tunnel may need to be rerouted due to conditions based on administrative policy, for example, when a more optimal route becomes available, when a resource fails along the LSP, or when a failed resource is reactivated. The SE reservation style allows a smooth transition from an old LSP to a new LSP with the make-before-break operation. This style also permits the old and new LSPs to share a single reservation over links they have in common, preventing double-counting of resources.

### *Step 1: Establish the Initial LSP Tunnel*

The ingress router uses the Path message to request that the egress router set up the initial LSP tunnel with the SE reservation style. When establishing the initial LSP tunnel, the ingress and egress routers perform the following actions:

1. The ingress router includes the following in the initial Path message:

   - A LSP Tunnel IPv4 Session object that contains the following:

     - IPv4 address of the egress node.
     - Tunnel ID that remains constant for the life of the LSP tunnel between the ingress and egress routers.
     - Extended Tunnel ID that identifies the ingress router's IPv4 address.

   - The SE reservation style in the Session Attribute object.
   - A Sender Template object that contains the following:

     - The IPv4 address of the sender (ingress) node.
     - A LSP ID that can change in the future, when the LSP needs to be rerouted, allowing the ingress router to appear as a different sender so it can share resources with itself (see the LSP ID field of the LSP Tunnel IPv4 C-type extension for the Sender Template and Filter Spec objects).

2. Upon receipt of the Path message, the egress router sends a Resv message with an SE reservation style toward the ingress router.

3. When the ingress router receives the Resv message, the initial LSP tunnel is established with an SE reservation style.

### *Step 2: Reroute an LSP Tunnel*

When the ingress router attempts to reroute an exiting LSP tunnel to increase the bandwidth or change the path, it transmits a new Path message. During the reroute operation, the ingress router must appear as two different senders to the RSVP session. This is achieved by including a new LSP ID in the Sender Template object and the Filter Spec object. The ingress and egress routers perform the following actions:

1. The ingress router includes the following in the new Path message:
   - An Explicit Route object (ERO) for the new LSP tunnel.
   - The existing LSP Tunnel IPv4 Session object to identify the LSP that will be rerouted.
   - A new LSP ID and a new Sender Template object, ensuring that the ingress router appears as a different sender to the RSVP session.
2. The ingress router transmits the new Path message toward the egress router, continuing to use the old LSP tunnel to forward traffic and continuing to refresh the original PATH message (make-before-break).
3. The egress router responds to the new Path message with a Resv message that contains a number of RSVP objects, including:
   - A Label object to support the upstream on-demand label distribution process
   - An SE reservation Style object

---

**NOTE:** On links that are not shared by the old and new LSP tunnels, the new Path/Resv message pair is treated as a new conventional LSP. However, on links that are traversed by both the old and new LSP tunnels, the LSP Tunnel IPv4 Session object and SE reservation style allow the new LSP tunnel to establish so that it shares resources with the old LSP tunnel, eliminating the double-counting problem on shared links.

---

4. The ingress router begins to use the new LSP tunnel after it receives the new Resv message.
5. The ingress router sends a Path Tear message to remove the old LSP tunnel from intermediate routers.

## Related Information

For additional information about MPLS fast reroute and MPLS protection methods, see the following:

- *JUNOS Feature Guide*
- *JUNOS MPLS Applications Configuration Guide*
- Semeria, Chuck. *RSVP Signaling Extensions for MPLS Traffic Engineering.* White paper. 2002
- Semeria, Chuck. *IP Dependability: Network Link and Node Protection*. White paper. 2002
- RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*