Chapter 9

Verifying the IP and IGP Layers

This chapter describes how to check the Internet Protocol (IP) and interior gateway protocol (IGP) layers of the layered Multiprotocol Label Switching (MPLS) model. (See Table 15.)

Table 15: Checklist for Verifying the IP and IGP Layers

Verifying the IP and IGP Layer Tasks	Command or Action
Verifying the IP and IGP Layers on page 115	
Verifying the IP Layer on page 117	
1. Verify the LSP on page 118	show mpls lsp extensive
2. Verify IP Addressing on page 119	show interfaces terse
3. Verify Neighbors or Adjacencies at the IP Layer on page 120	show ospf neighbor extensive show isis adjacency extensive
4. Take Appropriate Action on page 123	The following sequence of commands addresses the specific problem described in this section: [edit interfaces so-0/0/2] show rename unit 0 family inet address 10.1.13.2/30 to address 10.1.13.1/30 show commit
5. Verify the LSP Again on page 124	show mpls lsp extensive
Verifying the OSPF Protocol on page 128	
1. Verify the LSP on page 129	show mpls lsp extensive
2. Verify OSPF Interfaces on page 131	show ospf interface
3. Verify OSPF Neighbors on page 133	show ospf neighbor
4. Verify the OSPF Protocol Configuration on page 133	show configuration protocols ospf
5. Take Appropriate Action on page 134	The following sequence of commands addresses the specific problem described in this section: [edit] edit protocols ospf area 0.0.0.0 [edit protocols ospf area 0.0.0.0] set interface lo0 set interface lo0 passive up [edit protocols ospf] set traffic-engineering show commit

Verifying the IP and IGP Layer Tasks	Command or Action
6. Verify the LSP Again on page 136	show mpls lsp extensive
Verifying the IS-IS Protocol on page 139	
1. Verify the LSP on page 140	show mpls lsp extensive
2. Verify IS-IS Adjacencies and Interfaces on page 141	show isis adjacency show isis interface
3. Verify the IS-IS Configuration on page 142	show configuration protocols isis
4. Take Appropriate Action on page 143	The following sequence of commands addresses the specific problem described in this section: edit [edit] edit protocols isis [edit protocols isis] show delete level 2 set level 1 disable show commit run show isis adjacency
5. Verify the LSP Again on page 144	show mpls lsp extensive

Verifying the IP and IGP Layers

Purpose After you have configured the label-switched path (LSP), issued the show mpls lsp extensive command, and determined that there is an error, you might find that the error is not in the physical or data link layers. Continue investigating the problem at the IP and IGP layers of the network.

Figure 14 illustrates the IP and IGP layers of the layered MPLS model.

Figure 14: IP and IGP Layers

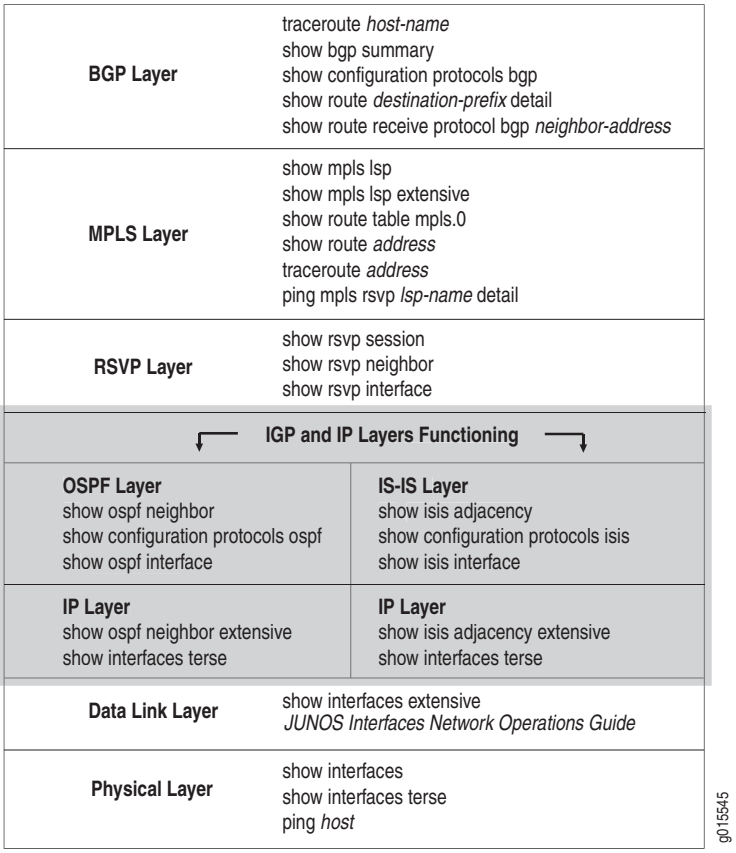

At the IP and IGP layers, you must check the following:

- Interfaces have correct IP addressing, and the IGP neighbors or adjacencies are established.
- Open Shortest Path First (OSPF) or Intermediate System-to-Intermediate System (IS-IS) protocols are configured and running correctly.
  - If the OSPF protocol is configured, check the IP layer first, then the OSPF configuration, making sure that the protocol, interfaces, and traffic engineering are configured correctly.

- If the IS-IS protocol is configured, it doesn't matter whether you check IS-IS or IP first because both protocols are independent of each other. Verify that IS-IS adjacencies are up, and that the interfaces and IS-IS protocol are configured correctly.

NOTE: The IS-IS protocol has traffic engineering enabled by default.

If the network is not functioning at the IP or IGP layers, the LSP does not work as configured.

Figure 15 illustrates the MPLS network used in this chapter.

Figure 15: MPLS Network Broken at the IP and IGP Layers

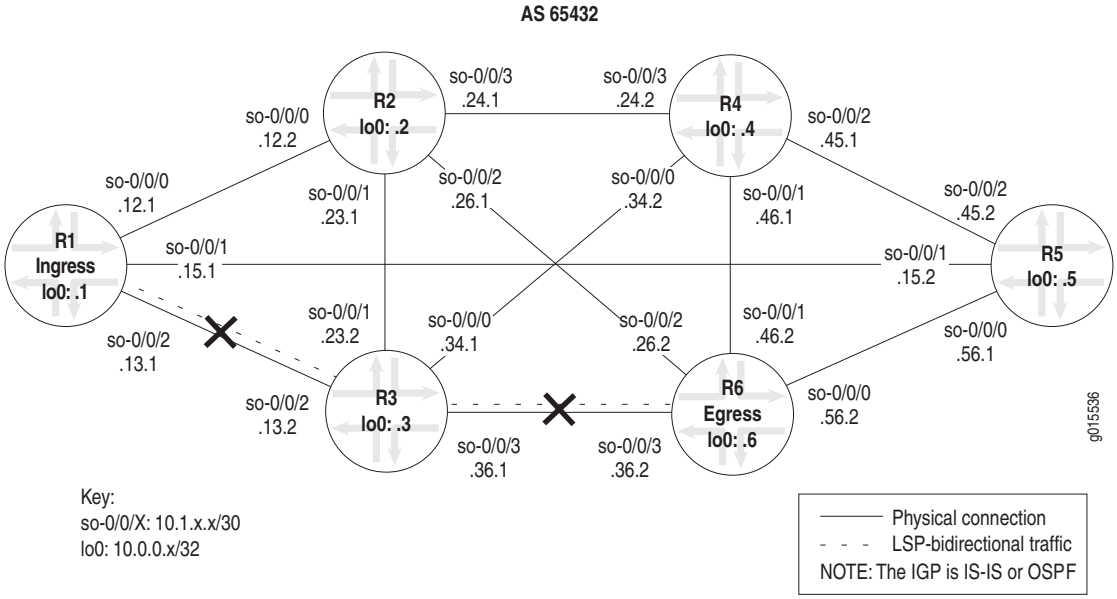

The network shown in Figure 15 is a fully meshed configuration where every directly connected interface can receive and send packets to every other similar interface. The LSP in this network is configured to run from ingress router R1, through transit router R3, to egress router R6. In addition, a reverse LSP is configured to run from R6, through R3, to R1, creating bidirectional traffic. The crosses in Figure 15 indicate where the LSP is not working because of the following problems at the IP and IGP layer:

- An IP address is configured incorrectly on the ingress router (R1).
- The OSPF protocol is configured with a router ID (RID) but without the loopback (lo0) interface, and traffic engineering is missing from the transit router (R3).
- Levels in the IS-IS network are mismatched.

Steps To Take To check the IP and IGP layers, follow these steps:

1. Verifying the IP Layer on page 117
2. Verifying the OSPF Protocol on page 128
3. Verifying the IS-IS Protocol on page 139

Verifying the IP Layer

Purpose You can check the IP layer before or after you check the IGP layer, depending on whether you have OSPF or IS-IS configured as the IGP. If your MPLS network is configured with OSPF as the IGP, you must first verify the IP layer, checking that the interfaces have correct IP addressing and that the OSPF neighbors are established before you check the OSPF layer.

If you have IS-IS configured as the IGP in your MPLS network, you can verify either the IP layer or the IS-IS protocol layer first. The order in which you check the IP or IS-IS layer does not affect the results.

Figure 16: MPLS Network Broken at the IP Layer

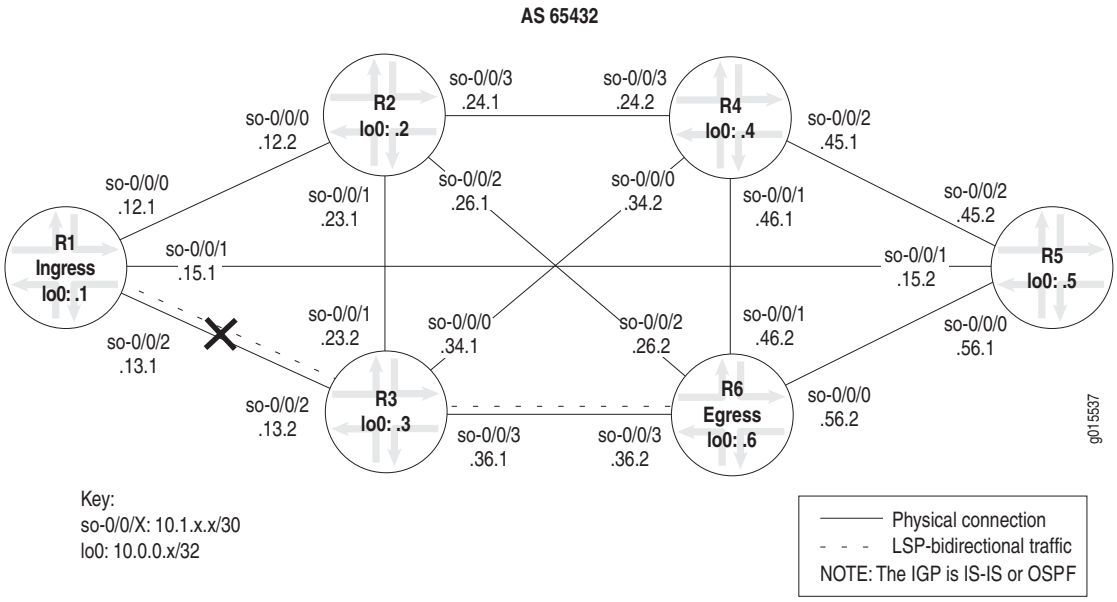

The cross in Figure 16 indicates where the LSP is broken because of the incorrect configuration of an IP address on ingress router R1.

Steps To Take To check the IP layer, follow these steps:

1. Verify the LSP on page 118
2. Verify IP Addressing on page 119
3. Verify Neighbors or Adjacencies at the IP Layer on page 120

4. Take Appropriate Action on page 123

5. Verify the LSP Again on page 136

Step 1: Verify the LSP

Purpose Typically, you use the show mpls lsp extensive command to verify the LSP. However for quick verification of the LSP state, use the show mpls lsp command. If the LSP is down, use the extensive option (show mpls lsp extensive) as a follow-up. If your network has numerous LSPs, you might consider specifying the name of the LSP, using the name option (show mpls lsp name *name* or show mpls lsp name *name* extensive).

Action To determine whether the LSP is up, enter the following command from the ingress router:

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

Sample Output 1

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

10.0.0.6
  From: 10.0.0.1, State: Dn, ActiveRoute: 0, LSPname: R1-to-R6
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary                   State: Dn
    Will be enqueued for recomputation in 25 second(s).
   44 Oct 15 16:56:11 CSPF failed: no route toward 10.0.0.6[2685 times]
   43 Oct 14 19:07:09 Clear Call
   42 Oct 14 19:06:56 Deselected as active
   41 Oct 14 19:06:56 10.1.12.1: MPLS label allocation failure
   40 Oct 14 19:06:56 Down
   39 Oct 14 18:43:43 Selected as active path
   38 Oct 14 18:43:43 Record Route:  10.1.13.2 10.1.36.2
   37 Oct 14 18:43:43 Up
   [...Output truncated...]
  Created: Thu Oct 14 16:04:33 2004
Total 1 displayed, Up 0, Down 1

Egress LSP: 0 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 from ingress router R1 shows that an MPLS label allocation failure occurred and the Constrained Shortest Path First (CSPF) algorithm failed, resulting in no route to destination 10.0.0.6 on R6.

Step 2: Verify IP Addressing

Purpose When you investigate the IP layer, you verify that interfaces have correct IP addressing, and that OSPF neighbors or IS-IS adjacencies are established. In this example, an IP address is configured incorrectly on the ingress router (R1).

Action To verify IP addressing, enter the following command from the ingress, transit, and egress routers:

```
user@host> show interfaces terse
```

Sample Output

```
user@R1> show interfaces terse
Interface               Admin Link Proto Local                 Remote
so-0/0/0                up    up
so-0/0/0.0              up    up   inet  10.1.12.1/30
                                  iso
                                  mpls
so-0/0/1                up    up
so-0/0/1.0              up    up   inet  10.1.15.1/30
                                  iso
                                  mpls
so-0/0/2                up    up
so-0/0/2.0              up    up   inet  10.1.13.2 <<< Incorrect IP address
                                  iso
                                  mpls
lo0                     up    up
lo0.0                   up    up   inet  10.0.0.1
                                  iso   49.0004.1000.0000.0001.00

user@R3> show interfaces terse
Interface               Admin Link Proto Local                 Remote
so-0/0/0                up    up
so-0/0/0.0              up    up   inet  10.1.34.1/30
                                  iso
                                  mpls
so-0/0/1                up    up
so-0/0/1.0              up    up   inet  10.1.23.2/30
                                  iso
                                  mpls
so-0/0/2                up    up
so-0/0/2.0              up    up   inet  10.1.13.2/30 <<< Identical to R1
                                  iso
                                  mpls
so-0/0/3                up    up
so-0/0/3.0              up    up   inet  10.1.36.1/30
                                  iso
                                  mpls
lo0                     up    up
lo0.0                   up    up   inet  10.0.0.3
                                  iso   49.0004.1000.0000.0003.00
```

```
user@R6> show interfaces terse
Interface               Admin Link Proto Local                 Remote
so-0/0/0                up    up
so-0/0/0.0              up    up   inet  10.1.56.2/30
                                  iso
                                  mpls
so-0/0/1                up    up
so-0/0/1.0              up    up   inet  10.1.46.2/30
                                  iso
                                  mpls
so-0/0/2                up    up
so-0/0/2.0              up    up   inet  10.1.26.2/30
                                  iso
                                  mpls
so-0/0/3                up    up
so-0/0/3.0              up    up   inet  10.1.36.2/30
                                  iso
                                  mpls
lo0.0                   up    up   inet  10.0.0.6
                                  iso   49.0004.1000.0000.0006.00
```

What It Means The sample output shows that the IP addresses for interface so-0/0/2.0 on R1 and interface so-0/0/2.0 on R3 are identical. Interface IP addresses within a network must be unique for the interface to be identified correctly.

Step 3: Verify Neighbors or Adjacencies at the IP Layer

Action To verify neighbors (OSPF) or adjacencies (IS-IS), enter the following commands from the ingress, transit, and egress routers:

```
user@host> show ospf neighbor extensive
user@host> show isis adjacency extensive
```

Sample Output 1

```
user@R1> show ospf neighbor extensive
Address         Interface             State      ID              Pri  Dead
10.1.12.2        so-0/0/0.0             Full      10.0.0.2          128   34
  area 0.0.0.0, opt 0x42, DR 0.0.0.0, BDR 0.0.0.0
  Up 1d 04:45:20, adjacent 1d 04:45:20
10.1.15.2        so-0/0/1.0             Full      10.0.0.5          128   35
  area 0.0.0.0, opt 0x42, DR 0.0.0.0, BDR 0.0.0.0
  Up 1d 04:45:20, adjacent 1d 04:45:10 <<< no adjacency with R3 so-0/0/2

user@R3> show ospf neighbor extensive
Address         Interface             State      ID              Pri  Dead
10.1.23.1        so-0/0/1.0             Full      10.0.0.2          128   35
  area 0.0.0.0, opt 0x42, DR 0.0.0.0, BDR 0.0.0.0
  Up 1w2d 04:54:30, adjacent 1w2d 04:54:21
10.1.36.2        so-0/0/3.0             Full      10.0.0.6          128   39
  area 0.0.0.0, opt 0x42, DR 0.0.0.0, BDR 0.0.0.0
  Up 1w2d 04:54:30, adjacent 1w2d 04:54:30 <<< no adjacency with R1 so-0/0/2
```

```
user@R6> show ospf neighbor extensive
Address         Interface             State      ID               Pri  Dead
10.1.56.1        so-0/0/0.0             Full      10.0.0.5         128   39
  area 0.0.0.0, opt 0x42, DR 0.0.0.0, BDR 0.0.0.0
  Up 1d 02:59:35, adjacent 1d 02:59:35
10.1.26.1        so-0/0/2.0             Full      10.0.0.2         128   36
  area 0.0.0.0, opt 0x42, DR 0.0.0.0, BDR 0.0.0.0
  Up 1w2d 04:57:30, adjacent 1w2d 04:57:30
10.1.36.1        so-0/0/3.0             Full      10.0.0.3         128   36
  area 0.0.0.0, opt 0x42, DR 0.0.0.0, BDR 0.0.0.0
  Up 1w2d 04:56:11, adjacent 1w2d 04:56:11
```

Sample Output 2

```
user@R1> show isis adjacency extensive
R2
  Interface: so-0/0/0.0, Level: 2, State: Up, Expires in 23 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 05:57:16 ago
  Circuit type: 2, Speaks: IP, IPv6
  Topologies: Unicast
  Restart capable: Yes
  IP addresses: 10.1.12.2
  Transition log:
  When                  State        Reason
  Fri Oct 15 14:58:35   Up           Seenself

R5
  Interface: so-0/0/1.0, Level: 2, State: Up, Expires in 26 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 05:56:52 ago
  Circuit type: 2, Speaks: IP, IPv6
  Topologies: Unicast
  Restart capable: Yes
  IP addresses: 10.1.15.2
  Transition log:
  When                  State        Reason
  Fri Oct 15 14:59:00   Up           Seenself

R3
  Interface: so-0/0/2.0, Level: 2, State: Up, Expires in 26 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 05:56:51 ago
  Circuit type: 2, Speaks: IP, IPv6
  Topologies: Unicast
  Restart capable: Yes
  IP addresses: 10.1.13.2
  Transition log:
  When                  State        Reason
  Fri Oct 15 14:59:01   Up           Seenself

user@R3> show isis adjacency extensive
R4
  Interface: so-0/0/0.0, Level: 2, State: Up, Expires in 25 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 1w1d 00:22:51 ago
  Circuit type: 2, Speaks: IP, IPv6
  Topologies: Unicast
  Restart capable: Yes
  IP addresses: 10.1.34.2
  Transition log:
  When                  State        Reason
  Thu Oct 28 15:13:12   Up           Seenself
```

```
R2
  Interface: so-0/0/1.0, Level: 2, State: Up, Expires in 25 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 2w2d 18:02:48 ago
  Circuit type: 2, Speaks: IP, IPv6
  Topologies: Unicast
  Restart capable: Yes
  IP addresses: 10.1.23.1
  Transition log:
  When                  State        Reason
  Tue Oct 19 21:33:15   Up           Seenself

R1
  Interface: so-0/0/2.0, Level: 2, State: Up, Expires in 22 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 2w2d 17:24:06 ago
  Circuit type: 2, Speaks: IP, IPv6
  Topologies: Unicast
  Restart capable: Yes
  IP addresses: 10.1.13.1
  Transition log:
  When                  State        Reason
  Tue Oct 19 22:11:57   Up           Seenself

R6
  Interface: so-0/0/3.0, Level: 2, State: Up, Expires in 21 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 2w1d 00:07:00 ago
  Circuit type: 2, Speaks: IP, IPv6
  Topologies: Unicast
  Restart capable: Yes
  IP addresses: 10.1.36.2
  Transition log:
  When                  State        Reason
  Thu Oct 21 15:29:03   Up           Seenself

user@R6> show isis adjacency extensive
R5
  Interface: so-0/0/0.0, Level: 2, State: Up, Expires in 23 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 1w2d 01:10:03 ago
  Circuit type: 2, Speaks: IP, IPv6
  Topologies: Unicast
  Restart capable: Yes
  IP addresses: 10.1.56.1
  Transition log:
  When                  State        Reason
  Wed Oct 27 14:35:32   Up           Seenself

R4
  Interface: so-0/0/1.0, Level: 2, State: Up, Expires in 25 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 1w1d 00:26:50 ago
  Circuit type: 2, Speaks: IP, IPv6
  Topologies: Unicast
  Restart capable: Yes
  IP addresses: 10.1.46.1
  Transition log:
  When                  State        Reason
  Thu Oct 28 15:18:45   Up           Seenself
```

```
R2
  Interface: so-0/0/2.0, Level: 2, State: Up, Expires in 24 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 2w1d 00:11:40 ago
  Circuit type: 2, Speaks: IP, IPv6
  Topologies: Unicast
  Restart capable: Yes
  IP addresses: 10.1.26.1
  Transition log:
  When                  State        Reason
  Thu Oct 21 15:33:55   Up           Seenself

R3
  Interface: so-0/0/3.0, Level: 2, State: Up, Expires in 19 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 2w1d 00:11:40 ago
  Circuit type: 2, Speaks: IP, IPv6
  Topologies: Unicast
  Restart capable: Yes
  IP addresses: 10.1.36.1
  Transition log:
  When                  State        Reason
  Thu Oct 21 15:33:55   Up           Seenself
```

What It Means Sample Output 1 from the ingress, transit, and egress routers shows that R1 and R3 are not established OSPF neighbors. Considering that the two interfaces so-0/0/2.0 (R1 and R3) are configured with identical IP addresses, you would expect this. The OSPF protocol routes IP packets based solely on the destination IP address contained in the IP packet header. Therefore, identical IP addresses in the autonomous system (AS) result in neighbors not establishing.

Sample Output 2 from the ingress, transit, and egress routers shows that R1 and R3 have established an IS-IS adjacency despite the identical IP addresses configured on interfaces so-0/0/2.0 on R1 and R3. The IS-IS protocol behaves differently from the OSPF protocol because it does not rely on IP to establish an adjacency. However, if the LSP is not up, it is still useful to check the IP subnet addressing in case there is a mistake in that layer. Correcting the addressing error might bring the LSP back up.

Step 4: Take Appropriate Action

Purpose Depending on the error you encountered in your investigation, you must take the appropriate action to correct the problem. In this example, the IP address of an interface on transit router R2 is incorrectly configured.

Action To correct the error in this example, enter the following commands:

```
[edit interfaces so-0/0/2]
user@R1# show
user@R1# rename unit 0 family inet address 10.1.13.2/30 to address
10.1.13.1/30
user@R1# show
user@R1# commit
```

Sample Output

```
[edit interfaces so-0/0/2]
user@R1# show
unit 0 {
    family inet {
        address 10.1.13.2/30; <<< Incorrect IP address
    }
    family iso;
    family mpls;
}

[edit interfaces so-0/0/2]
user@R1# rename unit 0 family inet address 10.1.13.2/30 to address 10.1.13.1/30

[edit interfaces so-0/0/2]
user@R1# show
unit 0 {
    family inet {
        address 10.1.13.1/30; <<< Correct IP address.
    }
    family iso;
    family mpls;
}

[edit interfaces so-0/0/2]
user@R1# commit
commit complete
```

What It Means The sample output shows that interface so-0/0/2 on ingress router R1 is now configured with the correct IP address. This correction results in unique subnet IP addresses for all interfaces in the MPLS network in Figure 15 on page 116, and the possibility that the LSP might come up.

Step 5: Verify the LSP Again

Action To verify the LSP again, enter the following command on the ingress, transit, and egress routers:

user@host> **show mpls lsp extensive**

Sample Output 1

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

10.0.0.6
  From: 10.0.0.1, State: Up, ActiveRoute: 1, LSPname: R1-to-R6
  ActivePath:  (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary                    State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
 10.1.13.2 S 10.1.36.2 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node
10=SoftPreempt):
          10.1.13.2 10.1.36.2
   54 Oct 15 21:28:16 Selected as active path
   53 Oct 15 21:28:16 Record Route:  10.1.13.2 10.1.36.2
   52 Oct 15 21:28:16 Up
   51 Oct 15 21:28:16 10.1.15.1: MPLS label allocation failure[2 times]
```

```
   50 Oct 15 21:28:11 CSPF: computation result accepted
   49 Oct 15 21:27:42 10.1.15.1: MPLS label allocation failure
   48 Oct 15 21:27:42 CSPF: computation result accepted
   47 Oct 15 21:27:31 10.1.15.1: MPLS label allocation failure[4 times]
   46 Oct 15 21:27:13 Originate Call
   45 Oct 15 21:27:13 CSPF: computation result accepted
   [...Output truncated...]
  Created: Thu Oct 14 16:04:34 2004
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, LSPstate: Up, ActiveRoute: 0
  LSPname: R6-to-R1, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left:  149, Since: Fri Oct 15 21:28:13 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 13 receiver 39024 protocol 0
  PATH rcvfrom: 10.1.13.2 (so-0/0/2.0) 10 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.1.36.2 10.1.13.2 <self>
Total 1 displayed, Up 1, Down 0

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

Sample Output 2

```
user@R3> show mpls lsp extensive
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

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

Transit LSP: 2 sessions

10.0.0.1
  From: 10.0.0.6, LSPstate: Up, ActiveRoute: 1
  LSPname: R6-to-R1, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: 100336, Label out: 3
  Time left:  156, Since: Fri Oct 15 21:15:47 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 13 receiver 39024 protocol 0
  PATH rcvfrom: 10.1.36.2 (so-0/0/3.0) 11 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 10.1.13.1 (so-0/0/2.0) 11 pkts
  RESV rcvfrom: 10.1.13.1 (so-0/0/2.0) 11 pkts
  Explct route: 10.1.13.1
  Record route: 10.1.36.2 <self> 10.1.13.1
```

```
10.0.0.6
  From: 10.0.0.1, LSPstate: Up, ActiveRoute: 1
  LSPname: R1-to-R6, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: 100352, Label out: 3
  Time left:  159, Since: Fri Oct 15 21:15:50 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 5 receiver 47901 protocol 0
  PATH rcvfrom: 10.1.13.1 (so-0/0/2.0) 11 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 10.1.36.2 (so-0/0/3.0) 11 pkts
  RESV rcvfrom: 10.1.36.2 (so-0/0/3.0) 11 pkts
  Explct route: 10.1.36.2
  Record route: 10.1.13.1 <self> 10.1.36.2
Total 2 displayed, Up 2, Down 0
```

Sample Output 3

```
user@R6> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, State: Up, ActiveRoute: 1, LSPname: R6-to-R1
  ActivePath:  (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary                    State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
 10.1.36.1 S 10.1.13.1 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node
10=SoftPreempt):
          10.1.36.1 10.1.13.1
   187 Oct 15 21:20:05 Selected as active path
   186 Oct 15 21:20:05 Record Route:  10.1.36.1 10.1.13.1
   185 Oct 15 21:20:05 Up
   184 Oct 15 21:20:05 Clear Call
   183 Oct 15 21:20:05 CSPF: computation result accepted
   182 Oct 15 21:20:05 CSPF: link down/deleted
10.1.13.2(R3.00/10.0.0.3)->10.1.13.2(R1.00/10.0.0.1)
   [...Output truncated...]
  Created: Tue Aug 17 12:18:33 2004
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, LSPstate: Up, ActiveRoute: 0
  LSPname: R1-to-R6, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left:  144, Since: Fri Oct 15 21:20:08 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 5 receiver 47901 protocol 0
  PATH rcvfrom: 10.1.36.1 (so-0/0/3.0) 11 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.1.13.1 10.1.36.1 <self>
Total 1 displayed, Up 1, Down 0

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

What It Means Sample Output 1 from ingress router R1 shows that LSP R1-to-R6 has an active route to R6 and the state is up. The output shows that the egress LSP session R6-to-R1 received and sent a recovery label.

Sample Output 2 from transit router R3 shows that there are two transit LSP sessions, one from R1 to R6 and the other from R6 to R1. Both LSPs are up.

Sample Output 3 from egress router R6 shows that the LSP is up and the active route is the primary route. The LSP is now traversing the network along the expected path, from R1 through R3 to R6, and the reverse LSP, from R6 through R3 to R1.

Verifying the OSPF Protocol

Purpose After you have verified that the LSP is down, and the cause is not in the physical, datalink, or IP layer, verify the OSPF configuration. Check the routers in your network to ensure that the interfaces and the OSPF protocol are configured correctly, and that the neighbors are established.

Figure 17: MPLS Network Broken at the OSPF Protocol Layer

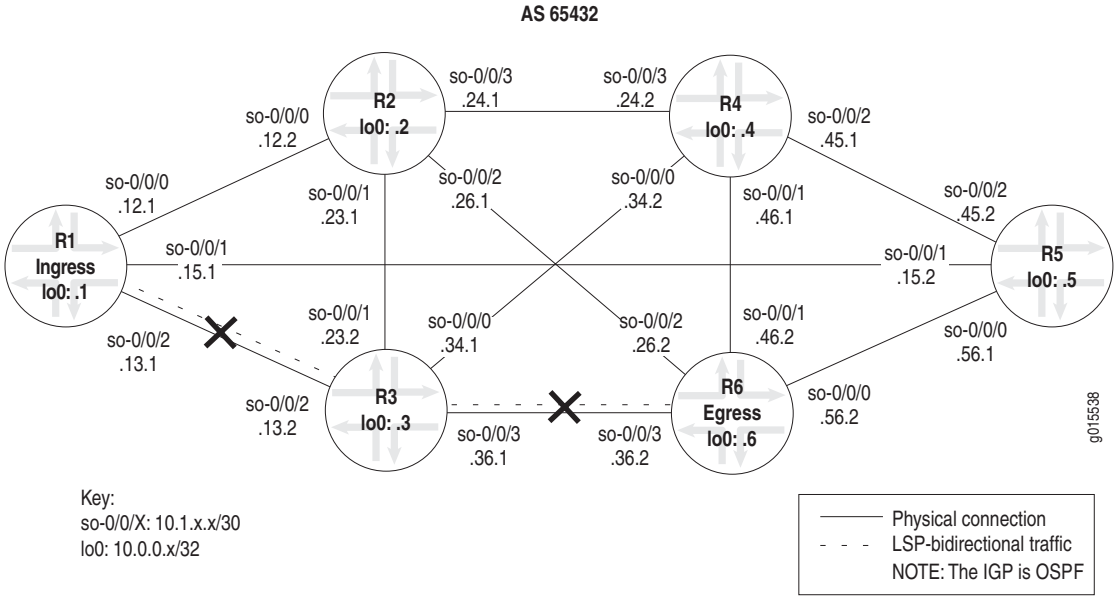

The crosses shown in Figure 17 indicate where the LSP is not working because of problems with the OSPF protocol configuration. The OSPF protocol is configured with a RID but without the loopback (lo0) interface, and traffic engineering is missing from the transit router (R3).

Steps To Take

1. Verify the LSP on page 129
2. Verify OSPF Interfaces on page 131
3. Verify OSPF Neighbors on page 133
4. Verify the OSPF Protocol Configuration on page 133
5. Take Appropriate Action on page 134
6. Verify the LSP Again on page 136

Step 1: Verify the LSP

Action To verify the LSP, enter the following command on the ingress, transit, and egress routers:

user@host> **show mpls lsp extensive**

Sample Output 1

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

10.0.0.6
  From: 10.0.0.1, State: Dn, ActiveRoute: 0, LSPname: R1-to-R6
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary                   State: Dn
   11 Oct 19 18:06:04 No Route toward dest[78 times]
   10 Oct 19 17:08:09 Deselected as active
  Created: Mon Oct 18 21:48:42 2004
Total 1 displayed, Up 0, Down 1

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

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

Sample Output 2

```
user@R3> show mpls lsp extensive
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

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

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

Sample Output 3

```
user@R6> show mpls lsp extensive
Ingress LSP: 1 sessions
To              From            State Rt ActivePath       P     LSPname
10.0.0.1        10.0.0.6        Dn     0 -                      R6-to-R1
Total 1 displayed, Up 0, Down 1

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

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

Sample Output 4

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

10.0.0.6
  From: 10.0.0.1, State: Up, ActiveRoute: 1, LSPname: R1-to-R6
  ActivePath:  (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary                   State: Up
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node
10=SoftPreempt):
```

```
          10.1.13.2 10.1.36.2
    5 Oct 19 10:37:55 Selected as active path
    4 Oct 19 10:37:55 Record Route:  10.1.13.2 10.1.36.2
    3 Oct 19 10:37:55 Up
    2 Oct 19 10:37:10 No Route toward dest[1029 times]
    1 Oct 18 21:48:42 Originate Call
  Created: Mon Oct 18 21:48:42 2004
Total 1 displayed, Up 1, Down 0

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

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

Sample Output 5

```
user@R3> show mpls lsp extensive
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

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

Transit LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, LSPstate: Up, ActiveRoute: 1
  LSPname: R1-to-R6, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: 100368, Label out: 3
  Time left:  154, Since: Tue Oct 19 10:25:24 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 47933 protocol 0
  PATH rcvfrom: 10.1.13.1 (so-0/0/2.0) 209 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 10.1.36.2 (so-0/0/3.0) 209 pkts
  RESV rcvfrom: 10.1.36.2 (so-0/0/3.0) 209 pkts
  Record route: 10.1.13.1 <self> 10.1.36.2
Total 1 displayed, Up 1, Down 0
```

Sample Output 6

```
user@R6> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, State: Dn, ActiveRoute: 0, LSPname: R6-to-R1
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary                   State: Dn
    2 Oct 19 13:01:54 10.1.56.2: MPLS label allocation failure[9 times]
    1 Oct 19 12:57:51 Originate Call
  Created: Tue Oct 19 12:57:51 2004
Total 1 displayed, Up 0, Down 1

Egress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, LSPstate: Up, ActiveRoute: 0
  LSPname: R1-to-R6, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
```

```
  Time left:  148, Since: Tue Oct 19 10:30:03 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 47933 protocol 0
  PATH rcvfrom: 10.1.36.1 (so-0/0/3.0) 206 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.1.13.1 10.1.36.1 <self>
Total 1 displayed, Up 1, Down 0

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

What It Means Sample Outputs 1, 2, and 3 show that the LSP and the reverse LSP are down:

- Sample Output 1 from ingress router R1 shows that LSP R1-to-R6 does not have a route towards the destination (R6).
- Sample Output 2 from transit router R3 shows that there are no LSP sessions.
- Sample Output 3 from egress router R6 also shows that reverse LSP R6-to-R1 is down.

Sample Outputs 4, 5, and 6 show that the LSP is up and the reverse LSP is down:

- Sample Output 4 from ingress router R1 shows that LSP R1-to-R6 is up and there are no egress LSP sessions.
- Sample Output 5 from transit router R3 shows that there is one ingress LSP session (R1-to-R6) and no egress LSP sessions.
- Sample Output 6 from egress router R6 shows that LSP R6-to-R1 is down due to an MPLS label allocation failure.

Step 2: Verify OSPF Interfaces

Purpose After you have verified that the LSP is down, and the cause is not in the physical, data link, or IP layer, check the routers in your network to determine that all relevant OSPF interfaces are configured correctly.

Action To verify OSPF interfaces, enter the following commands from the ingress, transit, and egress routers:

user@host> **show ospf interface**

Sample Output 1

```
user@R1> show ospf interface
Interface      State     Area            DR ID           BDR ID       Nbrs
so-0/0/0.0      PtToPt   0.0.0.0         0.0.0.0         0.0.0.0         1
so-0/0/1.0      PtToPt   0.0.0.0         0.0.0.0         0.0.0.0         1
so-0/0/2.0      PtToPt   0.0.0.0         0.0.0.0         0.0.0.0         1

user@R3> show ospf interface
Interface      State     Area            DR ID           BDR ID       Nbrs
so-0/0/0.0      PtToPt   0.0.0.0         0.0.0.0         0.0.0.0         1
so-0/0/1.0      PtToPt   0.0.0.0         0.0.0.0         0.0.0.0         1
so-0/0/2.0      PtToPt   0.0.0.0         0.0.0.0         0.0.0.0         1
so-0/0/3.0      PtToPt   0.0.0.0         0.0.0.0         0.0.0.0         1
```

```
user@R6> show ospf interface
Interface      State     Area            DR ID           BDR ID       Nbrs
so-0/0/0.0      PtToPt   0.0.0.0         0.0.0.0         0.0.0.0         1
so-0/0/1.0      PtToPt   0.0.0.0         0.0.0.0         0.0.0.0         1
so-0/0/2.0      PtToPt   0.0.0.0         0.0.0.0         0.0.0.0         1
so-0/0/3.0      PtToPt   0.0.0.0         0.0.0.0         0.0.0.0         1
```

Sample Output 2

```
user@R1> show ospf interface
Interface      State     Area            DR ID           BDR ID       Nbrs
lo0.0           DR       0.0.0.0         10.0.0.1        0.0.0.0         0
so-0/0/0.0      PtToPt   0.0.0.0         0.0.0.0         0.0.0.0         1
so-0/0/1.0      PtToPt   0.0.0.0         0.0.0.0         0.0.0.0         1
so-0/0/2.0      PtToPt   0.0.0.0         0.0.0.0         0.0.0.0         1

user@R3> show ospf interface
Interface      State     Area            DR ID           BDR ID       Nbrs
lo0.0           DR       0.0.0.0         10.0.0.3        0.0.0.0         0
so-0/0/0.0      Down     0.0.0.0         0.0.0.0         0.0.0.0         0
so-0/0/1.0      PtToPt   0.0.0.0         0.0.0.0         0.0.0.0         1
so-0/0/2.0      PtToPt   0.0.0.0         0.0.0.0         0.0.0.0         1
so-0/0/3.0      PtToPt   0.0.0.0         0.0.0.0         0.0.0.0         1

user@R6> show ospf interface
Interface      State     Area            DR ID           BDR ID       Nbrs
lo0.0           DR       0.0.0.0         10.0.0.6        0.0.0.0         0
so-0/0/0.0      PtToPt   0.0.0.0         0.0.0.0         0.0.0.0         1
so-0/0/1.0      Down     0.0.0.0         0.0.0.0         0.0.0.0         0
so-0/0/2.0      PtToPt   0.0.0.0         0.0.0.0         0.0.0.0         1
so-0/0/3.0      PtToPt   0.0.0.0         0.0.0.0         0.0.0.0         1
```

What It Means

Sample Output 1 shows that all interfaces on all routers are in the correct area (0.0.0.0), and the loopback (lo0) interface is missing from the list of interfaces on all routers. The missing loopback (lo0) interface is a problem in this configuration.

In an MPLS network configured with OSPF as the IGP, when you manually configure the RID, it is important to explicitly configure the loopback interface at the [edit protocols ospf] hierarchy level. If the RID is not manually configured, OSPF automatically advertises the loopback (lo0) interface. In the configuration of all the routers in this network, the RID is configured manually, therefore, the loopback (lo0) interface must be explicitly configured at the [edit protocols ospf] hierarchy level. In addition, the loopback (lo0) interface is configured with the passive statement to ensure that the protocols are not run over the loopback (lo0) interface and it is correctly advertised throughout the network.

Sample Output 2 shows that all the relevant interfaces on the ingress, transit, and egress routers, including the loopback (lo0) interface, are in the correct area (0.0.0.0). Because the configuration of the interfaces is correct, further investigation is required to determine the reason for the LSP problem.

Step 3: Verify OSPF Neighbors

Purpose After you have checked OSPF interfaces, check your network topology to determine that all relevant neighbors are established.

Action To verify OSPF neighbors, enter the following commands from the ingress, transit, and egress routers:

user@host> **show ospf neighbor**

Sample Output

```
user@R1> show ospf neighbor
  Address          Interface              State      ID               Pri  Dead
10.1.12.2          so-0/0/0.0             Full       10.0.0.2         128  39
10.1.15.2          so-0/0/1.0             Full       10.0.0.5         128  39
10.1.13.2          so-0/0/2.0             Full       10.0.0.3         128  33

user@R3> show ospf neighbor
Address          Interface              State      ID               Pri  Dead
10.1.34.2          so-0/0/0.0               Full       10.0.0.4         128  33
10.1.23.1          so-0/0/1.0               Full       10.0.0.2         128  33
10.1.13.1          so-0/0/2.0               Full       10.0.0.1         128  33
10.1.36.2          so-0/0/3.0               Full       10.0.0.6         128  33

user@R6> show ospf neighbor
Address          Interface              State      ID               Pri  Dead
10.1.56.1          so-0/0/0.0               Full       10.0.0.5         128  30
10.1.46.1          so-0/0/1.0               Full       10.0.0.4         128  38
10.1.26.1          so-0/0/2.0               Full       10.0.0.2         128  34
10.1.36.1          so-0/0/3.0               Full       10.0.0.3         128  35
```

What It Means The sample output shows that all neighbors are fully adjacent, indicating that each router has exchanged a full copy of its link-state database with the other routers, passed through several neighbor states, and become fully adjacent. These adjacencies are created by router link and network link advertisements.

Step 4: Verify the OSPF Protocol Configuration

Purpose After you have checked interfaces and neighbors, verify the OSPF protocol configuration.

Action To verify the OSPF protocol configuration, enter the following command from the ingress, transit, and egress routers:

user@host> **show configuration protocols ospf**

Sample Output 1

```
user@R1> show configuration protocols ospf
traffic-engineering;
area 0.0.0.0 {
    interface so-0/0/0.0;
    interface so-0/0/1.0;
    interface so-0/0/2.0;     <<< The loopback interface (lo0) is missing
}
```

Sample Output 2

```
user@R3> show configuration protocols ospf
area 0.0.0.0 {      <<< traffic engineering is missing
    interface so-0/0/0.0;
    interface so-0/0/1.0;
    interface so-0/0/2.0;
    interface so-0/0/3.0;    <<< The loopback interface (lo0) is missing
}
```

Sample Output 3

```
user@R6> show configuration protocols ospf
traffic-engineering;
area 0.0.0.0 {
    interface so-0/0/0.0;
    interface so-0/0/1.0;
    interface so-0/0/2.0;
    interface so-0/0/3.0;    <<< The loopback interface (lo0) is missing
}
```

What It Means All three sample outputs show that the loopback interface is not included on any of the routers. Including the loopback (lo0) interface is important when you have the RID manually configured.

In addition, Sample Output 2 from transit router R3 shows that traffic engineering is not configured. Traffic engineering must be manually enabled when you configure OSPF for an MPLS network.

Because the loopback interface and traffic engineering are missing from the OSPF protocol configuration, the LSP does not work as expected.

Step 5: Take Appropriate Action

Purpose Depending on the error you encountered in your investigation, you must take the appropriate action to correct the problem. In this example, the loopback (lo0) interface is missing from all routers, and traffic engineering is missing from the transit router (R3).

Action To correct the errors in this example, follow these steps:

1. Include the loopback (lo0) interface on all routers that have the RID manually configured. Enter the following configuration mode commands:

```
[edit]
user@R3# edit protocols ospf area 0.0.0.0
[edit protocols ospf area 0.0.0.0]
user@R3# set interface lo0
user@R3# set interface lo0 passive
```

2. Move up one level of the configuration hierarchy:

```
[edit protocols ospf area 0.0.0.0]
user@R3# up
[edit protocols ospf]
user @R3#
```

3. Include traffic engineering on the transit router (R3). Enter the following configuration mode command:

```
[edit protocols ospf]
user@R3# set traffic-engineering
```

4. On all routers, verify and commit the configuration:

```
user@R3# show
user@R3# commit
```

Sample Output

```
user@R3> edit
Entering configuration mode

[edit]
user@R3# edit protocols ospf area 0.0.0.0

[edit protocols ospf area 0.0.0.0]
user@R3# set interface lo0

[edit protocols ospf area 0.0.0.0]
user@R3# set interface lo0 passive

[edit protocols ospf area 0.0.0.0]
user@R3# up

[edit protocols ospf]
user@R3# set traffic-engineering

[edit protocols ospf]
user@R3# show
traffic-engineering;
area 0.0.0.0 {
    interface so-0/0/0.0;
    interface so-0/0/1.0;
    interface so-0/0/2.0;
    interface lo0.0; {
        passive
}
}

[edit protocols ospf]
user@R3# commit
commit complete
```

What It Means

The sample output shows that the loopback (lo0) interface and traffic engineering are now correctly configured on transit router R3. When traffic engineering is configured, OSPF advertises the traffic engineering capabilities of the links.

In the OSPF configuration, you must manually include the loopback (lo0) interface and set it to passive when you manually configure an RID. Setting the loopback (lo0) interface to passive ensures that protocols are not run over the loopback (lo0) interface and the loopback (lo0) interface is advertised correctly throughout the network.. If you do not manually configure an RID, there is no need to explicitly include the loopback interface because the OSPF protocol automatically includes the loopback (lo0) interface.

For more information about configuring LSPs and MPLS, see the *JUNOS MPLS Applications Configuration Guide*.

Step 6: Verify the LSP Again

Action To verify that the LSP is up and traversing the network as expected, enter the following command from the ingress, egress, and transit routers:

user@host> **show mpls lsp extensive**

Sample Output

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

10.0.0.6
  From: 10.0.0.1, State: Up, ActiveRoute: 1, LSPname: R1-to-R6
  ActivePath:  (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary                    State: Up
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node
10=SoftPreempt):
          10.1.13.2 10.1.36.2
    4 Oct 19 21:22:54 Selected as active path
    3 Oct 19 21:22:53 Record Route:  10.1.13.2 10.1.36.2
    2 Oct 19 21:22:53 Up
    1 Oct 19 21:22:53 Originate Call
  Created: Tue Oct 19 21:22:53 2004
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, LSPstate: Up, ActiveRoute: 0
  LSPname: R6-to-R1, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left:  117, Since: Tue Oct 19 21:17:42 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 2 receiver 39064 protocol 0
  PATH rcvfrom: 10.1.13.2 (so-0/0/2.0) 10 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.1.36.2 10.1.13.2 <self>
Total 1 displayed, Up 1, Down 0

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

user@R3> show mpls lsp extensive
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

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

Transit LSP: 2 sessions
```

```
10.0.0.1
  From: 10.0.0.6, LSPstate: Up, ActiveRoute: 1
  LSPname: R6-to-R1, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: 100416, Label out: 3
  Time left:  139, Since: Tue Oct 19 21:05:11 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 2 receiver 39064 protocol 0
  PATH rcvfrom: 10.1.36.2 (so-0/0/3.0) 11 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 10.1.13.1 (so-0/0/2.0) 11 pkts
  RESV rcvfrom: 10.1.13.1 (so-0/0/2.0) 11 pkts
  Explct route: 10.1.13.1
  Record route: 10.1.36.2 <self> 10.1.13.1

10.0.0.6
  From: 10.0.0.1, LSPstate: Up, ActiveRoute: 1
  LSPname: R1-to-R6, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: 100448, Label out: 3
  Time left:  135, Since: Tue Oct 19 21:10:22 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 47951 protocol 0
  PATH rcvfrom: 10.1.13.1 (so-0/0/2.0) 4 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 10.1.36.2 (so-0/0/3.0) 4 pkts
  RESV rcvfrom: 10.1.36.2 (so-0/0/3.0) 4 pkts
  Record route: 10.1.13.1 <self> 10.1.36.2
Total 2 displayed, Up 2, Down 0

user@R6> run show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, State: Up, ActiveRoute: 1, LSPname: R6-to-R1
  ActivePath:  (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary                    State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 2)
 10.1.36.1 S 10.1.13.1 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node
10=SoftPreempt):
          10.1.36.1 10.1.13.1
   19 Oct 19 21:09:52 Selected as active path
   18 Oct 19 21:09:52 Record Route:  10.1.36.1 10.1.13.1
   17 Oct 19 21:09:52 Up
   16 Oct 19 21:09:52 Originate Call
   15 Oct 19 21:09:52 CSPF: computation result accepted
  Created: Tue Oct 19 18:30:09 2004
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, LSPstate: Up, ActiveRoute: 0
  LSPname: R1-to-R6, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left:  120, Since: Tue Oct 19 21:15:03 2004
```

```
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 47951 protocol 0
  PATH rcvfrom: 10.1.36.1 (so-0/0/3.0) 4 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.1.13.1 10.1.36.1 <self>
Total 1 displayed, Up 1, Down 0

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

What It Means The sample output from ingress router R1 and egress router R6 shows that the LSP is now traversing the network along the expected path, from R1 through R3 to R6, and the reverse LSP, from R6 through R3 to R1. In addition, the sample output from transit router R3 shows that there are two transit LSP sessions, one from R1 to R6, and the other from R6 to R1.

Verifying the IS-IS Protocol

Purpose If your MPLS network is configured with IS-IS as the IGP, and the output of the show mpls lsp extensive command shows that there is a problem, check the IP and IS-IS layers. Because IS-IS and IP are independent of each other, you can check either layer first. For more information on checking the IP layer, see "Verifying the IP Layer" on page 117.

After you have checked the IP layer and determined that there is still a problem, check the IS-IS layer, verify that IS-IS adjacencies are up, and make sure that the interfaces and IS-IS protocol are configured correctly.

Figure 18: MPLS Network Broken at the IS-IS Protocol Layer

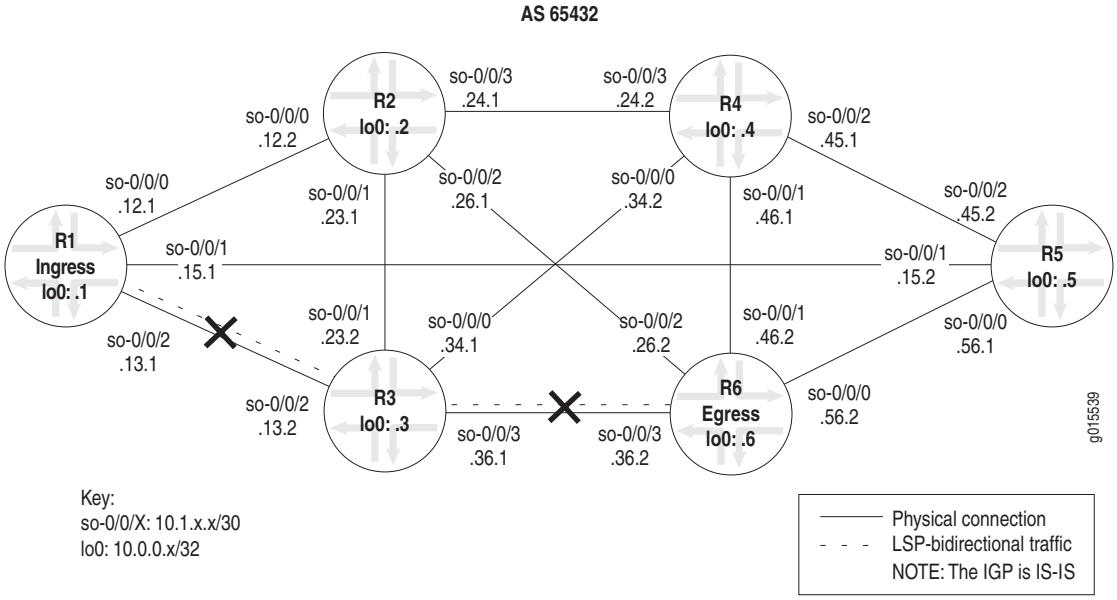

The crosses in Figure 18 indicate where the LSP is not working because IS-IS levels are mismatched.

Steps To Take To check the IS-IS protocol, follow these steps:

1. Verify the LSP on page 140
2. Verify IS-IS Adjacencies and Interfaces on page 141
3. Verify the IS-IS Configuration on page 142
4. Take Appropriate Action on page 143
5. Verify the LSP Again on page 144

Step 1: Verify the LSP

Action To verify the LSP, enter the following command on the ingress, transit, and egress routers:

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

Sample Output 1

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

10.0.0.6
  From: 10.0.0.1, State: Dn, ActiveRoute: 0, LSPname: R1-to-R6
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary                   State: Dn
   24 Oct 21 13:48:01 No Route toward dest[3 times]
   23 Oct 21 13:47:44 Deselected as active
   22 Oct 21 13:47:43 No Route toward dest[2 times]
   21 Oct 21 13:47:43 ResvTear received
   20 Oct 21 13:47:43 Down
   19 Oct 21 13:47:43 10.1.13.2: No Route toward dest[2 times]
   18 Oct 21 13:47:38 Record Route:  10.1.13.2 10.1.36.2
   [...Output truncated...]
  Created: Tue Oct 19 21:22:53 2004
Total 1 displayed, Up 0, Down 1

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

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

Sample Output 2

```
user@R3> show mpls lsp extensive
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

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

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

Sample Output 3

```
user@R6> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, State: Dn, ActiveRoute: 0, LSPname: R6-to-R1
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary                   State: Dn
    Will be enqueued for recomputation in 3 second(s).
   13 Oct 21 14:23:33 CSPF failed: no route toward 10.0.0.1[90 times]
   12 Oct 21 13:39:56 Deselected as active
   11 Oct 21 13:39:56 CSPF: could not determine self
   [...Output truncated...]
  Created: Tue Oct 19 22:28:30 2004
Total 1 displayed, Up 0, Down 1
```

```
Egress LSP: 0 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 shows that LSP R1-to-R6 and the reverse LSP R6-to-R1 are down, and there are no LSP sessions on transit router R3.

Step 2: Verify IS-IS Adjacencies and Interfaces

Purpose When you check the IS-IS layer, you verify that IS-IS adjacencies are up, and that the IS-IS interfaces are included at the protocol level.

Action To verify the functioning of adjacent interfaces, enter the following commands from the relevant routers:

```
user@host> show isis adjacency
user@host> show isis interface
```

Sample Output 1

```
user@R1> show isis adjacency
Interface             System         L State       Hold (secs) SNPA
so-0/0/0.0            R2             2 Up                  20
so-0/0/1.0            R5             2 Up                  23
so-0/0/2.0            R3             2 Up                  26

user@R3> show isis adjacency
Interface             System         L State       Hold (secs) SNPA
so-0/0/0.0            R4             2 Up                  23
so-0/0/1.0            R2             2 Up                  21
so-0/0/2.0            R1             2 Up                  19
so-0/0/3.0            R6             2 Down                 0

user@R6> show isis adjacency

user@R6>            <<< No IS-IS adjacencies are established
```

Sample Output 2

```
user@R1> show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR        Level 2 DR        L1/L2 Metric
lo0.0                 0   0x1 Passive           Passive                 0/0
so-0/0/0.0            2   0x1 Disabled          Point to Point        10/10
so-0/0/1.0            2   0x1 Disabled          Point to Point        10/10
so-0/0/2.0            2   0x1 Disabled          Point to Point        10/10

user@R3> show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR        Level 2 DR        L1/L2 Metric
lo0.0                 0   0x1 Passive           Passive                 0/0
so-0/0/0.0            2   0x1 Disabled          Point to Point        10/10
so-0/0/1.0            2   0x1 Disabled          Point to Point        10/10
so-0/0/2.0            2   0x1 Disabled          Point to Point        10/10
so-0/0/3.0            2   0x1 Disabled          Point to Point        10/10
```

```
user@R6> show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR        Level 2 DR        L1/L2 Metric
lo0.0                 0   0x1 Passive           Passive                0/0
so-0/0/0.0            1   0x1 Point to Point    Disabled              10/10
so-0/0/1.0            1   0x1 Down              Disabled              10/10
so-0/0/2.0            1   0x1 Point to Point    Disabled              10/10
so-0/0/3.0            1   0x1 Point to Point    Disabled              10/10
```

What It Means Sample Output 1 shows that ingress router R1 has established adjacencies with the relevant routers. Transit router R3 does not have an adjacency with egress router R6, and egress router R6 has no adjacencies established in the network shown in Figure 15 on page 116, indicating that the problem might be at the IS-IS protocol level.

Sample Output 2 shows that R1 and R2 are Level 2 routers, in contrast to R6 which is a Level 1 router. When a router is configured explicitly as a Level 1 or Level 2 router, it does not communicate with routers configured at a different level. Level 1 routers communicate with other Level 1 routers within their area, while Level 2 routers communicate with other Level 2 routers, and towards other autonomous systems. Because all the routers in this network are configured for Level 2, they cannot form an adjacency with R6, which is incorrectly configured as a Level 1 router.

Step 3: Verify the IS-IS Configuration

Purpose When you have determined that the problem is probably at the IS-IS protocol level, check the IS-IS configuration of the routers in your network.

Action To verify the IS-IS configuration, enter the following command from the relevant routers:

user@host> **show configuration protocols isis**

Sample Output

```
user@R1> show configuration protocols isis
level 1 disable;
interface so-0/0/0.0;
interface so-0/0/1.0;
interface so-0/0/2.0;
interface lo0.0; {
    passive

user@R3> show configuration protocols isis
level 1 disable;
interface all {
    level 2 metric 10;
}
interface fxp0.0 {
    disable;
}
interface lo0.0; {
    passive
```

```
user@R6> show configuration protocols isis
level 2 disable;        <<< Incorrect level disabled
interface all {
    level 2 metric 10;
}
interface fxp0.0 {
    disable;
}
interface lo0.0; {
    passive
```

What It Means The sample output shows that R6 has Level 2 disabled, while R1 and R3 have Level 1 disabled. For IS-IS adjacencies to establish, routers need to be at the same level. Another common configuration error is to omit the loopback (lo0) interface from the configuration at the [edit protocols isis] hierarchy level. IS-IS does not function correctly if the loopback (lo0) interface is not configured at this level. In addition, including the passive statement ensures that protocols are not run over the loopback (lo0) interface and the loopback (lo0) interface is advertised correctly throughout the network.

Step 4: Take Appropriate Action

Action Depending on the error you encountered in your investigation, you must take the appropriate action to correct the problem. In the example below, the routers are configured to function at different levels of the IS-IS protocol.

Action To correct the error in this example, enter the following commands:

```
user@R6> edit
[edit]
user@R6> edit protocols isis
[edit protocols isis]
user@R6# show
user@R6# delete level 2
user@R6# set level 1 disable
user@R6# show
user@R6# commit
user@R6# run show isis adjacency
```

Sample Output

```
user@R6> edit
Entering configuration mode

[edit]
user@R6# edit protocols isis

[edit protocols isis]
user@R6# show
level 2 disable;
interface all {
    level 2 metric 10;
}
interface fxp0.0 {
    disable;
}
interface lo0.0; {
    passive
```

```
[edit protocols isis]
user@R6# delete level 2

[edit protocols isis]
user@R6# set level 1 disable

[edit protocols isis]
user@R6# show
level 1 disable;
interface all {
    level 2 metric 10;
}
interface fxp0.0 {
    disable;
}
interface lo0.0; {
    passive

[edit protocols isis]
user@R6# commit
commit complete

[edit protocols isis]
user@R6# run show isis adjacency
Interface             System         L State       Hold (secs) SNPA
so-0/0/0.0            R5             2 Up                   22
so-0/0/1.0            R4             2 Up                   22
so-0/0/2.0            R2             2 Up                   22
so-0/0/3.0            R3             2 Up                   22
```

What It Means The sample output shows that the configuration error on egress router R6 has been corrected and IS-IS adjacencies are now established.

Step 5: Verify the LSP Again

Action To verify that the LSP is up and traversing the network as expected, enter the following command from the ingress, egress, and transit routers:

user@host> **show mpls lsp extensive**

Sample Output 1

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

10.0.0.6
  From: 10.0.0.1, State: Up, ActiveRoute: 1, LSPname: R1-to-R6
  ActivePath:  (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary                    State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
 10.1.13.2 S 10.1.36.2 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node
10=SoftPreempt):
          10.1.13.2 10.1.36.2
    5 Oct 21 15:52:07 Selected as active path
    4 Oct 21 15:52:07 Record Route:  10.1.13.2 10.1.36.2
    3 Oct 21 15:52:07 Up
```

```
    2 Oct 21 15:52:07 Originate Call
    1 Oct 21 15:52:07 CSPF: computation result accepted
  Created: Thu Oct 21 15:52:06 2004
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, LSPstate: Up, ActiveRoute: 0
  LSPname: R6-to-R1, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left:  142, Since: Thu Oct 21 15:41:59 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 2 receiver 39082 protocol 0
  PATH rcvfrom: 10.1.13.2 (so-0/0/2.0) 17 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.1.36.2 10.1.13.2 <self>
Total 1 displayed, Up 1, Down 0

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

Sample Output 2

```
user@R3> show mpls lsp extensive
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

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

Transit LSP: 2 sessions

10.0.0.1
  From: 10.0.0.6, LSPstate: Up, ActiveRoute: 1
  LSPname: R6-to-R1, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: 100528, Label out: 3
  Time left:  125, Since: Thu Oct 21 15:29:26 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 2 receiver 39082 protocol 0
  PATH rcvfrom: 10.1.36.2 (so-0/0/3.0) 17 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 10.1.13.1 (so-0/0/2.0) 17 pkts
  RESV rcvfrom: 10.1.13.1 (so-0/0/2.0) 17 pkts
  Explct route: 10.1.13.1
  Record route: 10.1.36.2 <self> 10.1.13.1

10.0.0.6
  From: 10.0.0.1, LSPstate: Up, ActiveRoute: 1
  LSPname: R1-to-R6, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: 100544, Label out: 3
  Time left:  147, Since: Thu Oct 21 15:39:33 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 47963 protocol 0
  PATH rcvfrom: 10.1.13.1 (so-0/0/2.0) 4 pkts
  Adspec: received MTU 1500 sent MTU 1500
```

```
  PATH sentto: 10.1.36.2 (so-0/0/3.0) 4 pkts
  RESV rcvfrom: 10.1.36.2 (so-0/0/3.0) 4 pkts
  Explct route: 10.1.36.2
  Record route: 10.1.13.1 <self> 10.1.36.2
Total 2 displayed, Up 2, Down 0
```

Sample Output 3

```
user@R6> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, State: Up, ActiveRoute: 1, LSPname: R6-to-R1
  ActivePath:  (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary                    State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
 10.1.36.1 S 10.1.13.1 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node
10=SoftPreempt):
          10.1.36.1 10.1.13.1
   18 Oct 21 15:34:18 Selected as active path
   17 Oct 21 15:34:17 Record Route:  10.1.36.1 10.1.13.1
   16 Oct 21 15:34:17 Up
   15 Oct 21 15:34:17 Originate Call
   14 Oct 21 15:34:17 CSPF: computation result accepted
   [...Output truncated...]
  Created: Tue Oct 19 22:28:30 2004
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, LSPstate: Up, ActiveRoute: 0
  LSPname: R1-to-R6, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left:  126, Since: Thu Oct 21 15:44:25 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 47963 protocol 0
  PATH rcvfrom: 10.1.36.1 (so-0/0/3.0) 4 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.1.13.1 10.1.36.1 <self>
Total 1 displayed, Up 1, Down 0

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

What It Means

Sample Outputs 1 and 3 from ingress router R1 and egress router R6 show that the LSP is now traversing the network along the expected path, from R1 through R3 to R6, and the reverse LSP, from R6 through R3 to R1. In addition, Sample Output 2 from transit router R3 shows that there are two transit LSP sessions, one from R1 to R6, and the other from R6 to R1.