Chapter 7

Troubleshooting Link Protection for Multiple Bypass LSPs Overview

This case study simulates a network problem with link protection for multiple bypass paths for Resource Reservation Protocol (RSVP)-signaled LSPs (LSPs). It includes a brief summary of link protection, an example network scenario, and commands to troubleshoot and resolve the problem. (See Table 10.)

The troubleshooting process described in this case study should not be followed rigidly; it is a basis from which you can develop your own process to suit your particular situation.

Table 10: Troubleshooting Link Protection for Multiple Bypass LSPs Checklist

Troubleshooting Tasks	Command or Action
Troubleshooting Link Protection for Multiple Bypass LSPs on page 152	
■ Symptom on page 153	One bypass LSP is pre-signaled instead of two. show mpls lsp bypass
■ Cause on page 153	The bandwidth reserved on the primary LSP is served by only one bypass path.
■ Troubleshooting Commands on page 153	show mpls lsp show mpls lsp bypass extensive show rsvp session ingress detail show rsvp interface show rsvp interface *type-fpc/pic/port* extensive show configuration *statement-path*
■ Solution on page 160	show configuration *statement-path* show mpls lsp bypass show mpls lsp bypass extensive
■ Conclusion on page 163	Multiple bypass paths are pre-signaled when the bandwidth values in the configuration require multiple bypass paths.
■ Router Configurations on page 163	show configuration | no-more

Troubleshooting Link Protection for Multiple Bypass LSPs

Link protection (many-to-one or facility backup) allows a router immediately upstream from a link failure to use an alternate interface to forward traffic to its downstream neighbor. This is accomplished by preestablishing a bypass path that is shared by all protected LSPs traversing the failed link. A single bypass path can safeguard a set of protected LSPs. When an outage occurs, the router immediately upstream from the link outage switches protected traffic to the bypass link, and then signals the link failure to the ingress router.

In this simulation, the network administrator mistakenly expects two bypass paths to be pre-signaled to protect four LSPs over two interfaces. However, because of the bandwidth configuration on both interfaces and the RSVP protocol, only one bypass path is pre-signaled. The second bypass path is not pre-signaled because the existing bandwidth reserved on the primary LSP is served by one bypass path.

Figure 15 illustrates the network topology used in this case study.

Figure 15: Link Protection for Multiple Bypass LSPs Network

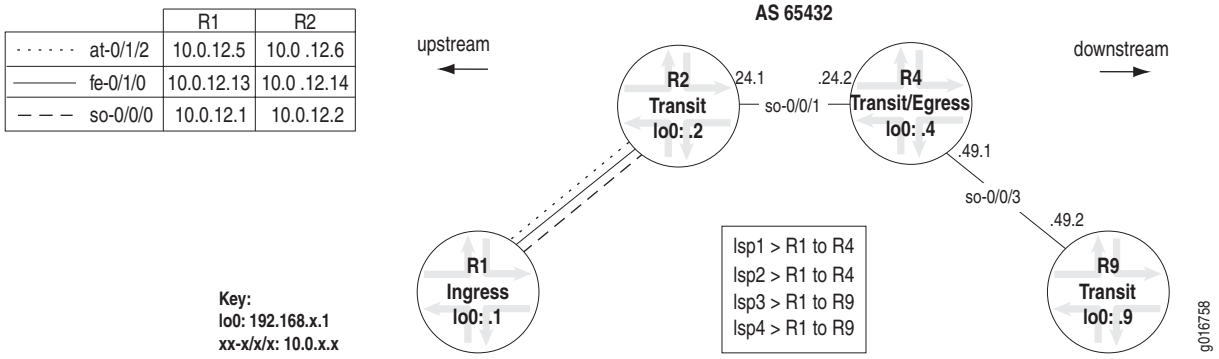

The MPLS network topology in Figure 15 shows a router-only network with SONET, Fast Ethernet, and ATM interfaces that consists of the following components:

- A full-mesh internal BGP (IBGP) topology using AS 65432
- MPLS and RSVP are enabled on all routers
- A send-statics policy on routers R1, R4, and R9 that allows a new route to be advertised into the network
- Six unidirectional LSPs between R1 and R4, and R1 and R9, with two LSPS running in the opposite direction to allow for bidirectional traffic
- Three interface connections between R1 and R4, which allows for a primary LSP and two bypass paths on different interfaces
- Bandwidth configured for interfaces, RSVP, and LSPs

Sample configurations for all four routers in the network shown in Figure 15 are provided at the end of this case study in “Router Configurations” on page 163.

Symptom

In the network shown in Figure 15 on page 152, only one bypass LSP is pre-signaled instead of two, as shown in the following sample output.

user@R1> **show mpls lsp bypass**

```
user@R1> show mpls lsp bypass
Ingress LSP: 5 sessions
To              From            State   Rt Style Labelin Labelout LSPname
192.168.2.1     192.168.1.1     Up       0  1 SE       -        3 Bypass->10.0.12.14
Total 1 displayed, Up 1, Down 0

Egress LSP: 2 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Cause

The cause of this problem is that bandwidth reserved for the primary LSPs is served by only one bypass path.

Troubleshooting Commands

The JUNOS software includes commands that are useful when troubleshooting a problem. This section provides a brief description of each command, followed by sample output, and a discussion of the output in relation to the network shown in Figure 15 on page 152.

The following commands can be used when troubleshooting:

user@host> **show mpls lsp**
user@host> **show mpls lsp bypass extensive**
user@host> **show rsvp session ingress detail**
user@host> **show rsvp interface**
user@host> **show rsvp interface** *type-fpc/pic/port* **extensive**
user@host> **show configuration** *statement-path*

Sample Output Use show mpls lsp command to display configured LSPs on the router, as well as all ingress, transit, and egress LSPs.

```
user@R1> show mpls lsp
Ingress LSP: 4 sessions
To              From            State Rt ActivePath       P     LSPname
192.168.4.1     192.168.1.1     Up     0 path1            *     lsp1
192.168.4.1     192.168.1.1     Up     0 path1            *     lsp2
192.168.9.1     192.168.1.1     Up     0 path1            *     lsp3
192.168.9.1     192.168.1.1     Up     0 path1            *     lsp4
Total 4 displayed, Up 4, Down 0

Egress LSP: 2 sessions
To              From            State   Rt Style Labelin Labelout LSPname
192.168.1.1     192.168.4.1     Up       0  1 FF       3        - r4-r1
192.168.1.1     192.168.9.1     Up       0  1 FF       3        - r9-r1
Total 2 displayed, Up 2, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

What It Means The sample output of the show mpls lsp command shows that four LSPs originating from this router R1 are up (ingress LSPs). The two LSPs originating at R4 and R9, and terminating at R1 are also up (egress LSPs). No LSPs are transiting this router (transit LSPs). In this case, all LSPs are up, indicating that the problem is not with the LSPs being in a down state.

Sample Output Use show mpls lsp bypass extensive command to display detailed information about LSPs used for protecting other LSPs (bypass LSPs).

```
user@R1> show mpls lsp bypass extensive
Ingress LSP: 5 sessions

192.168.2.1
  From: 192.168.1.1, LSPstate: Up, ActiveRoute: 0
  LSPname: Bypass->10.0.12.14 #This bypass path is from R1 to R2
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 SE, Label in: -, Label out: 3
  Time left:    -, Since: Fri Nov 10 08:29:27 2006
  Tspec: rate 100Mbps size 100Mbps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 45808 protocol 0
  Type: Bypass LSP
    Number of data route tunnel through: 4 #LSPs protected by this bypass path
    Number of RSVP session tunnel through: 0
  ActiveResv 4, PreemptionCnt 0, Update threshold 0%
  Subscription 100%,
  bc0 = ct0, StaticBW 100Mbps
  ct0: StaticBW 100Mbps, AvailableBW 60Mbps
    MaxAvailableBW 100Mbps = (bc0*subscription)
    ReservedBW [0] 40Mbps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7]0bps
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.12.2 (so-0/0/0.0) 57 pkts
  RESV rcvfrom: 10.0.12.2 (so-0/0/0.0) 57 pkts
  Explct route: 10.0.12.2
  Record route: <self> 10.0.12.2
Total 1 displayed, Up 1, Down 0

Egress LSP: 2 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

What It Means The sample output of the show mpls lsp bypass extensive command shows one bypass LSP (Bypass->10.0.12.14) from ingress router R1 to transit router R2. All four of the ingress LSPs are protected by this single bypass path, as indicated by the Number of data route tunnel through: 4 field. Interface so-0/0/0.0 is the interface on which the bypass is pre-signaled. In this case study, the problem is that interface at-0/1/2.0 is supposed to also have a pre-signaled bypass path, and the two LSPs should be protected by a bypass path on each interface (so-0/0/0.0 and at-0/1/2.0).

Sample Output Use show rsvp session ingress detail command to display detailed information about RSVP sessions.

```
user@R1> show rsvp session ingress detail
Ingress RSVP: 5 sessions

192.168.2.1
  From: 192.168.1.1, LSPstate: Up, ActiveRoute: 0
  LSPname: Bypass->10.0.12.14
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 SE, Label in: -, Label out: 3
  Time left:    -, Since: Fri Nov 10 08:29:27 2006
  Tspec: rate 100Mbps size 100Mbps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 45808 protocol 0
  Type: Bypass LSP
    Number of data route tunnel through: 4
    Number of RSVP session tunnel through: 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.12.2 (so-0/0/0.0) 60 pkts
  RESV rcvfrom: 10.0.12.2 (so-0/0/0.0) 61 pkts
  Explct route: 10.0.12.2
  Record route: <self> 10.0.12.2

192.168.4.1
  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: 101008
  Resv style: 1 SE, Label in: -, Label out: 101008
  Time left:    -, Since: Thu Nov  9 11:39:04 2006
  Tspec: rate 10Mbps size 10Mbps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 45673 protocol 0
  Link protection desired
  Type: Link protected LSP
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.12.14 (fe-0/1/0.0) 1880 pkts
  RESV rcvfrom: 10.0.12.14 (fe-0/1/0.0) 1838 pkts
  Explct route: 10.0.12.14 10.0.24.2
  Record route: <self> 10.0.12.14 10.0.24.2

192.168.4.1
  From: 192.168.1.1, LSPstate: Up, ActiveRoute: 0
  LSPname: lsp2, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 101104
  Resv style: 1 SE, Label in: -, Label out: 101104
  Time left:    -, Since: Thu Nov  9 20:34:02 2006
  Tspec: rate 10Mbps size 10Mbps peak Infbps m 20 M 1500
  Port number: sender 2 receiver 45675 protocol 0
  Link protection desired
  Type: Link protected LSP
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.12.14 (fe-0/1/0.0) 1076 pkts
  RESV rcvfrom: 10.0.12.14 (fe-0/1/0.0) 1068 pkts
  Explct route: 10.0.12.14 10.0.24.2
```

```
  Record route: <self> 10.0.12.14 10.0.24.2

192.168.9.1
  From: 192.168.1.1, LSPstate: Up, ActiveRoute: 0
  LSPname: lsp3, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 101120
  Resv style: 1 SE, Label in: -, Label out: 101120
  Time left:    -, Since: Thu Nov  9 20:34:02 2006
  Tspec: rate 10Mbps size 10Mbps peak Infbps m 20 M 1500
  Port number: sender 2 receiver 45685 protocol 0
  Link protection desired
  Type: Link protected LSP
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.12.14 (fe-0/1/0.0) 1080 pkts
  RESV rcvfrom: 10.0.12.14 (fe-0/1/0.0) 1072 pkts
  Explct route: 10.0.12.14 10.0.24.2 10.0.49.2
  Record route: <self> 10.0.12.14 10.0.24.2 10.0.49.2

192.168.9.1
  From: 192.168.1.1, LSPstate: Up, ActiveRoute: 0
  LSPname: lsp4, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 101136
  Resv style: 1 SE, Label in: -, Label out: 101136
  Time left:    -, Since: Thu Nov  9 20:34:02 2006
  Tspec: rate 10Mbps size 10Mbps peak Infbps m 20 M 1500
  Port number: sender 2 receiver 45687 protocol 0
  Link protection desired
  Type: Link protected LSP
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.12.14 (fe-0/1/0.0) 1076 pkts
  RESV rcvfrom: 10.0.12.14 (fe-0/1/0.0) 1068 pkts
  Explct route: 10.0.12.14 10.0.24.2 10.0.49.2
  Record route: <self> 10.0.12.14 10.0.24.2 10.0.49.2
Total 5 displayed, Up 5, Down 0
```

What It Means The sample output of the show rsvp session ingress detail command shows five RSVP sessions originating at ingress router R1. Each session is up and each LSP is protected by the bypass path Bypass->10.0.12.14 on interface so-0/0/0.0. In this case study, two of the LSPs should be protected by a second bypass path on interface at-0/1/2.0.

Sample Output Use show rsvp interface command to display the status of RSVP-enabled interfaces and packet statistics.

```
user@R1> show rsvp interface
RSVP interface: 3 active
                  Active Subscr- Static      Available   Reserved    Highwater
Interface   State resv   iption  BW          BW          BW          mark
at-0/2/1.0  Up         0   100%  50Mbps      50Mbps      0bps        0bps
fe-0/1/0.0  Up         4   100%  100Mbps     60Mbps      40Mbps      100Mbps
so-0/0/0.0  Up         1   100%  100Mbps     0bps        100Mbps     100Mbps
```

What It Means The sample output of the show rsvp interface command shows that all RSVP interfaces are up with four reservations on the Fast Ethernet interface (fe-0/1/0.0), one reservation on the SONET interface (so-0/0/0.0), and no reservations on the ATM interface at-0/2/1.0. The total interface bandwidth (Static BW) is 100 Mbps on the Fast Ethernet and SONET interfaces, and only 50 Mbps on the ATM interface, indicating that the SONET interface is providing enough bandwidth to satisfy the requirements of the primary path of all four LSPs. Therefore, there is no need for a second bypass path on the ATM interface with this configuration.

Sample Output Use show rsvp interface *interface-name* extensive command to display detailed information about a specific interface. The extensive option provides output for the latest 50 events on this interface.

```
user@R1> show rsvp interface fe-0/1/0.0 extensive
fe-0/1/0.0 Index 66, State Ena/Up
  NoAuthentication, NoAggregate, NoReliable, LinkProtection
  HelloInterval 9(second)
  Address 10.0.12.13
  ActiveResv 4, PreemptionCnt 0, Update threshold 10%
  Subscription 100%,
  bc0 = ct0, StaticBW 100Mbps
  ct0: StaticBW 100Mbps, AvailableBW 60Mbps
    MaxAvailableBW 100Mbps = (bc0*subscription)
    ReservedBW [0] 40Mbps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7]
0bps
  Protection: On, Bypass: 1, LSP: 4, Protected LSP: 0, Unprotected LSP: 4
      1 Nov 10 09:48:12 New bypass Bypass->10.0.12.14
    Bypass: Bypass->10.0.12.14, State: Up, Type: LP, LSP: 0, Backup: 0
      4 Nov 10 09:49:13 Record Route:  10.0.12.2
      3 Nov 10 09:49:13 Up
      2 Nov 10 09:49:13 CSPF: computation result accepted
      1 Nov 10 09:48:43 CSPF failed: no route toward 10.0.12.14[2 times]
```

What It Means The sample output of the show rsvp interface *interface-name* extensive command shows one bypass path protecting four LSPs. The bypass path is pre-signaled on 10.0.12.2, which is the SONET interface so-0/0/0.0. Also, the total amount of bandwidth that RSVP is allowed to reserve is 100 Mbps. 40 Mbps are reserved with 60 Mbps available, indicating that there is more than enough bandwidth available to meet the needs of the four LSPs with one bypass path.

Sample Output Use show configuration *statement-path* command to display a specific configuration hierarchy; for example, routing protocols..

```
user@R1> show configuration protocols rsvp
interface fe-0/1/0.0 {
    link-protection {
        bandwidth 100m;
        max-bypasses 2;
    }
}
interface so-0/0/0.0;
interface at-0/2/1.0;
interface fxp0.0 {
    disable;
}
```

What It Means The sample output of the show configuration protocols rsvp command shows that the Fast Ethernet interface is configured with link protection, 100 Mbps of bandwidth, and two bypass paths. In this case study, the amount of bandwidth may need to be adjusted until two bypass paths are pre-signaled.

Sample Output Use show configuration *statement-path* command to display a specific configuration hierarchy; for example, interfaces.

```
user@R1> show configuration interfaces
so-0/0/0 {
    unit 0 {
        bandwidth 100m;
        family inet {
            address 10.0.12.1/32;
        }
        family mpls;
    }
}
fe-0/1/0 {
    unit 0 {
        family inet {
            address 10.0.12.13/30;
        }
        family mpls;
    }
}
at-0/2/1 {
    atm-options {
        pic-type atm2;
        vpi 0;
    }
    unit 0 {
        bandwidth 50m;
        vci 0.128;
        family inet {
            address 10.0.12.5/32 {
                destination 10.0.12.6;
            }
        }
        family mpls;
    }
}
[...Output truncated...]
```

What It Means The sample output of the show configuration interface command shows that the SONET interface is configured with 100 Mbps, while the ATM interface is configured with 50 Mbps. In this case study, the amount of bandwidth for each interface may need to be adjusted until two bypass paths are pre-signaled.

Sample Output Use show configuration *statement-path* command to display a specific configuration hierarchy; for example, routing protocols.

```
user@R1> show configuration protocols mpls
mpls {
        label-switched-path lsp1 {
            from 192.168.1.1;
            to 192.168.4.1;
            bandwidth 10m;
            link-protection;
            primary path1;
        }
        label-switched-path lsp2 {
            from 192.168.1.1;
            to 192.168.4.1;
            bandwidth 20m;
            link-protection;
            primary path1;
        }
        label-switched-path lsp3 {
            from 192.168.1.1;
            to 192.168.9.1;
            bandwidth 30m;
            link-protection;
            primary path1;
        }
        label-switched-path lsp4 {
            from 192.168.1.1;
            to 192.168.9.1;
            bandwidth 40m;
            link-protection;
            primary path1;
        }
        path path1 {
            10.0.12.14 strict;
        }
        interface fe-0/1/0.0;
        interface so-0/0/0.0;
        interface at-0/2/1.0;
        interface fxp0.0 {
            disable;
        }
    }
```

What It Means The sample output of the show configuration protocols mpls command shows that the four LSPs are configured with different bandwidth values. In this case study, the bandwidth value for each LSP may need to be adjusted until two bypass paths are pre-signaled.

Solution

Adjust the bandwidth for the interfaces, RSVP link protection, and LSPs until two bypass LSPs are pre-signaled. In this case study, the bandwidth value for the SONET interface and the protected Fast Ethernet interface was reduced. An adjustment was not made to the bandwidth of the LSPs. The bandwidth adjustment described in this case study should not be followed rigidly; it is a basis from which you can develop your own process of adjusting bandwidth that suits your particular situation.

For information on adjusting the bandwidth for interfaces, see the *JUNOS Network Interfaces Configuration Guide*. For information on adjusting the bandwidth for link protection, see "Load Balancing in an MPLS Network" on page 65. For information on adjusting the LSP bandwidth, see "Path Protection in an MPLS Network" on page 9.

The JUNOS software includes commands that are useful when verifying the solution to a problem. This section provides a brief description of each command, followed by sample output, and a discussion of the output in relation to the network shown in Figure 15 on page 152.

You can use the following commands when verifying the solution to a problem:

user@host> **show configuration** *statement-path*
user@host> **show mpls lsp bypass**
user@host> **show mpls lsp bypass extensive**

Sample Output Use show configuration *statement-path* command to display a specific configuration hierarchy; for example, interfaces.

```
user@R1> show configuration interfaces
so-0/0/0 {
    unit 0 {
        bandwidth 50m;
        family inet {
            address 10.0.12.1/32;
        }
        family mpls;
    }
}
[...Output truncated...]
at-0/2/1 {
    atm-options {
        pic-type atm2;
        vpi 0;
    }
    unit 0 {
        bandwidth 50m;
        vci 0.128;
        family inet {
            address 10.0.12.5/32 {
                destination 10.0.12.6;
            }
        }
        family mpls;
    }
}
[...Output truncated...]
```

What It Means The sample output of the show configuration interfaces command shows that the bandwidth for the SONET interface has been adjusted down from 100 Mbps to 50 Mbps. This adjustment did not in itself result in two bypass paths coming up. A further adjustment to the link protection bandwidth was necessary before two bypass paths were pre-signaled.

Sample Output Use show configuration *statement-path* command to display a specific configuration hierarchy; for example, routing protocols.

```
user@R1> show configuration protocols
protocols {
    rsvp {
        interface fe-0/1/0.0 {
            link-protection {
                bandwidth 50m;
                max-bypasses 2;
            }
        }
        interface fe-0/1/2.0;
        interface so-0/0/0.0;
        interface at-0/2/1.0;
        interface fxp0.0 {
            disable;
        }
    }
[...Output truncated...]
```

What It Means The sample output of the show configuration interfaces command shows that the bandwidth for link protection on the Fast Ethernet interface has been adjusted down from 100 Mbps to 50 Mbps.

Sample Output Use show mpls lsp bypass command to display information about LSPs used for protecting other LSPs (bypass LSPs).

```
user@R1> show mpls lsp bypass
Ingress LSP: 6 sessions
To              From            State   Rt Style Labelin Labelout LSPname
192.168.2.1     192.168.1.1     Up       0  1 SE       -        3 Bypass->10.0.12.14
192.168.2.1     192.168.1.1     Up       0  1 SE       -        3 Bypass->10.0.12.14-1
Total 2 displayed, Up 2, Down 0

Egress LSP: 2 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

What It Means The sample output of the show configuration interfaces command shows that two bypass LSPs are pre-signaled (Up), 10.0.12.14 and 10.0.12.14-1, indicating that reducing the bandwidth of the link-protected interface and the SONET interface was successful. The bandwidth adjustment made in this case study may be different from the adjustment that is required in your network.

Sample Output Use show mpls lsp bypass extensive command to display detailed information about LSPs used for protecting other LSPs (bypass LSPs). The no-more option entered after the pipe (|) prevents the output from being paginated if the output is longer than the length of the terminal screen.

```
user@R1> show mpls lsp bypass extensive | no-more
Ingress LSP: 6 sessions

192.168.2.1
  From: 192.168.1.1, LSPstate: Up, ActiveRoute: 0
  LSPname: Bypass->10.0.12.14
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 SE, Label in: -, Label out: 3
  Time left:    -, Since: Thu Nov  9 17:47:17 2006
  Tspec: rate 50Mbps size 50Mbps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 45762 protocol 0
  Type: Bypass LSP
    Number of data route tunnel through: 2
    Number of RSVP session tunnel through: 0
  ActiveResv 2, PreemptionCnt 0, Update threshold 0%
  Subscription 100%,
  bc0 = ct0, StaticBW 50Mbps
  ct0: StaticBW 50Mbps, AvailableBW 0bps
    MaxAvailableBW 50Mbps = (bc0*subscription)
    ReservedBW [0] 50Mbps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7]0bps
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.12.6 (at-0/2/1.0) 213 pkts
  RESV rcvfrom: 10.0.12.6 (at-0/2/1.0) 213 pkts
  Explct route: 10.0.12.6
  Record route: <self> 10.0.12.6

192.168.2.1
  From: 192.168.1.1, LSPstate: Up, ActiveRoute: 0
  LSPname: Bypass->10.0.12.14-1
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 SE, Label in: -, Label out: 3
  Time left:    -, Since: Thu Nov  9 17:47:51 2006
  Tspec: rate 50Mbps size 50Mbps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 45764 protocol 0
  Type: Bypass LSP
    Number of data route tunnel through: 2
    Number of RSVP session tunnel through: 0
  ActiveResv 2, PreemptionCnt 0, Update threshold 0%
  Subscription 100%,
  bc0 = ct0, StaticBW 50Mbps
  ct0: StaticBW 50Mbps, AvailableBW 0bps
    MaxAvailableBW 50Mbps = (bc0*subscription)
    ReservedBW [0] 50Mbps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7]0bps
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.12.2 (so-0/0/0.0) 212 pkts
  RESV rcvfrom: 10.0.12.2 (so-0/0/0.0) 212 pkts
  Explct route: 10.0.12.2
  Record route: <self> 10.0.12.2
Total 2 displayed, Up 2, Down 0

Egress LSP: 2 sessions
```

```
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

What It Means The sample output of the show mpls lsp bypass extensive command shows two bypass LSPs (Bypass->10.0.12.14 and Bypass->10.0.12.14-1) from ingress router R1 to transit router R2. All four of the ingress LSPs are protected by the two bypass paths, as indicated by the Number of data route tunnel through: 2 field in the output for each bypass LSP. The SONET interface so-0/0/0.0 and the ATM interface at-0/1/2.0 are the interfaces on which the bypass paths are pre-signaled.

Conclusion

In this simulation, the network administrator mistakenly expected two bypass paths to be pre-signaled when the bandwidth configuration on the interfaces and the RSVP protocol required only one bypass path. After troubleshooting the example network scenario, and adjusting the bandwidth for the interfaces and link protection in the RSVP protocol, the second bypass path was pre-signaled, and the problem resolved.

In conclusion, multiple bypass paths are pre-signaled when the bandwidth values in the configuration require multiple bypass paths.

Router Configurations

Purpose Output that shows the configurations of all routers in the network. The no-more option entered after the pipe (|) prevents the output from being paginated if the output is longer than the length of the terminal screen.

Sample Output 1 The following sample output is for ingress router R1:

```
user@R1> show configuration | no-more
interfaces {
    so-0/0/0 {
        unit 0 {
            bandwidth 50m;
            family inet {
                address 10.0.12.1/32;
            }
            family mpls;
        }
    }
    fe-0/1/0 {
        unit 0 {
            family inet {
                address 10.0.12.13/30;
            }
            family mpls;
        }
    }
    at-0/2/1 {
        atm-options {
            pic-type atm2;
            vpi 0;
        }
        unit 0 {
            bandwidth 50m;
```

```
            vci 0.128;
            family inet {
                address 10.0.12.5/32 {
                    destination 10.0.12.6;
                }
            }
            family mpls;
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.70.143/21;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.1.1/32;
            }
        }
    }
}
routing-options {
    static {
        [...Output truncated...]
    }
    router-id 192.168.1.1;
    autonomous-system 65432;
}
protocols {
    rsvp {
        interface fe-0/1/0.0 {
            link-protection {
                bandwidth 50m;
                max-bypasses 2;
            }
        }
        interface fe-0/1/2.0;
        interface so-0/0/0.0;
        interface at-0/2/1.0;
        interface fxp0.0 {
            disable;
        }
    }
    mpls {
        label-switched-path lsp1 {
            from 192.168.1.1;
            to 192.168.4.1;
            bandwidth 10m;
            link-protection;
            primary path1;
        }
        label-switched-path lsp2 {
            from 192.168.1.1;
            to 192.168.4.1;
            bandwidth 20m;
            link-protection;
            primary path1;
        }
        label-switched-path lsp3 {
            from 192.168.1.1;
```

```
            to 192.168.9.1;
            bandwidth 30m;
            link-protection;
            primary path1;
        }
        label-switched-path lsp4 {
            from 192.168.1.1;
            to 192.168.9.1;
            bandwidth 40m;
            link-protection;
            primary path1;
        }
        path path1 {
            10.0.12.14 strict;
        }
        interface fe-0/1/0.0;
        interface so-0/0/0.0;
        interface at-0/2/1.0;
        interface fxp0.0 {
            disable;
        }
    }
    bgp {
        export send-statics;
        group internal {
            type internal;
            local-address 192.168.1.1;
            neighbor 192.168.2.1;
            neighbor 192.168.4.1;
            neighbor 192.168.9.1;
        }
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface fe-0/1/0.0;
            interface at-0/2/1.0;
            interface so-0/0/0.0;
            interface lo0.0 {
                passive;
            }
        }
    }
}
policy-options {
    policy-statement send-statics {
        term statics {
            from {
                route-filter 100.100.1.0/24 exact;
            }
            then accept;
        }
    }
}
```

Sample Output 2 The following sample output is for transit router R2:

```
user@R2> show configuration | no-more
interfaces {
    so-0/0/0 {
        unit 0 {
            family inet {
                address 10.0.12.2/30;
            }
            family mpls;
        }
    }
    so-0/0/1 {
        unit 0 {
            family inet {
                address 10.0.24.1/30;
            }
            family mpls;
        }
    }
    fe-0/1/0 {
        unit 0 {
            family inet {
                address 10.0.12.14/30;
            }
            family mpls;
        }
    }
    at-0/2/1 {
        atm-options {
            pic-type atm2;
            vpi 0;
        }
        unit 0 {
            vci 0.128;
            family inet {
                address 10.0.12.6/32 {
                    destination 10.0.12.5;
                }
            }
            family mpls;
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.70.144/21;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.2.1/32;
            }
        }
    }
}
routing-options {
    static {
        [...Output truncated...]
    }
    router-id 192.168.2.1;
```

```
    autonomous-system 65432;
}
protocols {
    rsvp {
        interface so-0/0/1.0;
        interface fe-0/1/0.0;
        interface so-0/0/0.0;
        interface at-0/2/1.0;
        interface fxp0.0;
    }
    mpls {
        interface fe-0/1/0.0;
        interface so-0/0/1.0;
        interface so-0/0/0.0;
        interface at-0/2/1.0;
        interface fxp0.0;
    }
    bgp {
        group internal {
            type internal;
            local-address 192.168.2.1;
            neighbor 192.168.1.1;
            neighbor 192.168.4.1;
            neighbor 192.168.9.1;
        }
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface fe-0/1/0.0;
            interface so-0/0/1.0;
            interface at-0/2/1.0;
            interface so-0/0/0.0;
            interface lo0.0;
        }
    }
}
```

Sample Output 3 The following sample output is for transit/egress router R4:

```
user@R4> show configuration | no-more
[...Output truncated...]
interfaces {
    so-0/0/1 {
        unit 0 {
            family inet {
                address 10.0.24.2/30;
            }
            family mpls;
        }
    }
    so-0/0/3 {
        unit 0 {
            family inet {
                address 10.0.49.1/30;
            }
            family mpls;
        }
    }
```

```
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.70.146/21;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.4.1/32;
            }
        }
    }
}
routing-options {
    static {
        [...Output truncated...]
    }
    router-id 192.168.4.1;
    autonomous-system 65432;
}
protocols {
    rsvp {
        interface so-0/0/1.0;
        interface so-0/0/3.0;
        interface fxp0.0 {
            disable;
        }
    }
    mpls {
        label-switched-path r4-r1 {
            to 192.168.1.1;
        }
        interface so-0/0/1.0;
        interface so-0/0/3.0;
        interface fxp0.0 {
            disable;
        }
    }
    bgp {
        group internal {
            type internal;
            local-address 192.168.4.1;
            neighbor 192.168.1.1;
            neighbor 192.168.2.1;
            neighbor 192.168.9.1;
        }
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface so-0/0/1.0;
            interface so-0/0/3.0;
            interface lo0.0 {
                passive;
            }
        }
    }
}
```

Sample Output 4 The following sample output is for egress router R9:

```
user@R9> show configuration | no-more
[...Output truncated...]
interfaces {
    so-0/0/3 {
        unit 0 {
            family inet {
                address 10.0.49.2/30;
            }
            family mpls;
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.69.206/21;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.9.1/32;
            }
        }
    }
}
routing-options {
    static {
        [...Output truncated...]
    }
    router-id 192.168.9.1;
    autonomous-system 65432;
}
protocols {
    rsvp {
        interface so-0/0/3.0;
        interface fxp0.0 {
            disable;
        }
    }
    mpls {
        label-switched-path r9-r1 {
            to 192.168.1.1;
        }
        interface so-0/0/3.0;
        interface fxp0.0 {
            disable;
        }
    }
    bgp {
        group internal {
            type internal;
            local-address 192.168.9.1;
            neighbor 192.168.1.1;
            neighbor 192.168.2.1;
            neighbor 192.168.4.1;
        }
    }
```

```
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface so-0/0/3.0;
            interface lo0.0 {
                passive;
            }
        }
    }
}
```