Chapter 5

Verifying LSP Use

This chapter describes how to verify the availability and valid use of a label-switched path (LSP) in your network. (See Table 10.)

Table 10: Checklist for Verifying LSP Use

Verifying LSP Use Tasks	Command or Action
Verifying LSP Use in Your Network on page 78	
Verifying an LSP on the Ingress Router on page 79	show route table inet.3
Verifying an LSP on a Transit Router on page 80	show route table mpls.0

Verifying LSP Use in Your Network

Purpose When you verify the valid use of an LSP on the ingress and transit routers in your network, you can determine if there is a problem with Multiprotocol Label Switching (MPLS) in your network. Figure 7 describes the example network used in this chapter.

Figure 7: MPLS Topology for Verifying LSP Use

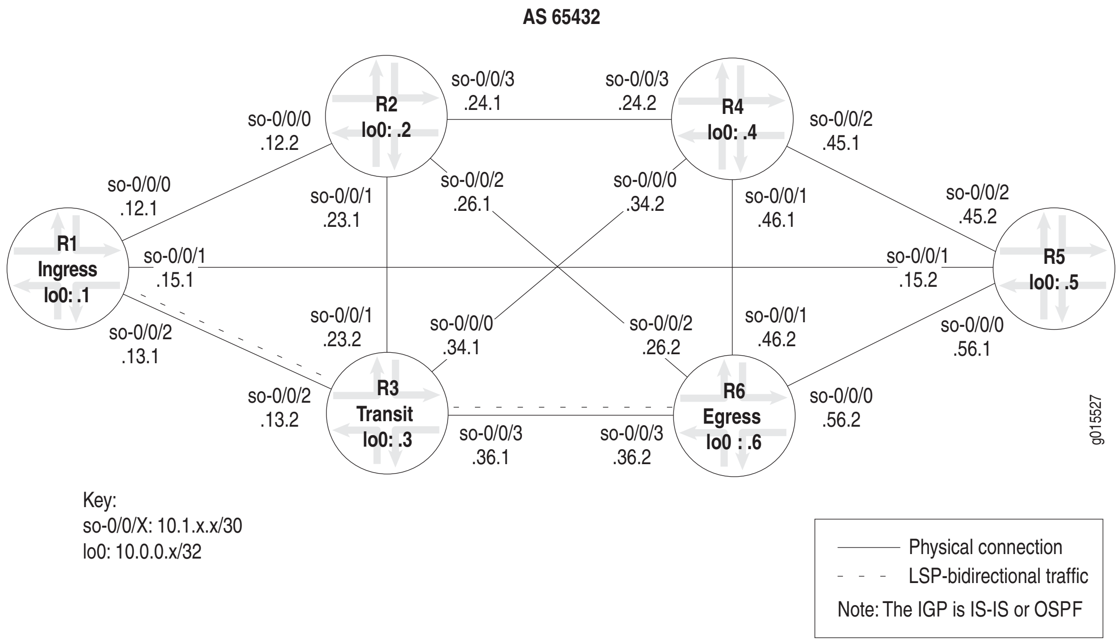

The MPLS network in Figure 7 illustrates a router-only network with SONET interfaces that consist of the following components:

- A full-mesh interior Border Gateway Protocol (IBGP) topology, using AS 65432
- MPLS and Resource Reservation Protocol (RSVP) enabled on all routers
- A send-statics policy on routers R1 and R6 that allows a new route to be advertised into the network
- An LSP between routers R1 and R6

The network shown in Figure 7 is a Border Gateway Protocol (BGP) full-mesh network. Since route reflectors and confederations are not used to propagate BGP learned routes, each router must have a BGP session with every other router running BGP. For the full configuration for each router in the example network, see “Configuring MPLS on a Network” on page 3.

Steps To Take To verify LSP use in your network, follow these steps:

1. Verifying an LSP on the Ingress Router on page 79
2. Verifying an LSP on a Transit Router on page 80

Verifying an LSP on the Ingress Router

Purpose

You can verify the availability of an LSP when it is up by examining the inet.3 routing table on the ingress router. The inet.3 routing table contains the host address of each LSP's egress router. This routing table is used on ingress routers to route BGP packets to the destination egress router. BGP uses the inet.3 routing table on the ingress router to help resolve next-hop addresses.

Action

To verify an LSP on an ingress router, enter the following JUNOS command-line interface (CLI) operational mode command:

user@host> **show route table inet.3**

Sample Output

```
user@R1> show route table inet.3

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.6/32        *[RSVP/7] 4w0d 22:40:57, metric 20
                    > via so-0/0/2.0, label-switched-path R1-to-R6
```

What It Means

The sample output shows the inet.3 routing table. By default, only BGP and MPLS virtual private networks (VPNs) can use the inet.3 route table to resolve next-hop information. One destination is listed in the route table, 10.0.0.6. This destination (10.0.0.6) is signaled by RSVP, and is the current active path, as indicated by the asterisk (*). The protocol preference for this route is 7, and the metric associated with it is 20. The label-switched path is R1-to-R6, through interface so-0/0/2.0, which is the physical next-hop transit interface.

Typically, the penultimate router in the LSP either pops the packet's label or changes the label to a value of 0. If the penultimate router pops the top label and an IPv4 packet is underneath, the egress router routes the IPv4 packet, consulting the IP routing table inet.0 to determine how to forward the packet. If another type of label (such as one created by Label Distribution Protocol (LDP) tunneling or VPNs, but not IPv4) is underneath the top label, the egress router does not examine the inet.0 routing table. Instead, it examines the mpls.0 routing table for forwarding decisions.

If the penultimate router changes the packet's label to a value of 0, the egress router strips off the 0 label, indicating that an IPv4 packet follows. The packet is examined by the inet.0 routing table for forwarding decisions.

When a transit or egress router receives an MPLS packet, information in the MPLS forwarding table is used to determine the next transit router in the LSP or whether this router is the egress router.

When BGP resolves a next-hop prefix, it examines both the inet.0 and inet.3 routing tables, seeking the next hop with the lowest preference; for example, RSVP preference 7 is preferred over OSPF preference 10. The RSVP signaled LSP is used to reach the BGP next hop. This is the default when the BGP next hop equals the LSP egress address. Once the BGP next hop is resolved through an LSP, the BGP traffic uses the LSP to forward BGP transit traffic.

Verifying an LSP on a Transit Router

Purpose You can verify the availability of an LSP when it is up by examining the mpls.0 routing table on a transit router. MPLS maintains the mpls.0 routing table, which contains a list of the next label-switched router in each LSP. This routing table is used on transit routers to route packets to the next router along an LSP.

Action To verify an LSP on a transit router, enter the following JUNOS CLI operational mode command:

user@host> **show route table mpls.0**

Sample Output

```
user@R3> show route table mpls.0

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                  *[MPLS/0] 7w3d 22:20:56, metric 1
                      Receive
1                  *[MPLS/0] 7w3d 22:20:56, metric 1
                      Receive
2                  *[MPLS/0] 7w3d 22:20:56, metric 1
                      Receive
100064             *[RSVP/7] 2w1d 04:17:36, metric 1
                    > via so-0/0/3.0, label-switched-path R1-to-R6
100064(S=0)        *[RSVP/7] 2w1d 04:17:36, metric 1
                    > via so-0/0/3.0, label-switched-path R1-to-R6
```

What It Means The sample output from transit router R3 shows route entries in the form of MPLS label entries, indicating that there is only one active route, even though there are five active entries.

The first three MPLS labels are reserved MPLS labels defined in RFC 3032. Packets received with these label values are sent to the Routing Engine for processing. Label 0 is the IPv4 explicit null label. Label 1 is the MPLS equivalent of the IP Router Alert label and Label 2 is the IPv6 explicit null label.

The two entries with the 100064 label are for the same LSP, R1-to-R6. There are two entries because the stack values in the MPLS header may be different. The second entry, 100064 (S=0), indicates that the stack depth is not 1 and additional label values are included in the packet. In contrast, the first entry of 100064 has an inferred S = 1 which indicates a stack depth of 1 and makes it the last label in the packet. The dual entry indicates that this is the penultimate router. For more information on MPLS label stacking, see RFC 3032, *MPLS Label Stack Encoding*.

The incoming label is the MPLS header of the MPLS packet, and is assigned by RSVP to the upstream neighbor. Juniper Networks routers dynamically assign labels for RSVP traffic-engineered LSPs in the range from 100,000 through 1,048,575.

The router assigns labels starting at label 100,000, in increments of 16. The sequence of label assignments is 100,000, 100,016, 100,032, 100,048, and so on. At the end of the assigned labels, the label numbers start over at 100001, incrementing in units of 16. Juniper Networks reserves labels for various purposes. Table 11 lists the various label range allocations for incoming labels.

Table 11: MPLS Label Range Allocations

Incoming Label	Status
0 through 15	Reserved by IETF
16 through 1023	Reserved for static LSP assignment
1024 through 9999	Reserved for internal use (for example, CCC labels)
10,000 through 99,999	Reserved for static LSP assignment
100,000 through 1,048,575	Reserved for dynamic label assignment