Chapter 12

Checking the BGP Layer

This chapter describes how to check the Border Gateway Protocol (BGP) layer of the layered Multiprotocol Label Switching (MPLS) model. (See Table 18.)

Table 18: Checklist for Checking the BGP Layer

Checking the BGP Layer Tasks	Command or Action
Checking the BPG Layer on page 180	
1. Check That BGP Traffic Is Using the LSP on page 182	traceroute *hostname*
2. Check BGP Sessions on page 182	show bgp summary
3. Verify the BGP Configuration on page 183	show configuration
4. Examine BGP Routes on page 189	show route *destination-prefix* detail
5. Verify Received BGP Routes on page 190	show route receive protocol bgp *neighbor-address*
6. Take Appropriate Action on page 191	The following sequence of commands addresses the specific problem described in this section: [edit] edit protocols bgp [edit protocols bgp] show set local-address 10.0.0.1 delete group internal neighbor 10.1.36.2 show commit
7. Check That BGP Traffic Is Using the LSP Again on page 192	traceroute *hostname*

Checking the BPG Layer

Purpose After you have configured the label-switched path (LSP) and determined that it is up, and configured BGP and determined that sessions are established, ensure that BGP is using the LSP to forward traffic.

Figure 23 illustrates the BGP layer of the layered MPLS model.

Figure 23: Checking the BGP Layer

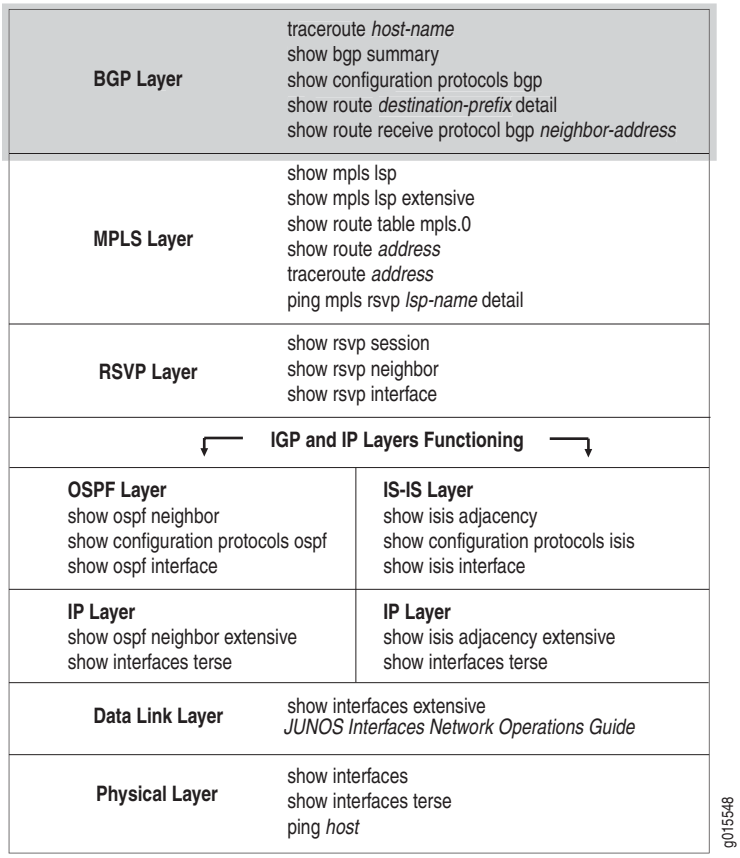

When you check the BGP layer, you verify that the route is present and active, and more importantly, you ensure that the next hop is the LSP. There is no point in checking the BGP layer unless the LSP is established, because BGP uses the MPLS LSP to forward traffic. If the network is not functioning at the BGP layer, the LSP does not work as configured.

Figure 24 illustrates the MPLS network used in this chapter.

Figure 24: MPLS Network Broken at the BGP Layer

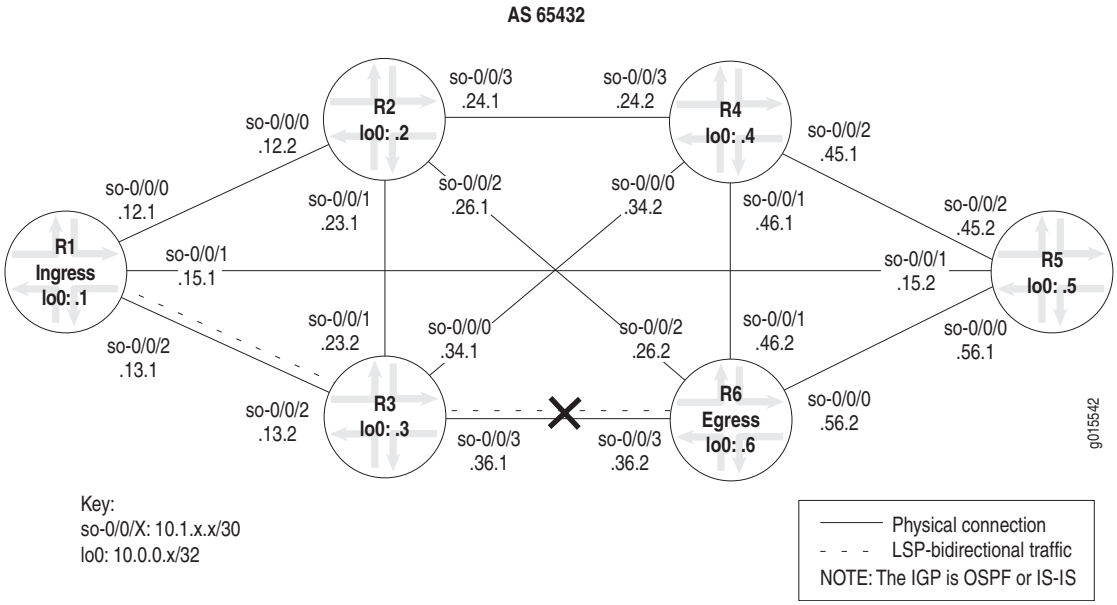

The network shown in Figure 24 is a fully meshed configuration where every directly connected interface can receive and send packets to every other similar interface. The LSP in this network is configured to run from ingress router R1, through transit router R3, to egress router R6. In addition, a reverse LSP is configured to run from R6 through R3 to R1, creating bidirectional traffic.

The cross shown in Figure 24 indicates where BGP is not being used to forward traffic through the LSP. Possible reasons for the LSP not working correctly are that the destination IP address of the LSP does not equal the BGP next hop or that BGP is not configured properly.

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

1. Check That BGP Traffic Is Using the LSP on page 182
2. Check BGP Sessions on page 182
3. Verify the BGP Configuration on page 183
4. Examine BGP Routes on page 189
5. Verify Received BGP Routes on page 190
6. Take Appropriate Action on page 191
7. Check That BGP Traffic Is Using the LSP Again on page 192

Step 1: Check That BGP Traffic Is Using the LSP

Purpose At this level of the troubleshooting model, BGP and the LSP may be up, however BGP traffic might not be using the LSP to forward traffic.

Action To verify that BGP traffic is using the LSP, enter the following JUNOS command-line interface (CLI) operational mode command from the ingress router:

user@host> **traceroute** *hostname*

Sample Output

```
user@R1> traceroute 100.100.6.1
traceroute to 100.100.6.1 (100.100.6.1), 30 hops max, 40 byte packets
 1  10.1.13.2 (10.1.13.2)  0.653 ms  0.590 ms  0.543 ms
 2  10.1.36.2 (10.1.36.2)  0.553 ms !N  0.552 ms !N  0.537 ms !N

user@R6> traceroute 100.100.1.1
traceroute to 100.100.1.1 (100.100.1.1), 30 hops max, 40 byte packets
 1  10.1.36.1 (10.1.36.1)  0.660 ms  0.551 ms  0.526 ms
 2  10.1.13.1 (10.1.13.1)  0.568 ms !N  0.553 ms !N  0.536 ms !N
```

What It Means The sample output shows that BGP traffic is not using the LSP, consequently MPLS labels do not appear in the output. Instead of using the LSP, BGP traffic is using the interior gateway protocol (IGP) (IS-IS or OSPF, in the example network shown in Figure 24 on page 181) to reach the BGP next-hop LSP egress address for R6 and R1. The JUNOS software default is to use LSPs for BGP traffic when the BGP next hop equals the LSP egress address.

Step 2: Check BGP Sessions

Purpose Display summary information about BGP and its neighbors to determine if routes are received from peers in the autonomous system (AS). When a BGP session is established, the peers are exchanging update messages.

Action To check that BGP sessions are up, enter the following JUNOS CLI operational mode command from the ingress router:

user@host> **show bgp summary**

Sample Output 1

```
user@R1> show bgp summary
Groups: 1 Peers: 6 Down peers: 1
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0                 1          1          0          0          0          0
Peer              AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn State|#Active/Received/Damped...
10.0.0.2        65432      11257      11259       0       0 3d 21:49:57 0/0/0                0/0/0
10.0.0.3        65432      11257      11259       0       0 3d 21:49:57 0/0/0                0/0/0
10.0.0.4        65432      11257      11259       0       0 3d 21:49:57 0/0/0                0/0/0
10.0.0.5        65432      11257      11260       0       0 3d 21:49:57 0/0/0                0/0/0
10.0.0.6        65432          4       4572       0       1 3d 21:46:59 Active
10.1.36.2       65432      11252      11257       0       0 3d 21:46:49 1/1/0                0/0/0
```

Sample Output 2

```
user@R1> show bgp summary
Groups: 1 Peers: 5 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0                 1          1          0          0          0          0
Peer              AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn State|#Active/Received/Damped...
10.0.0.2        65432         64         68       0       0       32:18 0/0/0                 0/0/0
10.0.0.3        65432         64         67       0       0       32:02 0/0/0                 0/0/0
10.0.0.4        65432         64         67       0       0       32:10 0/0/0                 0/0/0
10.0.0.5        65432         64         67       0       0       32:14 0/0/0                 0/0/0
10.0.0.6        65432         38         39       0       1       18:02 1/1/0                 0/0/0
```

What It Means Sample Output 1 shows that one peer (egress router 10.0.0.6) is not established, as indicated by the Down Peers: 1 field. The last column (State|#Active/Received/Damped) shows that peer 10.0.0.6 is active, indicating that is it not established. All other peers are established as indicated by the number of active, received, and damped routes. For example, 0/0/0 for peer 10.0.0.2 indicates that no BGP routes were active or received in the routing table, and no BGP routes were damped; 1/1/0 for peer 10.1.36.2 indicates that one BGP route was active and received in the routing table, and no BGP routes were damped.

If the output of the show bgp summary command of an ingress router shows that a neighbor is down, check the BGP configuration. For information on checking the BGP configuration, see "Verify the BGP Configuration" on page 183.

Sample Output 2 shows output from ingress router R1 after the BGP configurations on R1 and R6 were corrected in "Take Appropriate Action" on page 191. All BGP peers are established and one route is active and received. No BGP routes were damped.

If the output of the show bgp summary command shows that a neighbor is up but packets are not being forwarded, check for received routes from the egress router. For information on checking the egress router for received routes, see "Verify Received BGP Routes" on page 190.

Step 3: Verify the BGP Configuration

Purpose For BGP to run on the router, you must define the local AS number, configure at least one group, and include information about at least one peer in the group (the peer's IP address and AS number). When BGP is part of an MPLS network, you must ensure that the LSP is configured with a destination IP address equal to the BGP next hop in order for BGP routes to be installed with the LSP as the next hop for those routes.

Action To verify the BGP configuration, enter the following JUNOS CLI operational mode command:

```
user@host> show configuration
```

Sample Output 1

```
user@R1> show configuration
[...Output truncated...]
interfaces {
    so-0/0/0 {
        unit 0 {
            family inet {
                address 10.1.12.1/30;
            }
            family iso;
            family mpls;
        }
    }
    so-0/0/1 {
        unit 0 {
            family inet {
                address 10.1.15.1/30;
            }
            family iso;
            family mpls;
        }
    }
    so-0/0/2 {
        unit 0 {
            family inet {
                address 10.1.13.1/30;
            }
            family iso;
            family mpls;
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.70.143/21;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 10.0.0.1/32;
            }
            family iso {
                address 49.0004.1000.0000.0001.00;
            }
        }
    }
}
routing-options {
    [...Output truncated...]
        route 100.100.1.0/24 reject;
    }
    router-id 10.0.0.1;
    autonomous-system 65432;
}
protocols {
    rsvp {
        interface so-0/0/0.0;
        interface so-0/0/1.0;
        interface so-0/0/2.0;
        interface fxp0.0 {
            disable;
        }
```

```
    }
    mpls {
        label-switched-path R1-to-R6 {
            to 10.0.0.6;     <<< destination address of the LSP
        }
        inactive: interface so-0/0/0.0;
        inactive: interface so-0/0/1.0;
        interface so-0/0/2.0;
        interface fxp0.0 {
            disable;
        }
    }
    bgp {
        export send-statics;        <<< missing local-address statement
        group internal {
            type internal;
            neighbor 10.0.0.2;
            neighbor 10.0.0.5;
            neighbor 10.0.0.4;
            neighbor 10.0.0.6;
            neighbor 10.0.0.3;
            neighbor 10.1.36.2;     <<< incorrect interface address
        }
    }
    isis {
        level 1 disable;
        interface so-0/0/0.0;
        interface so-0/0/1.0;
        interface so-0/0/2.0;
        interface all {
            level 2 metric 10;
        }
        interface fxp0.0 {
            disable;
        }
        interface lo0.0 {
            passive;
        }
    }
    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 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

```
user@R6> show configuration
[...Output truncated...]
interfaces {
    so-0/0/0 {
        unit 0 {
            family inet {
                address 10.1.56.2/30;
            }
            family iso;
            family mpls;
        }
    }
    so-0/0/1 {
        unit 0 {
            family inet {
                address 10.1.46.2/30;
            }
            family iso;
            family mpls;
        }
    }
    so-0/0/2 {
        unit 0 {
            family inet {
                address 10.1.26.2/30;
            }
            family iso;
            family mpls;
        }
    }
    so-0/0/3 {
        unit 0 {
            family inet {
                address 10.1.36.2/30;
            }
            family iso;
            family mpls;
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.70.148/21;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 10.0.0.6/32;
                address 127.0.0.1/32;
            }
            family iso {
                address 49.0004.1000.0000.0006.00;
            }
        }
    }
}
```

```
routing-options {
    [...Output truncated...]
        route 100.100.6.0/24 reject;
    }
    router-id 10.0.0.6;
    autonomous-system 65432;
}
protocols {
    rsvp {
        interface so-0/0/0.0;
        interface so-0/0/1.0;
        interface so-0/0/2.0;
        interface so-0/0/3.0;
        interface fxp0.0 {
            disable;
        }
    }
    mpls {
        label-switched-path R6-to-R1 {
            to 10.0.0.1;    <<< destination address of the reverse LSP
        }
        inactive: interface so-0/0/0.0;
        inactive: interface so-0/0/1.0;
        inactive: interface so-0/0/2.0;
        interface so-0/0/3.0;
    }
    bgp {
        group internal {
            type internal;
            export send-statics;    <<< missing local-address statement
            neighbor 10.0.0.2;
            neighbor 10.0.0.3;
            neighbor 10.0.0.4;
            neighbor 10.0.0.5;
            neighbor 10.0.0.1;
            neighbor 10.1.13.1;     <<< incorrect interface address
        }
    }
    isis {
        level 1 disable;
        interface all {
            level 2 metric 10;
        }
        interface fxp0.0 {
            disable;
        }
        interface lo0.0 {
            passive;
        }
    }
    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;
            interface lo0.0 {
                passive;
            }
        }
    }
}
```

```
policy-options {
    policy-statement send-statics {
        term statics {
            from {
                route-filter 100.100.6.0/24 exact;
            }
            then accept;
        }
    }
}
```

What It Means The sample output shows the BGP configurations on ingress router R1 and egress router R6. Both configurations show the local AS (65432), one group (internal), and six peers configured. The underlying interior gateway protocol is IS-IS, and the relevant interfaces are configured to run IS-IS.

NOTE: In this configuration, the RID is manually configured to avoid any duplicate RID problems, and all interfaces configured with BGP include the family inet statement at the [edit interfaces *type-fpc/pic/port* unit *logical-unit-number*] hierarchy level.

Sample output for ingress router R1 and egress router R6 shows that the BGP protocol configuration is missing the local-address statement for the internal group. When the local-address statement is configured, BGP packets are forwarded from the local router loopback (lo0) interface address, which is the address to which BGP peers are peering. If the local-address statement is not configured, BGP packets are forwarded from the outgoing interface address, which does not match the address to which BGP peers are peering, and BGP does not come up.

On the ingress router, the IP address (10.0.0.1) in the local-address statement should be the same as the address configured for the LSP on the egress router (R6) in the to statement at the [edit protocols mpls label-switched-path *lsp-path-name*] hierarchy level. BGP uses this address, which is identical to the LSP address, to forward BGP traffic through the LSP.

In addition, the BGP configuration on R1 includes two IP addresses for R6, an interface address (10.1.36.2) and a loopback (lo0) interface address (10.0.0.6), resulting in the LSP destination address (10.0.0.6) not matching the BGP next-hop address (10.1.36.2). The BGP configuration on R6 also includes two IP addresses for R1, an interface address (10.1.13.1) and a loopback (lo0) interface address, resulting in the reverse LSP destination address (10.0.0.1) not matching the BGP next-hop address (10.1.13.1).

In this instance, because the local-address statement is missing in the BGP configurations of both routers and the LSP destination address does not match the BGP next-hop address, BGP is not using the LSP to forward traffic.

Step 4: Examine BGP Routes

Purpose You can examine the BGP path selection process to determine the single, active path when BGP receives multiple routes to the same destination. In this step, we examine the reverse LSP R6-to-R1, making R6 the ingress router for that LSP.

Action To examine BGP routes and route selection, enter the following JUNOS CLI operational mode command:

user@host> **show route** *destination-prefix* **detail**

Sample Output 1

```
user@R6> show route 100.100.1.1 detail

inet.0: 30 destinations, 46 routes (29 active, 0 holddown, 1 hidden)
100.100.1.0/24 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Source: 10.1.13.1
                Next hop: via so-0/0/3.0, selected
                Protocol next hop: 10.1.13.1 Indirect next hop: 8671594 304
                State: <Active Int Ext>
                Local AS: 65432 Peer AS: 65432
                Age: 4d 5:15:39         Metric2: 2
                Task: BGP_65432.10.1.13.1+3048
                Announcement bits (2): 0-KRT 4-Resolve inet.0
                AS path: I
                Localpref: 100
                Router ID: 10.0.0.1
```

Sample Output 2

```
user@R6> show route 100.100.1.1 detail

inet.0: 30 destinations, 46 routes (29 active, 0 holddown, 1 hidden)
100.100.1.0/24 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Source: 10.0.0.1
                Next hop: via so-0/0/3.0 weight 1, selected
                Label-switched-path R6-to-R1
                Label operation: Push 100000
                Protocol next hop: 10.0.0.1 Indirect next hop: 8671330 301
                State: <Active Int Ext>
                Local AS: 65432 Peer AS: 65432
                Age: 24:35      Metric2: 2
                Task: BGP_65432.10.0.0.1+179
                Announcement bits (2): 0-KRT 4-Resolve inet.0
                AS path: I
                Localpref: 100
                Router ID: 10.0.0.1
```

What It Means Sample Output 1 shows that the BGP next hop (10.1.13.1) does not equal the LSP destination address (10.0.0.1) in the to statement at the [edit protocols mpls label-switched-path *label-switched-path-name*] hierarchy level when the BGP configuration of R6 and R1 is incorrect.

Sample Output 2, taken after the configurations on R1 and R6 are corrected, shows that the BGP next hop (10.0.0.1) and the LSP destination address (10.0.0.1) are the same, indicating that BGP can use the LSP to forward BGP traffic.

Step 5: Verify Received BGP Routes

Purpose Display the routing information received on router R6, the ingress router for the reverse LSP R6-to-R1.

Action To verify that a particular BGP route is received on the egress router, enter the following JUNOS CLI operational mode command:

user@host> **show route receive protocol bgp** *neighbor-address*

Sample Output 1

```
user@R6> show route receive-protocol bgp 10.0.0.1

inet.0: 30 destinations, 46 routes (29 active, 0 holddown, 1 hidden)
    <<< missing route
inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

__juniper_private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0
hidden)
```

Sample Output 2

```
user@R6> show route receive-protocol bgp 10.0.0.1

inet.0: 30 destinations, 46 routes (29 active, 0 holddown, 1 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 100.100.1.0/24          10.0.0.1                     100        I

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

__juniper_private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0
hidden)
```

What It Means Sample Output 1 shows that ingress router R6 (reverse LSP R6-to-R1) does not receive any BGP routes into the inet.0 routing table when the BGP configurations of R1 and R6 are incorrect.

Sample Output 2 shows a BGP route installed in the inet.0 routing table after the BGP configurations on R1 and R6 are corrected using “Take Appropriate Action” on page 191.

Step 6: 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 ingress and egress routers are incorrectly configured for BGP to forward traffic using the LSP.

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

1. On ingress router R1, include the local-address statement and delete the incorrect interface address (repeat these steps on egress router R6):

```
[edit]
user@R1# edit protocols bgp
[edit protocols bgp]
user@R1# show
user@R1# set local-address 10.0.0.1
user@R1# delete group internal neighbor 10.1.36.2
```

2. Verify and commit the configuration:

```
[edit protocols bgp]
user@R1# show
user@R1# commit
```

Sample Output

```
[edit]
user@R1# edit protocols bgp

[edit protocols bgp]
user@R1# show
export send-statics;
group internal {
    type internal;
    neighbor 10.0.0.2;
    neighbor 10.0.0.5;
    neighbor 10.0.0.4;
    neighbor 10.0.0.6;
    neighbor 10.0.0.3;
    neighbor 10.1.36.2;
}

[edit protocols bgp]
user@R1# set local-address 10.0.0.1

[edit protocols bgp]
user@R1# delete group internal neighbor 10.1.36.2

[edit protocols bgp]
user@R1# show
local-address 10.0.0.1;
export send-statics;
group internal {
    type internal;
    neighbor 10.0.0.2;
    neighbor 10.0.0.5;
    neighbor 10.0.0.4;
    neighbor 10.0.0.6;
    neighbor 10.0.0.3;
}
```

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

What It Means The sample output shows that the configuration of BGP on ingress router R1 is now correct. BGP can now forward BGP traffic through the LSP.

Step 7: Check That BGP Traffic Is Using the LSP Again

Action To verify that BGP traffic is using the LSP, enter the following JUNOS CLI operational mode command from the ingress router:

user@host> **traceroute** *hostname*

Sample Output

```
user@R1> traceroute 100.100.6.1
traceroute to 100.100.6.1 (100.100.6.1), 30 hops max, 40 byte packets
 1  10.1.13.2 (10.1.13.2)  0.858 ms  0.740 ms  0.714 ms
     MPLS Label=100016 CoS=0 TTL=1 S=1
 2  10.1.36.2 (10.1.36.2)  0.592 ms !N  0.564 ms !N  0.548 ms !N

user@R6> traceroute 100.100.1.1
traceroute to 100.100.1.1 (100.100.1.1), 30 hops max, 40 byte packets
 1  10.1.36.1 (10.1.36.1)  0.817 ms  0.697 ms  0.771 ms
     MPLS Label=100000 CoS=0 TTL=1 S=1
 2  10.1.13.1 (10.1.13.1)  0.581 ms !N  0.567 ms !N  0.544 ms !N
```

What It Means The sample output shows that MPLS labels are used to forward packets through the LSP. Included in the output is a label value (MPLS Label=100016), the time-to-live value (TTL=1), and the stack bit value (S=1).

The MPLS Label field is used to identify the packet to a particular LSP. It is a 20-bit field, with a maximum value of (2^^20-1), approximately 1,000,000.

The time-to-live (TTL) value contains a limit on the number of hops that this MPLS packet can travel through the network (1). It is decremented at each hop, and if the TTL value drops below one, the packet is discarded.

The bottom of the stack bit value (S=1) indicates that is the last label in the stack and that this MPLS packet has one label associated with it. The MPLS implementation in the JUNOS software supports a stacking depth of 3 on the M-series routers and up to 5 on the T-series routing platforms. For more information on MPLS label stacking, see RFC 3032, *MPLS Label Stack Encoding*.

MPLS labels appear in the sample output because the traceroute command is issued to a BGP destination where the BGP next hop for that route is the LSP egress address. The JUNOS software by default uses LSPs for BGP traffic when the BGP next hop equals the LSP egress address.

If the BGP next hop does not equal the LSP egress address, the BGP traffic does not use the LSP, and consequently MPLS labels do not appear in the output for the traceroute command, as indicated in the sample output in “Check BGP Sessions” on page 182.