Chapter 2

Path Protection in an MPLS Network

The JUNOS software implementation of Mutiprotocol Label Switching (MPLS) provides several complementary mechanisms for protecting against Resource Reservation Protocol (RSVP)-signaled LSP failures, including path protection (primary and secondary paths), and local protection (the fast reroute statement, link protection, and node-link protection). This chapter describes path protection supported by the JUNOS software. (See Table 5.)

Table 5: Checklist for Path Protection

Path Protection Tasks	Command or Action
Path Protection Overview on page 11	
Configuring and Verifying a Primary Path on page 12	
1. Configure a Primary Path on page 13	[edit] edit protocols mpls [edit protocols mpls] set path *path-name address* < strict | loose > set label-switched-path *lsp-path-name* to *destination* [edit protocols mpls label-switched-path *lsp-path-name*] set primary *primary-name* set primary *primary-name* bandwidth *bandwidth* set primary *primary-name* priority *reservation-priority setup-priority* show commit
2. Verify That the Primary Path Is Operational on page 15	show mpls lsp extensive ingress show rsvp interface
Configuring and Verifying a Secondary Path on page 17	
1. Configure a Standby Secondary Path on page 18	[edit] edit protocols mpls [edit protocols mpls] set path *path-name destination* loose set label-switched-path *lsp-path-name* secondary *secondary-name* standby show commit

Path Protection Tasks	Command or Action
2. Verify That the Secondary Path Is Established on page 19	Deactivate a link or node critical to the primary path. show mpls lsp extensive
Ensuring That Secondary Paths Establish When Resources Are Diminished on page 21	Configure different bandwidth values for the primary and secondary paths. For example: [edit protocols mpls] edit label-switched-path *lsp-path-name* set primary *primary-name* bandwidth *bandwidth* show commit In this example, no bandwidth is configured for the secondary path.
Preventing Use of a Path That Previously Failed on page 22	Configure only multiple secondary paths.

Path Protection Overview

The main advantages of path protection are control over where the traffic goes after a failure and minimum packet loss when combined with fast reroute (one-to-one backup or link protection). Path protection is the configuration, within a label-switched path (LSP), of two types of paths: a primary path, used in normal operations, and a secondary path used when the primary fails, as shown in Figure 1.

In Figure 1, an MPLS network consisting of eight routers has a primary path between R1 and R5 which is protected by the secondary path between R1 and R5. When a failure is detected, such as an interface down event, an Resource Reservation Protocol (RSVP) error message is sent to the ingress router which switches traffic to the secondary path, maintaining traffic flow.

Figure 1: Path Protection

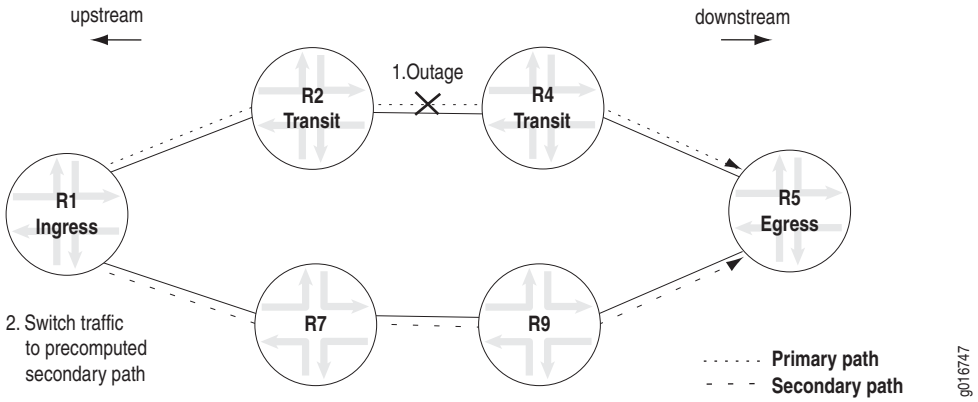

If the secondary path is pre-signaled or on standby, recovery time from a failure is faster than if the secondary path is not pre-signaled. When the secondary path is not pre-signaled a call-setup delay occurs during which the new physical path for the LSP is established, extending the recovery time. If the failure in the primary path is corrected, and after a few minutes of hold time, the ingress router switches traffic back from the secondary path to the primary path.

Because path protection is provided by the ingress router for the entire path, there can be some disadvantages, for example, double-booking of resources and unnecessary protection of links. By protecting a single resource at a time, local protection can remedy these disadvantages.

Configuring and Verifying a Primary Path

Primary paths are optional and when configured, limit the RSVP calculation of the complete path to the routers specified in the primary Explicit Route Object (ERO) list, which determines the physical path for the LSP. When primary paths are not configured, the ingress router determines the path to the egress router. Only one primary path is permitted per LSP, as shown in Figure 2.

Figure 2: Primary Path

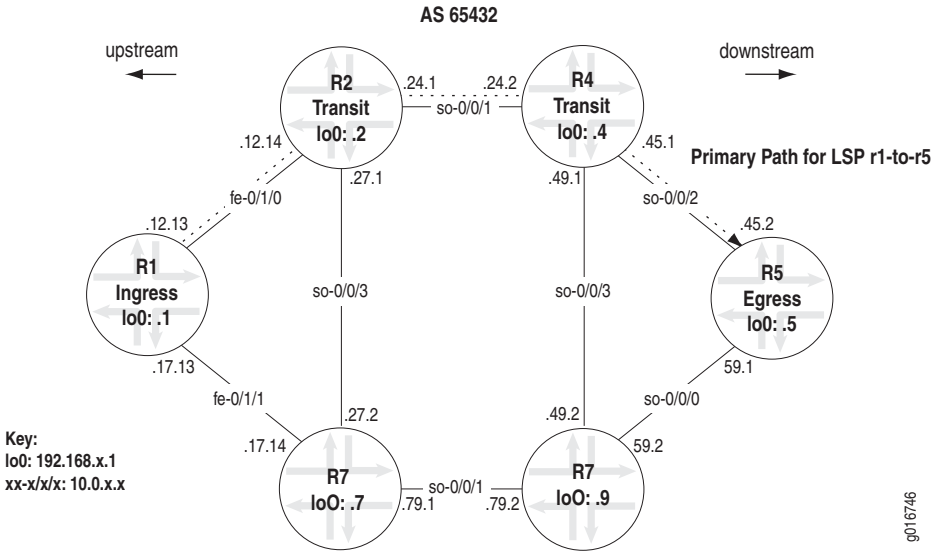

Within the configuration of the primary physical path, you can specify strict or loose ERO values and parameters that affect only the primary physical path, such as bandwidth or priority. The ERO list for the primary path includes an address for each transit router. Specifying the ingress and/or egress routers is optional. For each router address, you can specify the type, which can be one of the following:

- Strict—The route taken from the previous router to this router is a direct path and cannot include any other routers. This is the default. If the address is an interface address, this router also ensures that the incoming interface is the one specified. Specifying the incoming interface is important when there are parallel links between the previous router and this router, and because it ensures that routing can be enforced on a per-link basis.

  For strict addresses, you must ensure that the router immediately preceding the router you are configuring has a direct connection to that router. The address can be a loopback interface address, in which case the incoming interface is not checked.

- Loose—The route taken from the previous router to this router need not be a direct path, can include other routers, and can be received on any interface. The address can be any interface address or the address of the loopback interface.

If you are listing more than one address, specify the addresses in order, starting with the ingress router (optional) or the first transit router, and continuing sequentially along the path up to the egress router (optional) or the router immediately before the egress router. You need to specify only one address per router hop. If you specify more than one address for the same router, only the first address is used; the additional addresses are ignored and truncated.

When configuring a primary path, you can specify the bandwidth and priority values associated with that primary path.

The bandwidth value is included in the sender's Tspec field in RSVP path setup messages. You specify the bandwidth value in bits per second, with a higher value implying a greater user traffic volume. The default bandwidth is 0 bits per second. A nonzero bandwidth requires transit routers to reserve capacity along the outbound links for the path. The RSVP reservation scheme is used to reserve this capacity. Any failure in bandwidth reservation (such as failures at RSVP policy control or admission control) might cause the LSP setup to fail.

The priority value is composed of two distinct values: a setup and a hold priority. The setup priority value is used to determine if there is enough bandwidth available at that priority level to establish the primary path. The priority level is between 0 (best) and 7 (worst).The hold priority value is used by an established primary path to retain its bandwidth reservations in the network. If insufficient link bandwidth is available during session establishment, the setup priority is compared to the hold priorities of other established sessions to determine whether some of them should be preempted to accommodate the new session. Sessions with worse hold priorities are preempted.

Steps To Take To configure and verify a primary path, follow these steps:

1. Configure a Primary Path on page 13
2. Verify That the Primary Path Is Operational on page 15

Step 1: Configure a Primary Path

Action To configure a primary path with an ERO list, bandwidth, and priority, follow these steps:

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

   ```
   [edit]
   user@host# edit protocols mpls
   ```

2. Configure the primary ERO list:

   ```
   [edit protocols mpls]
   user@host# set path path-name address strict
   ```

   For example:

   ```
   [edit protocols mpls]
   user@R1# set path via-r2 10.0.12.14 strict
   user@R1# set path via-r2 10.0.24.2 strict
   ```

3. Configure the LSP:

   [edit protocols mpls]
   user@host# **set label-switched-path** *lsp-path-name* **to** *destination*;

   For example:

   [edit protocols mpls]
   user@R1# **set label-switched-path r1-to-r5 to 192.168.5.1;**

4. Configure the primary path:

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

   For example:

   [edit protocols mpls label-switched-path r1-to-r5]
   user@R1# **set primary via-r2**

5. Configure the bandwidth:

   [edit protocols mpls label-switched-path *lsp-path-name*]
   user@host# **set primary** *primary-name* **bandwidth** *bandwidth*

   For example:

   [edit protocols mpls label-switched-path r1-to-r5]
   user@R1# **set primary via-r2 bandwidth 35m**

6. Configure the priority value:

   [edit protocols mpls label-switched-path *lsp-path-name*]
   user@host# **set primary** *primary-name* **priority** *reservation-priority setup-priority*

   For example:

   [edit protocols mpls label-switched-path r1-to-r5]
   user@R1# **set primary via-r2 priority 6 6**

7. Verify and commit the configuration:

   [edit protocols mpls label-switched-path *lsp-path-name*]
   user@host# **show**
   user@host# **commit**

Sample Output The sample output below illustrates the configuration of the primary path on ingress router R1 in the network shown in Figure 2 on page 12.

```
edit protocols mpls]
user@R1# show
label-switched-path r1-to-r5 {
    to 192.168.5.1;
    primary via-r2 {    # Bandwidth and priority configured at the primary path
        bandwidth 35m; # level of the hierarchy
        priority 6 6; # Priority setup and hold values
    }
}
path via-r2 { # Primary ERO list
    10.0.12.14 strict;
    10.0.24.2 strict;
[...Output truncated...]

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

What It Means The sample output shows a label-switched path (LSP) with bandwidth and priority applied to only one primary path. The same parameters specified one level up in the hierarchy, at the [edit protocols mpls label-switched-path *lsp-path-name*] hierarchy level, affect all paths.

The path, via-r2, specifies the complete strict path from the ingress to the egress routers through 10.0.12.14, 10.0.24.2, in that order. There cannot be any intermediate routers except the ones specified. However, there can be intermediate routers between 10.0.24.2 and the egress router because the egress router is not specifically listed in the path statement. To prevent intermediate routers before egress, configure the egress router as the last router, with a strict type.

For more information on configuring a primary path, see the *JUNOS MPLS Applications Configuration Guide*.

Step 2: Verify That the Primary Path Is Operational

Primary paths must always be used in the network if they are available, therefore an LSP always moves back to the primary path after a failure, unless the configuration is adjusted. For more information on adjusting the configuration to prevent a failed primary path from reestablishing, see "Preventing Use of a Path That Previously Failed" on page 22.

Action To verify that the primary path is operational, enter the following JUNOS command-line interface (CLI) operational mode commands:

```
user@host> show mpls lsp extensive ingress
user@host> show rsvp interface
```

Sample Output 1

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

192.168.5.1
  From: 192.168.1.1, State: Up, ActiveRoute: 0, LSPname: r1-to-r5
  ActivePath: via-r2 (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
```

```
 *Primary   via-r2           State: Up
    Priorities: 6 6
    Bandwidth: 35Mbps
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 11)
 10.0.12.14 S 10.0.24.2 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node
10=SoftPreempt):
          10.0.12.14 10.0.24.2
    5 Apr 29 14:40:43 Selected as active path
    4 Apr 29 14:40:43 Record Route:  10.0.12.14 10.0.24.2
    3 Apr 29 14:40:43 Up
    2 Apr 29 14:40:43 Originate Call
    1 Apr 29 14:40:43 CSPF: computation result accepted
  Standby   via-r7           State: Dn
    SmartOptimizeTimer: 180
        No computed ERO.
  Created: Sat Apr 29 14:40:43 2006
Total 1 displayed, Up 1, Down 0
```

Sample Output 2

```
user@R1> show rsvp interface
RSVP interface: 3 active
                  Active Subscr- Static      Available   Reserved    Highwater
Interface   State resv   iption  BW          BW          BW          mark
fe-0/1/0.0  Up         2   100%  100Mbps     100Mbps     0bps        0bps
fe-0/1/1.0  Up         1   100%  100Mbps     100Mbps     0bps        0bps
so-0/0/3.0  Up         1   100%  155.52Mbps  155.52Mbps  0bps        0bps
```

What It Means

Sample output 1 shows that the LSP is operational and is using the primary path (via-r2) with R2 (10.0.12.14) and R4 (10.0.24.2) as transit routers. The priority values are the same for setup and hold, 6 6. Priority 0 is the highest (best) priority and 7 is the lowest (worst) priority. The JUNOS software default for setup and hold priority is 7:0. Unless some LSPs are more important than others, preserving the default is a good practice. Configuring a setup priority that is better than the hold priority is not allowed, resulting in a failed commit in order to avoid preemption loops.

Configuring and Verifying a Secondary Path

Secondary paths (also known as secondary LSPs) are optional and protect against link and transit node failures. If the primary path can no longer reach the egress router, the alternative, secondary path is used, as shown in Figure 3.

Figure 3: Standby Secondary Paths

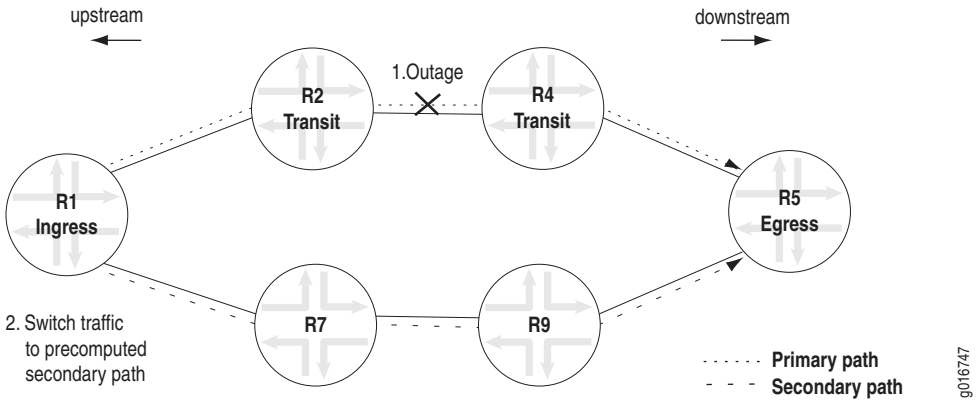

In Figure 3, a secondary path R1-R7-R9-R5 is activated when the primary path R1-R2-R4-R5 fails. R2 notifies R1 of the outage and R1 switches traffic to the precomputed secondary path.

Two types of secondary paths, standby and non-standby, can become active when a primary path fails, depending on which is configured. A standby secondary path, configured with the standby statement, is precomputed and pre-signaled. A non-standby secondary path, configured without the standby statement, is precomputed but is *not* pre-signaled.

Secondary paths configured with the standby statement consume more resources because the router must maintain state when the secondary path is not active. However, standby secondary paths do reduce recovery time by eliminating the call-setup delay that is required to establish a new physical path for the LSP.

If the problem with the primary path is corrected, after a few minutes of hold-down to ensure that the primary path remains stable, the ingress router switches traffic from the secondary path back to the primary path. It may not be always prudent for the router to switch back to the primary path. For information on how to keep the router from switching back to the primary path, see “Preventing Use of a Path That Previously Failed” on page 22.

To configure and verify a secondary path, follow these steps:

1. Configure a Standby Secondary Path on page 18
2. Verify That the Secondary Path Is Established on page 19

Step 1: Configure a Standby Secondary Path

Configuring a standby secondary path is a two-part process. In the first part, you define the path, and in the second part, you specify a secondary path for the LSP that refers to the defined path.

NOTE: To configure a non-standby secondary path, simply omit the standby statement from the secondary path configuration.

To configure a standby secondary path, follow these steps:

Action

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

   [edit]
   user@host# **edit protocols mpls**

2. Configure the secondary ERO list:

   [edit protocols mpls]
   user@host# **set path** *path-name destination* **loose**

   For example:

   [edit protocols mpls]
   user@R1# **set path via-r7 10.0.17.14 loose**

3. Configure the LSP and the secondary path:

   [edit protocols mpls]
   user@host# **set label-switched-path** *lsp-path-name* **secondary** *secondary-name* **standby**

   For example:

   [edit protocols mpls]
   user@R1# **set label-switched-path r1-to-r4 secondary via-r7 standby**

4. Verify and commit the configuration:

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

The sample output below illustrates the configuration of the standby secondary path on ingress router R1 in the network shown in Figure 2 on page 12.

Sample Output

```
[edit protocols mpls]
user@R1# show
label-switched-path r1-to-r4 {
    to 192.168.4.1;
    ldp-tunneling;
    fast-reroute;
    primary via-r2;
    secondary via-r7 {
        standby; # Omit the standby statement to configure a non-standby secondary path
    }
}
```

```
path via-r2 {
    10.0.12.14 loose;
}
path via-r7 {
    10.0.17.14 loose;
}
[...Output truncated...]
```

What It Means The sample output shows one standby secondary path via-r7, which includes the standby statement at the [edit protocols mpls label-switched-path *lsp-path-name* secondary *secondary-name*] hierarchy level. The standby secondary path is defined in the path statement path via-r7 and specifies a loose hop, indicating that the route taken from the previous router to this router need not be a direct path, can include other routers, and can be received on any interface.

If you have many secondary paths configured for an LSP, and you want them all to be standby, include the standby statement one level up in the hierarchy, at the [edit protocols mpls label-switched-path *lsp-path-name*] hierarchy level, as shown in the sample output below.

```
[edit protocols mpls]
user@R1# show
label-switched-path r1-to-r4 {
    to 192.168.4.1;
    standby; # Standby configured at the label-switched-path level of the hierarchy
    primary via-r2;
    }
    secondary via-r7;
    }
}
[...Output truncated...]
```

For more information on configuring a secondary path, see the *JUNOS MPLS Applications Configuration Guide*.

Step 2: Verify That the Secondary Path Is Established

When the secondary path is configured with the standby statement, the secondary path should be *up* but *not active*; it will become active if the primary path fails. A secondary path configured without the standby statement will not come up unless the primary path fails. To test that the secondary path is correctly configured and would come up if the primary path were to fail, you must deactivate a link or node critical to the primary path, then issue the show mpls lsp *lsp-path-name* extensive command.

Action To verify that the secondary path is established, enter the following JUNOS CLI operational mode command:

```
user@R1>show mpls lsp extensive
```

Sample Output The following sample output shows a correctly configured secondary path before and after it comes up. In the example, interface fe-0/1/0 on R2 is deactivated, which brings down the primary path via-r2. The ingress router R1 switches traffic to the secondary path via-r7.

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

192.168.5.1
  From: 192.168.1.1, State: Up, ActiveRoute: 0, LSPname: r1-to-r5
  ActivePath: via-r2 (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   via-r2           State: Up
    Priorities: 6 6
    Bandwidth: 35Mbps
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
 10.0.12.14 S 10.0.24.2 S 10.0.45.2 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node
10=SoftPreempt):
          10.0.12.14 10.0.24.2 10.0.45.2
    5 Apr 29 14:40:43 Selected as active path
    4 Apr 29 14:40:43 Record Route:  10.0.12.14 10.0.24.2
    3 Apr 29 14:40:43 Up
    2 Apr 29 14:40:43 Originate Call
    1 Apr 29 14:40:43 CSPF: computation result accepted
  Secondary via-r7          State: Dn
    SmartOptimizeTimer: 180
        No computed ERO.
  Created: Sat Apr 29 14:40:43 2006
Total 1 displayed, Up 1, Down 0

[edit interfaces]
user@R2# deactivate fe-0/1/0

[edit interfaces]
user@R2# show
inactive: fe-0/1/0 {
    unit 0 {
        family inet {
            address 10.0.12.14/30;
        }
        family iso;
        family mpls;
    }
}

user@R1> show mpls lsp name r1-to-r4 extensive
Ingress LSP: 1 sessions

192.168.4.1
  From: 192.168.1.1, State: Up, ActiveRoute: 0, LSPname: r1-to-r4
  ActivePath: via-r7 (secondary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary   via-r2          State: Dn
    Priorities: 6 6
    Bandwidth: 35Mbps
    SmartOptimizeTimer: 180
    Will be enqueued for recomputation in 14 second(s).
   10 Apr 29 14:52:33 CSPF failed: no route toward 10.0.12.14[21 times]
```

```
    9 Apr 29 14:42:48 Clear Call
    8 Apr 29 14:42:48 Deselected as active
    7 Apr 29 14:42:48 Session preempted
    6 Apr 29 14:42:48 Down
    5 Apr 29 14:40:43 Selected as active path
    4 Apr 29 14:40:43 Record Route:  10.0.12.14 10.0.24.2
    3 Apr 29 14:40:43 Up
    2 Apr 29 14:40:43 Originate Call
    1 Apr 29 14:40:43 CSPF: computation result accepted
 *Standby   via-r7           State: Up
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 11)
 10.0.17.14 S 10.0.47.1 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node
10=SoftPreempt):
          10.0.17.14 10.0.47.1
    5 Apr 29 14:42:48 Selected as active path
    4 Apr 29 14:41:12 Record Route:  10.0.17.14 10.0.47.1
    3 Apr 29 14:41:12 Up
    2 Apr 29 14:41:12 Originate Call
    1 Apr 29 14:41:12 CSPF: computation result accepted
  Created: Sat Apr 29 14:40:43 2006
Total 1 displayed, Up 1, Down 0
```

What It Means The sample output from egress router R1 shows a correctly configured standby secondary path in a down state because the primary path is still up. Upon deactivation of an interface (interface fe-0/1/0 on R2) critical to the primary path, the primary path via-r2 goes down and the standby secondary path via-r7 comes up, allowing R1 to switch traffic to the standby secondary path.

Ensuring That Secondary Paths Establish When Resources Are Diminished

The JUNOS software does not require that a primary and secondary path share the same parameters. You may decide to configure your primary paths with strict resource requirements, and configure your secondary paths with less strict requirements, allowing your secondary paths to establish more readily during periods of diminished resources.

Action To ensure that secondary paths establish when resources are diminished, follow these steps:

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

   [edit protocols mpls]
   user@host# **edit label-switched-path** *lsp-path-name*

   For example:

   [edit protocols mpls]
   user@R1# **edit label-switched-path r1-to-r4**

2. Configure the bandwidth for the primary path, and do not configure any bandwidth for the secondary path:

   [edit protocols mpls label-switched-path *lsp-path-name*]
   user@host# **set primary** *primary-name* **bandwidth** *bandwidth*

For example:

```
[edit protocols mpls label-switched-path r1-to-r4]
user@R1# set primary via-r2 bandwidth 35m
```

3. Verify and commit the configuration:

```
[edit protocols mpls label-switched-path lsp-path-name]
user@host# show
user@host# commit
```

Sample Output The sample output below illustrates a bandwidth configuration on ingress router R1 in the network shown in Figure 2 on page 12.

```
[edit protocols mpls]
user@R1# show
label-switched-path r1-to-r4 {
    to 192.168.4.1;
    primary via-r2 {
        bandwidth 35m;
    }
    secondary via-r7 { # In this example, bandwidth is not configured for the secondary path.
        standby;       # However you could configure a bandwidth value different from
    }                  # that on the primary path.
}
[...Output Truncated...]
```

What It Means The sample output shows the primary path via-r2 requires 35 Mbps of bandwidth, while secondary path via-r7 has no constraints. The primary path is configured with strict resource requirements, while the secondary path is configured with no bandwidth requirements, allowing the secondary path to establish more readily during periods of diminished resources. One thing to keep in mind when configuring a secondary path without bandwidth requirements is that it can be subject to traffic loss due to congestion.

Preventing Use of a Path That Previously Failed

If you configure an alternate path through the network in case the active path fails, you may not want traffic to revert back to the failed path, even if it is no longer failing. When you configure a primary path, the traffic switches over to the secondary path during a failure, and reverts back to the primary path when it returns.

At times, switching traffic back to a primary path that has previously failed may not be a particularly sound idea. In this case, only configure secondary paths, resulting in the next configured secondary path establishing when the first secondary path fails. Later, if the first secondary path becomes operational, the JUNOS software will not revert to it, but will continue using the second secondary path.

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*