

Chapter 9

RSVP LSP Tunnels

A Resource Reservation Protocol (RSVP) label-switched path (LSP) tunnel allows you to send RSVP LSPs inside other RSVP LSPs. This allows a network administrator to provide traffic engineering from one end of the network to the other. A useful application for this feature is to connect customer edge (CE) routers with provider edge (PE) routers by using an RSVP LSP, and then tunnel this edge LSP inside a second RSVP LSP traveling across the network core.

This document assumes you have a general understanding of Multiprotocol Label Switching (MPLS) and label switching concepts. For more information about MPLS, see the *JUNOS MPLS Applications Configuration Guide*.

This feature guide covers these topics:

- Overview on page 354
- System Requirements on page 354
- Terms and Acronyms on page 355
- RSVP LSP Tunneling Operation on page 355
- Configuring an RSVP LSP Tunnel on page 356
 - Example: RSVP LSP Tunnel Configuration on page 360
 - Checking Your Work on page 373
- For More Information on page 380
- Revision History on page 380

Overview

An RSVP LSP tunnel adds the concept of a forwarding adjacency, similar to the one used for generalized Multiprotocol Label Switching (GMPLS). (For more information about GMPLS, see “GMPLS” on page 199.)

The forwarding adjacency creates a tunneled path for sending data between peer devices in an RSVP LSP network. Once a forwarding adjacency LSP (FA-LSP) has been established, other LSPs can be sent over the FA-LSP by using Constrained Shortest Path First (CSPF), Link Management Protocol (LMP), Open Shortest Path First (OSPF), and RSVP.

To enable an RSVP LSP tunnel, the JUNOS software uses the following mechanisms:

- LMP—Originally designed for GMPLS, LMP establishes forwarding adjacencies between RSVP LSP tunnel peers, and maintains and allocates resources for traffic engineering links (TE links).
- OSPF extensions—OSPF was designed to route packets to physical and logical interfaces related to a Physical Interface Card (PIC). This protocol has been extended to route packets to virtual peer interfaces defined in an LMP configuration.
- RSVP-TE extensions—RSVP-TE was designed to signal the setup of packet LSPs to physical interfaces. The protocol has been extended to request path setup for packet LSPs travelling to virtual peer interfaces defined in an LMP configuration.

The following limitations exist for LSP hierarchies:

- Circuit cross-connect (CCC)-based LSPs are not supported.
- Graceful restart is not supported.
- Link protection is not available for FA-LSPs or at the egress point of the forwarding adjacency.
- Point-to-multipoint LSPs are not supported across FA-LSPs.

System Requirements

To implement an RSVP LSP tunnel, your system must meet these minimum requirements:

- JUNOS Release 8.2 or later for support on MX-series routing platforms
- JUNOS Release 7.4 or later for RSVP LSP tunnel support on M-series and T-series routing platforms
- Two Juniper Networks M-series, MX-series, or T-series routing platforms for the ingress and egress points of the forwarding adjacency, and two Juniper Networks M-series, MX-series, or T-series routing platforms for the ingress and egress PE routers

Terms and Acronyms

- **forwarding adjacency**—A forwarding path for sending data between peer devices in an RSVP LSP tunnel network.
- **forwarding adjacency LSP (FA-LSP)**—A hierarchical RSVP LSP that provides a tunnel-like forwarding adjacency for other packet-based LSPs.
- **Link Management Protocol**—A protocol used to define forwarding adjacencies between RSVP LSP tunnel peers and to maintain and allocate resources on traffic engineering links (TE links).
- **traffic engineering link (TE link)**—A logical connection between devices used to support RSVP LSP tunnels. TE links can have addresses or IDs and are associated with certain resources or interfaces. Each TE link represents a forwarding adjacency between a pair of devices.

RSVP LSP Tunneling Operation

An RSVP LSP tunnel requires close interaction between LMP, RSVP, and OSPF. The following sequence of events describes how this works:

1. LMP notifies RSVP and OSPF of the control peer, the control adjacency, and resources for the TE link.
2. GMPLS extracts the LSP attributes from the configuration and requests RSVP to signal one or more specific paths, specified by the TE link addresses.
3. RSVP determines the local TE link, corresponding control adjacency and active control channel, and transmission parameters (such as IP destination). It requests that LMP allocate a resource from the TE link with the specified attributes. If LMP successfully finds a resource matching the attributes, label allocation succeeds. RSVP sends a **PathMsg** hop-by-hop until it reaches the target router.
4. The target router, on receiving the RSVP **PathMsg**, requests that LMP allocate a resource based on the signaled parameters. If label allocation succeeds, it sends back a **ResvMsg**.
5. If the signaling is successful, an RSVP LSP tunnel is provisioned.

Configuring an RSVP LSP Tunnel

You must complete the following tasks to implement an RSVP LSP tunnel:

- Configuring Link Management Protocol Traffic Engineering Links on page 356
- Configuring Link Management Protocol Peers on page 357
- Configuring Peer Interfaces in OSPF and RSVP on page 357
- Establishing FA-LSP Path Information on page 358
- Defining Label-Switched Paths for the FA-LSP on page 358
- Creating End-to-End LSPs to Traverse the FA-LSP on page 359

The following is an optional LSP task:

- Option: Tearing Down RSVP LSPs Gracefully on page 359

To view an RSVP LSP tunnel configuration example, see the following sections:

- Example: RSVP LSP Tunnel Configuration on page 360
- Checking Your Work on page 373

Configuring Link Management Protocol Traffic Engineering Links

To begin your RSVP LSP tunnel configuration, configure LMP TE links on both the ingress and egress routing platforms. Because TE links define a unidirectional connection between peer devices, you must configure TE links in both directions between peers to enable the bidirectional transport of packets.

To configure TE links in LMP, include the `te-link te-link-name` statement at the [edit protocols link-management] hierarchy level. Define the TE link options shown below, especially the label-switched path to be used as the FA-LSP to reach the peer. Optionally, you can specify the traffic engineering metric for the TE link. By default, the traffic engineering metric is derived from the CSPF computation.

```
[edit]
protocols {
  link-management {
    te-link te-link-name { # Name of the TE link.
      label-switched-path lsp-name; # LSP used for the forwarding adjacency.
      local-address ip-address; # Local IP address associated with the TE link.
      remote-address ip-address; # Remote IP address mapped to the TE link.
      te-metric value; # Traffic engineering metric used for the TE link.
    }
  }
}
```

Configuring Link Management Protocol Peers

After you set up TE links, configure LMP network peers with the `peer` statement at the `[edit protocols link-management]` hierarchy level. A peer is the network device with which your routing platform communicates and establishes an FA-LSP. Designate a peer name, configure the peer router ID as the address (often a loopback address), and apply the TE link to be associated with this peer. Remember to configure both sides of a peering relationship to enable bidirectional communication.

Unlike GMPLS, you must not configure a control channel for a peer. If you include a control channel, the configuration will fail to commit.

```
[edit]
protocols {
  link-management {
    peer peer-name { # Configure the name of your network peer.
      address ip-address; # Include the router ID of the peer.
      te-link te-link-name; # Assign a TE link to this peer.
    }
  }
}
```

Configuring Peer Interfaces in OSPF and RSVP

After you establish LMP peers, you must add peer interfaces to OSPF and RSVP. A peer interface is a virtual interface used to support the control adjacency between two peers.

The peer interface name must match the `peer-name` statement configured in LMP at the `[edit protocols link-management peer]` hierarchy level. Because actual protocol packets are sent and received by peer interfaces, the peer interfaces can be signaled and advertised to peers like any other physical interface configured for OSPF and RSVP. To configure OSPF routing for LMP peers, include the `peer-interface` statement at the `[edit protocols ospf area area-number]` hierarchy level. To configure RSVP signaling for LMP peers, include the `peer-interface` statement at the `[edit protocols rsvp]` hierarchy level.

```
[edit]
protocols {
  rsvp {
    peer-interface peer-name { # Configure the name of your LMP peer.
    }
  }
  ospf {
    area area-number {
      peer-interface peer-name { # Configure the name of your LMP peer.
      }
    }
  }
}
```

Establishing FA-LSP Path Information

When you configure explicit LSP paths for an FA-LSP, you must use the TE link remote address as your next-hop address. When CSPF is supported, you can use any path option you wish. However, when CSPF is disabled with the `no-cspf` statement at the `[edit protocols mpls label-switched-path lsp-name]` hierarchy level, you must use strict paths.

```
[edit]
protocols {
  mpls {
    path path-name {
      next-hop-address (strict | loose);
    }
  }
}
```



NOTE: If the end-to-end LSP originates on the same routing platform as the FA-LSP, you must disable CSPF and use strict paths.

Defining Label-Switched Paths for the FA-LSP

Next, define your FA-LSP by including the `label-switched-path` statement at the `[edit protocols mpls]` hierarchy level. Include the router ID of the peer in the `to` statement at the `[edit protocols mpls label-switched-path]` hierarchy level. Because packet LSPs are unidirectional, you must create one FA-LSP to reach the peer and a second FA-LSP to return from the peer.

```
[edit]
protocols {
  mpls {
    label-switched-path lsp-name {
      from ip-address;
      to ip-address;
      primary path-name;
      secondary path-name;
      no-cspf;           # This statement to disable CSPF is optional.
    }
  }
}
```

Creating End-to-End LSPs to Traverse the FA-LSP

After you create the FA-LSPs and TE links, you can configure end-to-end LSPs to travel across the FA-LSPs. Here are some general guidelines for this step:

- In general, it is best to start your end-to-end LSP at least one hop before the ingress of the FA-LSP and terminate your LSP at least one hop after the egress of the FA-LSP. If you do originate both the end-to-end LSP and the FA-LSP on the same routing platform, you must disable CSPF and configure strict paths for the LSPs.
- If you configure CSPF, OSPF, and RSVP to create the FA-LSP, the LSP should automatically use the FA-LSP because of the preferred lower metric generated by the reduced number of hops.
- If you disable CSPF on the FA-LSP with the `no-cspf` option at the `[edit protocols mpls label-switched-path]` hierarchy level, you must configure the peer interface as a strict next hop in the path for your regular LSP. To manually configure an LSP to travel over the FA-LSP, include the `peer-name` statement at the `[edit protocols mpls path path-name]` hierarchy level and include the `strict` option.

For more information about creating RSVP LSPs, see the *JUNOS MPLS Applications Configuration Guide*.

Option: Tearing Down RSVP LSPs Gracefully

You can tear down an RSVP LSP in a two-step process that gracefully withdraws the RSVP session used by the LSP. For all neighbors that support graceful teardown, a request for the teardown is sent by the routing platform to the destination endpoint for the LSP and all RSVP neighbors in the path. The request is included within the `ADMIN_STATUS` field of the RSVP packet. When neighbors receive the request, they prepare for the RSVP session to be withdrawn. A second message is sent by the routing platform to complete the teardown of the RSVP session. If a neighbor does not support graceful teardown, the request is handled as a standard session teardown rather than a graceful one.

To perform a graceful teardown of an RSVP session, issue the `clear rsvp session gracefully` command. Optionally, you can specify the source and destination address of the RSVP session, the LSP identifier of the RSVP sender, and the tunnel identifier of the RSVP session. To use these qualifiers, include the `connection-source`, `connection-destination`, `lsp-id`, and `tunnel-id` options when you issue the `clear rsvp session gracefully` command.

You can also configure the amount of time that the routing platform waits for neighbors to receive the graceful teardown request before initiating the actual teardown. To configure, include the `graceful-deletion-timeout` statement at the `[edit protocols rsvp]` hierarchy level. The default graceful deletion timeout value is 30 seconds, with a minimum value of 1 second and a maximum value of 300 seconds. To view the current value configured for graceful deletion timeout, issue the `show rsvp version operational mode` command.

Example: RSVP LSP Tunnel Configuration

Figure 30: RSVP LSP Tunnel Topology Diagram




Figure 30 shows an end-to-end RSVP LSP called `e2e_lsp_r0r5` that originates on Router 0 and terminates on Router 5. In transit, this LSP traverses the FA-LSP `fa_lsp_r1r4`. The return path is represented by the end-to-end RSVP LSP `e2e_lsp_r5r0` that travels over the FA-LSP `fa_lsp_r4r1`.

On Router 0, configure the end-to-end RSVP LSP that travels to Router 5. Use a strict path that traverses Router 1 and the LMP TE link traveling from Router 1 to Router 4.

```

Router 0 [edit]
interfaces {
  so-0/0/3 {
    unit 0 {
      family inet {
        address 10.1.2.1/30;
      }
      family mpls;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.255.41.222/32;
      }
      family mpls;
    }
  }
}
  
```

```

routing-options {
  forwarding-table {
    export pplb;
  }
}
protocols {
  rsvp {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  mpls {
    admin-groups {
      fa 1;
      backup 2;
      other 3;
    }
    label-switched-path e2e_lsp_r0r5 { # An end-to-end LSP traveling to Router 5.
      to 10.255.41.221;
      bandwidth 30k;
      primary path-fa; # Reference the requested path here.
    }
    path path-fa { # Configure the strict path here.
      10.1.2.2 strict;
      172.16.30.2 strict; # This traverses the TE link heading to Router 4.
    }
    interface all;
    interface fxp0.0 {
      disable;
    }
    interface so-3/2/1.0 {
      admin-group other;
    }
    interface so-0/0/3.0 {
      admin-group other;
    }
  }
  ospf {
    traffic-engineering;
    area 0.0.0.0 {
      interface fxp0.0 {
        disable;
      }
      interface all;
    }
  }
}
policy-options {
  policy-statement pplb {
    then {
      load-balance per-packet;
    }
  }
}

```

On Router 1, configure an FA-LSP to reach Router 4. Establish an LMP TE link and LMP peer relationship with Router 4. Reference the FA-LSP in the TE link and add the peer interface into both OSPF and RSVP.

When the return path end-to-end LSP arrives at Router 1, the routing platform performs a routing lookup and can forward traffic to Router 0. Make sure you configure OSPF correctly between Routers 0 and 1.

```

Router 1 [edit]
interfaces {
  so-0/0/1 {
    unit 0 {
      family inet {
        address 10.2.3.1/30;
      }
      family mpls;
    }
  }
  so-0/0/2 {
    unit 0 {
      family inet {
        address 10.2.4.1/30;
      }
      family mpls;
    }
  }
  so-0/0/3 {
    unit 0 {
      family inet {
        address 10.1.2.2/30;
      }
      family mpls;
    }
  }
  fe-0/1/2 {
    unit 0 {
      family inet {
        address 10.2.5.1/30;
      }
      family mpls;
    }
  }
  at-1/0/0 {
    atm-options {
      vpi 1;
    }
    unit 0 {
      vci 1.100;
      family inet {
        address 10.2.3.5/30;
      }
      family mpls;
    }
  }
}

```

```

routing-options {
  forwarding-table {
    export [ pplb choose_lsp ];
  }
}
protocols {
  rsvp {
    interface all;
    interface fxp0.0 {
      disable;
    }
    peer-interface r4; # Apply the LMP peer interface here.
  }
  mpls {
    admin-groups {
      fa 1;
      backup 2;
      other 3;
    }
    label-switched-path fa_lsp_r1r4 { # Configure your FA-LSP to Router 4 here.
      to 10.255.41.217;
      bandwidth 400k;
      primary path_r1r4; # Apply the FA-LSP path here.
    }
    path path_r1r4 { # Configure the FA-LSP path here.
      10.2.4.2;
      10.4.5.2;
      10.3.5.1;
    }
    interface so-0/0/3.0 {
      admin-group other;
    }
    interface so-0/0/1.0 {
      admin-group fa;
    }
    interface at-1/0/0.0 {
      admin-group backup;
    }
    interface fe-0/1/2.0 {
      admin-group backup;
    }
    interface so-0/0/2.0 {
      admin-group fa;
    }
  }
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface fxp0.0 {
      disable;
    }
    interface all;
    peer-interface r4; # Apply the LMP peer interface here.
  }
}

```

```

link-management { # Configure LMP statements here.
  te-link link_r1r4 { # Assign a name to the TE link here.
    local-address 172.16.30.1; # Configure a local address for the TE link.
    remote-address 172.16.30.2; # Configure a remote address for the TE link.
    te-metric 1; # Manually set a metric here if you are not relying on CSPF.
    label-switched-path fa_lsp_r1r4; # Reference the FA-LSP here.
  }
  peer r4 { # Configure LMP peers here.
    address 10.255.41.217; # Configure the loopback address of your peer here.
    te-link link_r1r4; # Apply the LMP TE link here.
  }
}
}
policy-options {
  policy-statement choose_lsp {
    term A {
      from community choose_e2e_lsp;
      then {
        install-nexthop strict lsp e2e_lsp_r1r4;
        accept;
      }
    }
    term B {
      from community choose_fa_lsp;
      then {
        install-nexthop strict lsp fa_lsp_r1r4;
        accept;
      }
    }
  }
  policy-statement pplb {
    then {
      load-balance per-packet;
    }
  }
  community choose_e2e_lsp members 1000:1000;
  community choose_fa_lsp members 2000:2000;
  community set_e2e_lsp members 1000:1000;
  community set_fa_lsp members 2000:2000;
}

```

On Router 2, configure OSPF, MPLS, and RSVP on all interfaces that transport the FA-LSPs across the core network.

```
Router 2 [edit]
interfaces {
  so-0/0/0 {
    unit 0 {
      family inet {
        address 10.4.5.1/30;
      }
      family mpls;
    }
  }
  so-0/0/1 {
    unit 0 {
      family inet {
        address 10.1.4.2/30;
      }
      family mpls;
    }
  }
  so-0/0/2 {
    unit 0 {
      family inet {
        address 10.2.4.2/30;
      }
      family mpls;
    }
  }
  fe-0/1/2 {
    unit 0 {
      family inet {
        address 10.3.4.2/30;
      }
      family mpls;
    }
  }
}
routing-options {
  forwarding-table {
    export pplb;
  }
}
```

protocols { # OSPF, MPLS, and RSVP form the core backbone for the FA-LSPs.

```

rsvp {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
mpls {
  admin-groups {
    fa 1;
    backup 2;
    other 3;
  }
  path path_r1 {
    10.2.4.1;
  }
  path path_r3r4 {
    10.4.5.2;
    10.3.5.1;
  }
  interface all;
  interface fxp0.0 {
    disable;
  }
  interface so-0/0/1.0 {
    admin-group other;
  }
  interface fe-0/1/2.0 {
    admin-group backup;
  }
  interface so-0/0/2.0 {
    admin-group fa;
  }
  interface so-0/0/0.0 {
    admin-group fa;
  }
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface fxp0.0 {
      disable;
    }
    interface all;
  }
}
}
  policy-options {
    policy-statement pplb {
      then {
        load-balance per-packet;
      }
    }
  }
}

```

On Router 3, configure OSPF, MPLS, and RSVP on all interfaces that transport the FA-LSPs across the core network.

```

Router 3 [edit]
interfaces {
  so-0/0/0 {
    unit 0 {
      family inet {
        address 10.4.5.2/30;
      }
      family mpls;
    }
  }
  so-0/0/1 {
    unit 0 {
      family inet {
        address 10.5.6.1/30;
      }
      family mpls;
    }
  }
  so-0/0/2 {
    unit 0 {
      family inet {
        address 10.3.5.2/30;
      }
      family mpls;
    }
  }
  fe-0/1/2 {
    unit 0 {
      family inet {
        address 10.2.5.2/30;
      }
      family mpls;
    }
  }
}
routing-options {
  forwarding-table {
    export pplb;
  }
}

```

protocols { # OSPF, MPLS, and RSVP form the core backbone for the FA-LSPs.

```

rsvp {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
mpls {
  admin-groups {
    fa 1;
    backup 2;
    other 3;
  }
  path path_r4 {
    10.3.5.1;
  }
  path path_r2r1 {
    10.4.5.1;
    10.2.4.1;
  }
  interface all;
  interface fxp0.0 {
    disable;
  }
  interface so-0/0/2.0 {
    admin-group fa;
  }
  interface fe-0/1/2.0 {
    admin-group backup;
  }
  interface so-0/0/1.0 {
    admin-group other;
  }
  interface so-0/0/0.0 {
    admin-group fa;
  }
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface fxp0.0 {
      disable;
    }
    interface all;
  }
}
}
policy-options {
  policy-statement pplb {
    then {
      load-balance per-packet;
    }
  }
}

```

On Router 4, configure a return path FA-LSP to reach Router 1. Establish an LMP TE link and LMP peer relationship with Router 1. Reference the FA-LSP in the TE link and add the peer interface into both OSPF and RSVP.

When the initial end-to-end LSP arrives at Router 4, the routing platform performs a routing lookup and can forward traffic to Router 5. Make sure you configure OSPF correctly between Routers 4 and 5.

```

Router 4 [edit]
interfaces {
  so-0/0/0 {
    unit 0 {
      family inet {
        address 10.3.6.1/30;
      }
      family mpls;
    }
  }
  so-0/0/1 {
    unit 0 {
      family inet {
        address 10.2.3.2/30;
      }
      family mpls;
    }
  }
  so-0/0/2 {
    unit 0 {
      family inet {
        address 10.3.5.1/30;
      }
      family mpls;
    }
  }
  fe-0/1/2 {
    unit 0 {
      family inet {
        address 10.3.4.1/30;
      }
      family mpls;
    }
  }
  at-1/0/0 {
    atm-options {
      vpi 1;
    }
    unit 0 {
      vci 1.100;
      family inet {
        address 10.2.3.6/30;
      }
      family mpls;
    }
  }
}

```

```

routing-options {
  forwarding-table {
    export [ pplb choose_lsp ];
  }
}
protocols {
  rsvp {
    interface all;
    interface fxp0.0 {
      disable;
    }
    peer-interface r1; # Apply the LMP peer interface here.
  }
  mpls {
    admin-groups {
      fa 1;
      backup 2;
      other 3;
    }
    label-switched-path fa_lsp_r4r1 { # Configure your FA-LSP here.
      to 10.255.41.216;
      bandwidth 400k;
      primary path_r4r1; # Apply the FA-LSP path here.
    }
    path path_r4r1 { # Configure the FA-LSP path here.
      10.3.5.2;
      10.4.5.1;
      10.2.4.1;
    }
    interface all;
    interface fxp0.0 {
      disable;
    }
    interface at-1/0/0.0 {
      admin-group backup;
    }
    interface so-0/0/2.0 {
      admin-group fa;
    }
    interface fe-0/1/2.0 {
      admin-group backup;
    }
    interface so-0/0/0.0 {
      admin-group other;
    }
    interface so-0/0/1.0 {
      admin-group fa;
    }
  }
}

```

```

ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface fxp0.0 {
      disable;
    }
    interface all;
    peer-interface r1; # Apply the LMP peer interface here.
  }
}
link-management { # Configure LMP statements here.
  te-link link_r4r1 { # Assign a name to the TE link here.
    local-address 172.16.30.2; # Configure a local address for the TE link.
    remote-address 172.16.30.1; # Configure a remote address for the TE link.
    te-metric 1; # Manually set a metric here if you are not relying on CSPF.
    label-switched-path fa_lsp_r4r1; # Reference the FA-LSP here.
  }
  peer r1 { # Configure LMP peers here.
    address 10.255.41.216; # Configure the loopback address of your peer here.
    te-link link_r4r1; # Apply the LMP TE link here.
  }
}
}
policy-options {
  policy-statement choose_lsp {
    term A {
      from community choose_e2e_lsp;
      then {
        install-nexthop strict lsp e2e_lsp_r4r1;
        accept;
      }
    }
    term B {
      from community choose_fa_lsp;
      then {
        install-nexthop strict lsp fa_lsp_r4r1;
        accept;
      }
    }
  }
  policy-statement pplb {
    then {
      load-balance per-packet;
    }
  }
  community choose_e2e_lsp members 1000:1000;
  community choose_fa_lsp members 2000:2000;
  community set_e2e_lsp members 1000:1000;
  community set_fa_lsp members 2000:2000;
}

```

On Router 5, configure the return path end-to-end RSVP LSP that travels to Router 0. Use a strict path that traverses Router 4 and the LMP TE link traveling from Router 4 to Router 1.

```
Router 5 [edit]
interfaces {
  so-0/0/2 {
    unit 0 {
      family inet {
        address 10.3.6.2/30;
      }
      family mpls;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.255.41.221/32;
      }
    }
  }
}
routing-options {
  forwarding-table {
    export pplb;
  }
}
protocols {
  rsvp {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
}
```

```

mpls {
  admin-groups {
    fa 1;
    backup 2;
    other 3;
  }
  label-switched-path e2e_lsp_r5r0 { # An end-to-end LSP returning to Router 0.
    to 10.255.41.222;
    bandwidth 30k;
    primary path-fa; # Reference the requested path here.
  }
  path path-fa { # Configure the strict path here.
    10.3.6.1 strict;
    172.16.30.1 strict; # This traverses the TE link heading to Router 1.
  }
  interface all;
  interface fxp0.0 {
    disable;
  }
  interface so-0/0/2.0 {
    admin-group other;
  }
  interface so-0/0/1.0 {
    admin-group other;
  }
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface fxp0.0 {
      disable;
    }
    interface all;
  }
}
}
policy-options {
  policy-statement pplb {
    then {
      load-balance per-packet;
    }
  }
}
}

```

Checking Your Work

To verify that your RSVP LSP tunnel is working correctly, issue the following commands:

- show ted database (extensive)
- show rsvp session name (extensive)
- show link-management
- show link-management te-link name (detail)

To see these commands used with the configuration example, see the following sections:

- Router 0 on page 374
- Router 1 on page 378

Router 0

On Router 0, you can verify that the FA-LSPs appear as valid paths in the traffic engineering database. In this case, look for the paths from Router 1 (10.255.41.216) and Router 4 (10.255.41.217) that reference the LMP TE link addresses of 172.16.30.1 and 172.16.30.2. You can also issue the `show rsvp session extensive` command to look for the path of the end-to-end LSP as it travels to Router 5 over the FA-LSP.

```
user@router0> show ted database
TED database: 0 ISIS nodes 8 INET nodes
ID                               Type Age(s) LnkIn LnkOut Protocol
10.255.41.214                     Rtr   486     4     4 OSPF(0.0.0.0)
  To: 10.255.41.222, Local: 10.1.4.2, Remote: 10.1.4.1
  To: 10.255.41.216, Local: 10.2.4.2, Remote: 10.2.4.1
  To: 10.255.41.215, Local: 10.4.5.1, Remote: 10.4.5.2
  To: 10.3.4.1-1, Local: 10.3.4.2, Remote: 0.0.0.0
ID                               Type Age(s) LnkIn LnkOut Protocol
10.255.41.215                     Rtr   187     4     4 OSPF(0.0.0.0)
  To: 10.255.41.214, Local: 10.4.5.2, Remote: 10.4.5.1
  To: 10.255.41.217, Local: 10.3.5.2, Remote: 10.3.5.1
  To: 10.255.41.221, Local: 10.5.6.1, Remote: 10.5.6.2
  To: 10.2.5.1-1, Local: 10.2.5.2, Remote: 0.0.0.0
ID                               Type Age(s) LnkIn LnkOut Protocol
10.255.41.216                     Rtr   396     6     6 OSPF(0.0.0.0)
  To: 10.255.41.222, Local: 10.1.2.2, Remote: 10.1.2.1
  To: 10.255.41.214, Local: 10.2.4.1, Remote: 10.2.4.2
  To: 10.255.41.217, Local: 10.2.3.1, Remote: 10.2.3.2
  To: 10.255.41.217, Local: 172.16.30.1, Remote: 172.16.30.2
  To: 10.255.41.217, Local: 10.2.3.5, Remote: 10.2.3.6
  To: 10.2.5.1-1, Local: 10.2.5.1, Remote: 0.0.0.0
ID                               Type Age(s) LnkIn LnkOut Protocol
10.255.41.217                     Rtr   404     6     6 OSPF(0.0.0.0)
  To: 10.255.41.216, Local: 10.2.3.2, Remote: 10.2.3.1
  To: 10.255.41.216, Local: 172.16.30.2, Remote: 172.16.30.1
  To: 10.255.41.216, Local: 10.2.3.6, Remote: 10.2.3.5
  To: 10.255.41.215, Local: 10.3.5.1, Remote: 10.3.5.2
  To: 10.255.41.221, Local: 10.3.6.1, Remote: 10.3.6.2
  To: 10.3.4.1-1, Local: 10.3.4.1, Remote: 0.0.0.0
ID                               Type Age(s) LnkIn LnkOut Protocol
10.255.41.221                     Rtr   481     2     2 OSPF(0.0.0.0)
  To: 10.255.41.215, Local: 10.5.6.2, Remote: 10.5.6.1
  To: 10.255.41.217, Local: 10.3.6.2, Remote: 10.3.6.1
ID                               Type Age(s) LnkIn LnkOut Protocol
10.255.41.222                     Rtr  2883     2     2 OSPF(0.0.0.0)
  To: 10.255.41.216, Local: 10.1.2.1, Remote: 10.1.2.2
  To: 10.255.41.214, Local: 10.1.4.1, Remote: 10.1.4.2
```

```

user@router0> show ted database 10.255.41.216 extensive
TED database: 0 ISIS nodes 8 INET nodes
NodeID: 10.255.41.216
  Type: Rtr, Age: 421 secs, LinkIn: 6, LinkOut: 6
  Protocol: OSPF(0.0.0.0)
    To: 10.255.41.222, Local: 10.1.2.2, Remote: 10.1.2.1
      Color: 0x8 other
      Metric: 1
      Static BW: 155.52Mbps
      Reservable BW: 155.52Mbps
      Available BW [priority] bps:
        [0] 155.4Mbps [1] 155.4Mbps [2] 155.4Mbps [3] 155.4Mbps
        [4] 155.4Mbps [5] 155.4Mbps [6] 155.4Mbps [7] 155.4Mbps
      Interface Switching Capability Descriptor(1):
        Switching type: Packet
        Encoding type: Packet
        Maximum LSP BW [priority] bps:
          [0] 155.4Mbps [1] 155.4Mbps [2] 155.4Mbps [3] 155.4Mbps
          [4] 155.4Mbps [5] 155.4Mbps [6] 155.4Mbps [7] 155.4Mbps
    To: 10.255.41.214, Local: 10.2.4.1, Remote: 10.2.4.2
      Color: 0x2 fa
      Metric: 1
      Static BW: 155.52Mbps
      Reservable BW: 155.52Mbps
      Available BW [priority] bps:
        [0] 155.12Mbps [1] 155.12Mbps [2] 155.12Mbps [3] 155.12Mbps
        [4] 155.12Mbps [5] 155.12Mbps [6] 155.12Mbps [7] 155.12Mbps
      Interface Switching Capability Descriptor(1):
        Switching type: Packet
        Encoding type: Packet
        Maximum LSP BW [priority] bps:
          [0] 155.12Mbps [1] 155.12Mbps [2] 155.12Mbps [3] 155.12Mbps
          [4] 155.12Mbps [5] 155.12Mbps [6] 155.12Mbps [7] 155.12Mbps
    To: 10.255.41.217, Local: 10.2.3.1, Remote: 10.2.3.2
      Color: 0x2 fa
      Metric: 1
      Static BW: 155.52Mbps
      Reservable BW: 155.52Mbps
      Available BW [priority] bps:
        [0] 155.52Mbps [1] 155.52Mbps [2] 155.52Mbps [3] 155.52Mbps
        [4] 155.52Mbps [5] 155.52Mbps [6] 155.52Mbps [7] 155.52Mbps
      Interface Switching Capability Descriptor(1):
        Switching type: Packet
        Encoding type: Packet
        Maximum LSP BW [priority] bps:
          [0] 155.52Mbps [1] 155.52Mbps [2] 155.52Mbps [3] 155.52Mbps
          [4] 155.52Mbps [5] 155.52Mbps [6] 155.52Mbps [7] 155.52Mbps
    To: 10.255.41.217, Local: 172.16.30.1, Remote: 172.16.30.2
      Metric: 1
      Static BW: 400kbps
      Reservable BW: 400kbps
      Available BW [priority] bps:
        [0] 370kbps [1] 370kbps [2] 370kbps [3] 370kbps
        [4] 370kbps [5] 370kbps [6] 370kbps [7] 370kbps
      Interface Switching Capability Descriptor(1):
        Switching type: Packet
        Encoding type: Packet
        Maximum LSP BW [priority] bps:
          [0] 370kbps [1] 370kbps [2] 370kbps [3] 370kbps
          [4] 370kbps [5] 370kbps [6] 370kbps [7] 370kbps

```

```

To: 10.255.41.217, Local: 10.2.3.5, Remote: 10.2.3.6
Color: 0x4 backup
Metric: 1
Static BW: 155.52Mbps
Reservable BW: 155.52Mbps
Available BW [priority] bps:
  [0] 155.52Mbps  [1] 155.52Mbps  [2] 155.52Mbps  [3] 155.52Mbps
  [4] 155.52Mbps  [5] 155.52Mbps  [6] 155.52Mbps  [7] 155.52Mbps
Interface Switching Capability Descriptor(1):
Switching type: Packet
Encoding type: Packet
Maximum LSP BW [priority] bps:
  [0] 155.52Mbps  [1] 155.52Mbps  [2] 155.52Mbps  [3] 155.52Mbps
  [4] 155.52Mbps  [5] 155.52Mbps  [6] 155.52Mbps  [7] 155.52Mbps
To: 10.2.5.1-1, Local: 10.2.5.1, Remote: 0.0.0.0
Color: 0x4 backup
Metric: 1
Static BW: 100Mbps
Reservable BW: 100Mbps
Available BW [priority] bps:
  [0] 100Mbps    [1] 100Mbps    [2] 100Mbps    [3] 100Mbps
  [4] 100Mbps    [5] 100Mbps    [6] 100Mbps    [7] 100Mbps
Interface Switching Capability Descriptor(1):
Switching type: Packet
Encoding type: Packet
Maximum LSP BW [priority] bps:
  [0] 100Mbps    [1] 100Mbps    [2] 100Mbps    [3] 100Mbps
  [4] 100Mbps    [5] 100Mbps    [6] 100Mbps    [7] 100Mbps

```

```

user@router0> show ted database 10.255.41.217 extensive
TED database: 0 ISIS nodes 8 INET nodes
NodeID: 10.255.41.217
Type: Rtr, Age: 473 secs, LinkIn: 6, LinkOut: 6
Protocol: OSPF(0.0.0.0)
To: 10.255.41.216, Local: 10.2.3.2, Remote: 10.2.3.1
Color: 0x2 fa
Metric: 1
Static BW: 155.52Mbps
Reservable BW: 155.52Mbps
Available BW [priority] bps:
  [0] 155.52Mbps  [1] 155.52Mbps  [2] 155.52Mbps  [3] 155.52Mbps
  [4] 155.52Mbps  [5] 155.52Mbps  [6] 155.52Mbps  [7] 155.52Mbps
Interface Switching Capability Descriptor(1):
Switching type: Packet
Encoding type: Packet
Maximum LSP BW [priority] bps:
  [0] 155.52Mbps  [1] 155.52Mbps  [2] 155.52Mbps  [3] 155.52Mbps
  [4] 155.52Mbps  [5] 155.52Mbps  [6] 155.52Mbps  [7] 155.52Mbps
To: 10.255.41.216, Local: 172.16.30.2, Remote: 172.16.30.1
Metric: 1
Static BW: 400kbps
Reservable BW: 400kbps
Available BW [priority] bps:
  [0] 370kbps    [1] 370kbps    [2] 370kbps    [3] 370kbps
  [4] 370kbps    [5] 370kbps    [6] 370kbps    [7] 370kbps
Interface Switching Capability Descriptor(1):
Switching type: Packet
Encoding type: Packet
Maximum LSP BW [priority] bps:
  [0] 370kbps    [1] 370kbps    [2] 370kbps    [3] 370kbps
  [4] 370kbps    [5] 370kbps    [6] 370kbps    [7] 370kbps

```

```

To: 10.255.41.216, Local: 10.2.3.6, Remote: 10.2.3.5
Color: 0x4 backup
Metric: 1
Static BW: 155.52Mbps
Reservable BW: 155.52Mbps
Available BW [priority] bps:
  [0] 155.52Mbps  [1] 155.52Mbps  [2] 155.52Mbps  [3] 155.52Mbps
  [4] 155.52Mbps  [5] 155.52Mbps  [6] 155.52Mbps  [7] 155.52Mbps
Interface Switching Capability Descriptor(1):
Switching type: Packet
Encoding type: Packet
Maximum LSP BW [priority] bps:
  [0] 155.52Mbps  [1] 155.52Mbps  [2] 155.52Mbps  [3] 155.52Mbps
  [4] 155.52Mbps  [5] 155.52Mbps  [6] 155.52Mbps  [7] 155.52Mbps
To: 10.255.41.215, Local: 10.3.5.1, Remote: 10.3.5.2
Color: 0x2 fa
Metric: 1
Static BW: 155.52Mbps
Reservable BW: 155.52Mbps
Available BW [priority] bps:
  [0] 155.12Mbps  [1] 155.12Mbps  [2] 155.12Mbps  [3] 155.12Mbps
  [4] 155.12Mbps  [5] 155.12Mbps  [6] 155.12Mbps  [7] 155.12Mbps
Interface Switching Capability Descriptor(1):
Switching type: Packet
Encoding type: Packet
Maximum LSP BW [priority] bps:
  [0] 155.12Mbps  [1] 155.12Mbps  [2] 155.12Mbps  [3] 155.12Mbps
  [4] 155.12Mbps  [5] 155.12Mbps  [6] 155.12Mbps  [7] 155.12Mbps
To: 10.255.41.221, Local: 10.3.6.1, Remote: 10.3.6.2
Color: 0x8 other
Metric: 1
Static BW: 155.52Mbps
Reservable BW: 155.52Mbps
Available BW [priority] bps:
  [0] 155.52Mbps  [1] 155.52Mbps  [2] 155.52Mbps  [3] 155.52Mbps
  [4] 155.52Mbps  [5] 155.52Mbps  [6] 155.52Mbps  [7] 155.52Mbps
Interface Switching Capability Descriptor(1):
Switching type: Packet
Encoding type: Packet
Maximum LSP BW [priority] bps:
  [0] 155.52Mbps  [1] 155.52Mbps  [2] 155.52Mbps  [3] 155.52Mbps
  [4] 155.52Mbps  [5] 155.52Mbps  [6] 155.52Mbps  [7] 155.52Mbps
To: 10.3.4.1-1, Local: 10.3.4.1, Remote: 0.0.0.0
Color: 0x4 backup
Metric: 1
Static BW: 100Mbps
Reservable BW: 100Mbps
Available BW [priority] bps:
  [0] 100Mbps    [1] 100Mbps    [2] 100Mbps    [3] 100Mbps
  [4] 100Mbps    [5] 100Mbps    [6] 100Mbps    [7] 100Mbps
Interface Switching Capability Descriptor(1):
Switching type: Packet
Encoding type: Packet
Maximum LSP BW [priority] bps:
  [0] 100Mbps    [1] 100Mbps    [2] 100Mbps    [3] 100Mbps
  [4] 100Mbps    [5] 100Mbps    [6] 100Mbps    [7] 100Mbps

```

```

user@router0> show rsvp session name e2e_lsp_r0r5 extensive
Ingress RSVP: 1 sessions

10.255.41.221
  From: 10.255.41.222, LSPstate: Up, ActiveRoute: 2
  LSPname: e2e_lsp_r0r5, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 101584
  Resv style: 1 FF, Label in: -, Label out: 101584
  Time left: -, Since: Wed Sep 7 19:02:56 2005
  Tspec: rate 30kbps size 30kbps peak Infbps m 20 M 1500
  Port number: sender 2 receiver 29458 protocol 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.1.2.2 (so-0/0/3.0) 15 pkts
  RESV rcvfrom: 10.1.2.2 (so-0/0/3.0) 16 pkts
  Explct route: 10.1.2.2 172.16.30.2 10.3.6.2
  Record route: <self> 10.1.2.2 172.16.30.2 10.3.6.2
Total 1 displayed, Up 1, Down 0

Egress RSVP: 1 sessions
Total 0 displayed, Up 0, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Router 1

On Router 1, verify that your LMP TE link configuration is working and that the end-to-end LSP is traversing the TE link by issuing the `show link-management set` of commands. You can also issue the `show rsvp session extensive` command to confirm that the FA-LSP is operational.

```

user@router1> show link-management
Peer name: r4, System identifier: 10758
State: Up, Control address: 10.255.41.217
  TE links:
    link_r1r4

  TE link name: link_r1r4, State: Up
    Local identifier: 16299, Remote identifier: 0, Local address: 172.16.30.1,
    Remote address: 172.16.30.2,
    Encoding: Packet, Switching: Packet, Minimum bandwidth: 0bps, Maximum
    bandwidth: 400kbps,
    Total bandwidth: 400kbps, Available bandwidth: 370kbps
    Name      State Local ID Remote ID   Bandwidth Used LSP-name
    fa_lsp_r1r4 Up      22642      0      400kbps Yes  e2e_lsp_r0r5

```

```

user@router1> show link-management te-link name link_r1r4 detail
  TE link name: link_r1r4, State: Up
    Local identifier: 16299, Remote identifier: 0, Local address: 172.16.30.1,
  Remote address: 172.16.30.2,
    Encoding: Packet, Switching: Packet, Minimum bandwidth: 0bps, Maximum
  bandwidth: 400kbps,
    Total bandwidth: 400kbps, Available bandwidth: 370kbps
    Resource: fa_lsp_r1r4, Type: LSP, System identifier: 2147483683, State: Up,
  Local identifier: 22642,
    Remote identifier: 0
    Total bandwidth: 400kbps, Unallocated bandwidth: 370kbps
    Traffic parameters: Encoding: Packet, Switching: Packet, Granularity:
  Unknown
    Number of allocations: 1, In use: Yes
      LSP name: e2e_lsp_r0r5, Allocated bandwidth: 30kbps

user@router1> show rsvp session name fa_lsp_r1r4 extensive
Ingress RSVP: 1 sessions

10.255.41.217
  From: 10.255.41.216, LSPstate: Up, ActiveRoute: 0
  LSPname: fa_lsp_r1r4, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 100816
  Resv style: 1 FF, Label in: -, Label out: 100816
  Time left: -, Since: Wed Sep 7 19:02:33 2005
  Tspec: rate 400kbps size 400kbps peak Infbps m 20 M 1500
  Port number: sender 2 receiver 5933 protocol 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.2.4.2 (so-0/0/2.0) 28 pkts
  RESV rcvfrom: 10.2.4.2 (so-0/0/2.0) 26 pkts
  Explct route: 10.2.4.2 10.4.5.2 10.3.5.1
  Record route: <self> 10.2.4.2 10.4.5.2 10.3.5.1
  Total 1 displayed, Up 1, Down 0

Egress RSVP: 1 sessions
Total 0 displayed, Up 0, Down 0

Transit RSVP: 2 sessions
Total 0 displayed, Up 0, Down 0

```

For More Information

For additional information about implementing RSVP LSP tunnels, see the following:

- *JUNOS MPLS Applications Configuration Guide*
- RFC 2205, *Resource ReSerVation Protocol (RSVP)*
- Internet draft draft-ietf-mpls-bundle-04.txt, *Link Bundling in MPLS Traffic Engineering* (expires January 2003)
- Internet draft draft-ietf-ccamp-imp-10.txt, *Link Management Protocol (LMP)* (expires April 2004)
- Internet draft draft-ietf-mpls-lsp-hierarchy-08.txt, *LSP Hierarchy with Generalized MPLS TE* (expires March 2003)

Revision History

12 January 2007—Added support for MX960 Ethernet Services Routers. 8.2R1
Fawn Damitio.

15 September 2006—8.1R1 Release. Richard Hendricks.

29 June 2006—8.0R1 Release. Richard Hendricks.

27 March 2006—7.6R1 Release. Richard Hendricks.

9 January 2006—7.5R1 Release. Richard Hendricks.

14 September 2005—Initial document written, 7.4R1 Release. Richard Hendricks.