Chapter 7

Static and Explicit-Path LSP Configuration Guidelines

The following sections describe how to configure static and explicit-path label-switched paths (LSPs):

Configuring Static LSPs

To configure static LSPs, configure the ingress router and each router along the path up to and including the egress router.

For the ingress router, configure which packets to tag (based on the packet's IP destination address), the next router in the LSP, and the tag to apply to the packet. Manually assigned labels can have values in the range 16 through 1023. Optionally, you can apply preference and class-of-service (CoS) values to the packets.

For the intermediate routers in the path, configure the next router in the path and the tag to apply to the packet. Again, you can optionally apply preference and CoS values to the packets.

For the egress router, you generally just remove the label and continue forwarding the packet to the next hop. However, if the previous router removed the label, the egress router examines the packet's IP header and forwards the packet toward its IP destination.

To configure static Multiprotocol Label Switching (MPLS), perform the following tasks:

Configuring the Ingress Router for Static LSPs

The ingress router checks the IP address in the incoming packet's destination address field and, if it finds a match in the routing table, applies the label associated with that address to the packets. The label has forwarding information associated with it, including the address of the next-hop router, and the route preference and CoS values.

To configure static LSPs on the ingress router, include the static-path statement:

```
static-path inet {
    prefix {
        next-hop (address | interface-name | address/interface-name)
        push out-label;
        class-of-service value;
        preference preference;
    }
}
```

You can include this statement at the following hierarchy levels:

[edit protocols mpls]

[edit logical-routers *logical-router-name* protocols mpls]

The next-hop and push statements are required; the other statements are optional.

Each static-path statement consists of the following parts:

Criteria to use to analyze an incoming packet:

The inet option creates an LSP that handles IPv4 packets. All static MPLS routes created using the inet option are installed in the default IPv4 routing table (inet.0), and the creating protocol is identified as static. This process is no different from creating static IPv4 routes at the [edit routing-options static] hierarchy level.

In the *prefix* option, you configure the IP destination address to check when incoming packets are analyzed. If the address matches, the specified label, *out-label*, is assigned to the packet, and the packet enters an LSP. Each prefix that you specify is installed as a static route in the routing table. You can specify one or more *prefix* statements at the [edit protocols mpls static-path] hierarchy level.

The next-hop statement:

This statement supplies the IP address of the next hop to the destination. You can specify this as the IP address of the next hop, the interface name (for point-to-point interfaces only), or as *address*/*interface-name* to specify an IP address on an operational interface. When the next hop is on a directly attached interface, the route is installed in the routing table. You cannot configure a LAN or nonbroadcast multiaccess (NMBA) interface as a next-hop interface.

Label properties applied to the packet in the LSP:

Label to apply to the packet (push *out-label*)—The label is a 20-bit integer, so it can be a number in the range 0 through 1,048,575 (2^{20}– 1). Dynamic MPLS assigns the labels 100,000 through 1,048,575, so if your network uses both static and dynamic MPLS, we recommend that you use labels 16 through 1023 and 10,000 through 99,999 only for static MPLS. (Labels 0 through 15 are reserved and require special semantics. Labels 1024 through 9999 are reserved for future applications.)

Preference of this route (preference *preference*).

CoS value to apply to the packet (class-of-service *cos-value*).

To determine whether a static ingress route is installed, use the command show route table inet.0 protocol static. The following is a sample output. The push keyword identifies that a label is to be added in front of IP packet.

```
10.0.0.0/8              *[Static/5] 00:01:48
                        > to 11.1.1.1 via so-0/0/0, push 123
```

Example: Configuring the Ingress Router

Configure the ingress router for a static LSP that consists of three routers (see Figure 20).

Figure 20: Static MPLS Configuration

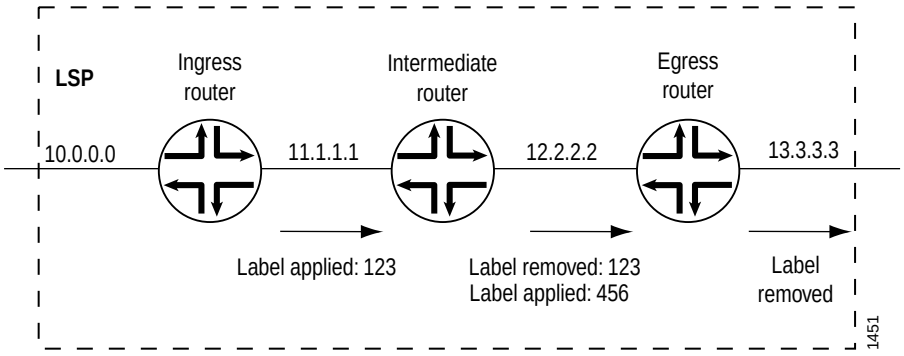

For packets addressed to 10.0.0.0/8, assign label 123 and transmit them to the next-hop router at 11.1.1.1:

```
[edit]
interfaces {
    so-0/0/0 {
        unit 0 {
            family mpls;
        }
    }
}
protocols {
    mpls {
        static-path inet {
            10.0.0.0/8 {
                next-hop 11.1.1.1;
                push 123;
            }
        }
        interface so-0/0/0;
    }
}
```

To determine whether the static ingress route is installed, use the following command:

```
user@host> show route table inet.0 protocol static
```

The following is a sample of the output. The push 123 keyword identifies the route.

```
10.0.0.0/8          *[Static/5] 00:01:48
                    > to 11.1.1.1 via so-0/0/0, push 123
```

Configuring the Intermediate and Egress Routers for Static LSPs

Intermediate and egress routers perform similar functions—they modify the label that has been applied to a packet. An intermediate router can change the label. An egress router removes the label (if the packet still contains a label) and continues forwarding the packet to its destination.

To configure static LSPs on intermediate and egress routers, include the interface statement:

```
interface interface-name {
    disable;
    admin-group {
        group-name;
    }
    label-map in-label {
        class-of-service value;
        next-hop (<address | interface-name>) | (reject | discard);
        (pop | swap <out-label> | swap-push <swap-label>);
        preference preference;
        type type;
    }
}
```

You can configure these statements at the following hierarchy levels:

[edit logical-routers *logical-router-name* protocols mpls]

[edit protocols mpls]

For the label-map statement configuration, the next-hop | (reject | discard) and pop | swap | swap-push statements are required. The remaining statements are optional.

Each statement within the interface statement consists of the following parts:

Criteria to use to analyze the labeled packet. Two criteria are used: the interface on which the packet was received (specified in the opening interface statement itself) and the packet's label (specified in the label-map statement).

The next-hop statement. This statement supplies the IP address of the next hop to the destination, specified as the IP address of the next hop, or the interface name (for point-to-point interfaces only), or *address* and *interface-name* to specify an IP address on an operational interface. When the specified next hop is on a directly attached interface, this route is installed in the routing table. You cannot configure a LAN or NBMA interface as a next-hop interface.

Operation to perform on the labeled packet:

For egress routers, remove the packet's label altogether (pop).

For intermediate routers only, exchange the label for another label (swap *out-label*).

Discard the packet, sending an ICMP unreachable message to the packet's originator (reject).

Discard the packet without sending an ICMP unreachable message to the packet's originator (discard).

Label properties to apply to the packet (all are optional):

Type of traffic in the LSP. Currently, the type can be IPv4 only (type inet), which is the default.

Preference value for this route (preference *preference*).

For intermediate routers only, the CoS value to apply to the packet (class-of-service *cos-value*).

You can specify any number of label-map statements at the [edit protocols mpls interface *interface-name*] hierarchy level.

The static routes are installed in the default MPLS routing table, mpls.0, and the creating protocol is identified as static. To verify that a static route is properly installed, use the command show route table mpls.0 protocol static. The following is an example of the output:

```
123                *[Static/5] 00:00:38
                    > to 12.2.2.2 via so-5/0/0.0, swap 456
```

Example: Configuring an Intermediate Router

For packets labeled 123 arriving on interface so-0/0/0, assign the label 456, and transmit them to the next-hop router at 12.2.2.2:

```
[edit]
interfaces {
    so-0/0/0 {
        unit 0 {
            family mpls;
        }
    }
}
protocols {
    mpls {
        interface so-0/0/0 {
            label-map 123 {
                next-hop 12.2.2.2;
                swap 456;
            }
        }
    }
}
```

To determine whether the static intermediate route is installed, use the following command:

```
user@host> show route table mpls.0 protocol static
```

The following is a sample of the output. The swap 456 keyword identifies the route.

```
123                *[Static/5] 00:01:48
                    > to 12.2.2.2 via so-0/0/0, swap 456
```

Example: Configuring an Egress Router

For packets labeled 456 arriving on interface so-0/0/0, remove the label and transmit the packets to the next-hop router at 13.3.3.3:

```
[edit]
interfaces {
    so-0/0/0 {
        unit 0 {
            family mpls;
        }
    }
}
protocols {
    mpls {
        interface so-0/0/0 {
            label-map 456 {
                next-hop 13.3.3.3;
                pop;
            }
        }
    }
}
```

To determine whether the static egress route is installed, use the following command:

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

The following is a sample of the output. The pop keyword identifies the egress route.

```
456                *[Static/5] 00:01:48
                    > to 13.3.3.3 via so-0/0/0, pop
```

Configuring Static Unicast Routes for Point-to-Multipoint LSPs

You can configure a static unicast IP route with a point-to-multipoint LSP as the next hop. For more information on point-to-multipoint LSPs, see "Point-to-Multipoint LSPs" on page 54, "Configuring Point-to-Multipoint LSPs" on page 107, and "Configuring CCC Switching for Point-to-Multipoint LSPs" on page 383.

To configure a static unicast route for a point-to-multipoint LSP, complete the following steps:

1. On the ingress PE router, configure a static IP unicast route with the point-to-multipoint LSP name as the next hop by including the p2mp-lsp-next-hop statement:

   p2mp-lsp-next-hop *point-to-multipoint-lsp-next-hop*;

   This statement can be configured at the following hierarchy levels:

   [edit logical-routers *logical-router-name* routing-options static route *route-name*]

   [edit routing-options static route *route-name*]

2. On the egress PE router, configure a static IP unicast route with the same destination address configured in Step 1 (the address configured at the [edit routing-options static route] hierarchy level) using the next-hop option of the static statement:

   next-hop *address*;

   This option can be configured at the following hierarchy levels:

   [edit logical-routers *logical-router-name* routing-options static route *route-name*]

   [edit routing-options static route *route-name*]

NOTE: CCC and static routes cannot use the same point-to-multipoint LSP.

For more information on static routes, see the *JUNOS Routing Protocols Configuration Guide*.

The following show route command output displays a unicast static route pointing to a point-to-multipoint LSP on the ingress PE router where the LSP has two branch next hops:

```
rahul@pro1-d> show route 5.5.5.5 detail
inet.0: 29 destinations, 30 routes (28 active, 0 holddown, 1 hidden)
5.5.5.5/32 (1 entry, 1 announced)
        *Static Preference: 5
                Next hop type: Flood
                Next hop: via so-0/3/2.0 weight 1
                Label operation: Push 100000
                Next hop: via t1-0/1/1.0 weight 1
                Label operation: Push 100064
                State: <Active Int Ext>
                Local AS: 10458
                Age: 2:41:15
                Task: RT
                Announcement bits (2): 0-KRT 3-BGP.0.0.0.0+179
                AS path: I
```

Configuring Explicit-Path LSPs

If you disable constrained-path label-switched path (LSP) computation, as described in “Disabling Constrained-Path LSP Computation” on page 92, you must configure LSPs manually. Experimenting with particular explicit paths can familiarize you with MPLS.

When explicit-path LSPs are configured, the LSP is established along the path you specified. If the path is topologically not feasible, either because the network is partitioned or insufficient resources are available along some parts of the path, the LSP will fail. No alternative paths can be used. If the setup succeeds, the LSP stays on the defined path indefinitely.

To configure an explicit-path LSP, follow these steps:

1. Configure the path information in a named path, as described in “Creating a Named Path” on page 64. To configure complete path information, specify every router hop between the ingress and egress routers, preferably using the strict attribute. To configure incomplete path information, specify only a subset of router hops, using the loose attribute in places where the path is incomplete.

   For incomplete paths, the MPLS routers complete the path by querying the local routing table. This query is done on a hop-by-hop basis, and each router can figure out only enough information to reach the next explicit hop. It might be necessary to traverse a number of routers to reach the next (loose) explicit hop.

   Configuring incomplete path information creates portions of the path that depend on the current routing table, and this portion of the path can reroute itself as the topology changes. Therefore, an explicit-path LSP that contains incomplete path information is not completely fixed. These types of LSPs have only a limited ability to repair themselves, and they tend to create loops or flaps depending on the contents of the local routing table.

2. To configure the LSP and point it to the named path, use either the primary or secondary statement, as described in “Configuring the Primary and Secondary LSPs” on page 72.

3. Disable constrained-path LSP computation by including the no-cspf statement either as part of the LSP or as part of a primary or secondary statement. For more information, see “Disabling Constrained-Path LSP Computation” on page 92.

4. Configure any other LSP properties.

Using explicit-path LSPs has the following drawbacks:

More configuration effort is required.

Configured path information cannot take into account dynamic network bandwidth reservation, so the LSPs tend to fail when resources become depleted.

When an explicit-path LSP fails, you might need to manually repair it.

Because of these limitations, we recommend that you use explicit-path LSPs only in controlled situations, such as to enforce an optimized LSP placement strategy resulting from computations with an offline simulation software package.