

Chapter 1

Configuring MPLS on a Network

This chapter describes how to configure a network to run Multiprotocol Label Switching (MPLS), including the components and supporting protocols. (See Table 5.)

Table 5: Checklist for Verifying a Network Configured with MPLS

| Verifying a Network Configured with MPLS Tasks | Command or Action |
|---|---|
| Configuring MPLS on Your Network on page 6 | |
| 1. Configure IP Addresses on Router Interfaces on page 8 | [edit] edit interfaces <i>type-fpc/pic/port</i> unit <i>logical-unit-number</i> set family inet address <i>address</i> show commit |
| 2. Configure IS-IS as the IGP on page 9 | |
| a. Enable IS-IS on Routers in Your Network on page 10 | [edit] edit protocols isis set level 1 disable set interface <i>type-fpc/pic/port</i> level <i>level-number</i> metric <i>metric</i> set interface fxp0.0 disable set interface lo0.0 set interface lo0 passive show commit |
| b. Configure ISO Addressing on page 12 | [edit] edit interfaces set lo0 unit <i>number</i> family iso address <i>address</i> show commit |
| c. Enable IS-IS on Router Interfaces on page 14 | [edit] edit interfaces set <i>type-fpc/pic/port</i> unit <i>number</i> family iso show commit |
| d. Verify That IS-IS Adjacencies Are Established on page 15 | show isis adjacency |

| Verifying a Network Configured with MPLS Tasks | Command or Action |
|--|--|
| 3. Configure OSPF as the IGP on page 16 | |
| a. Enable OSPF on Routers in Your Network on page 17 | <pre>[edit] edit protocols ospf [edit protocols ospf] set area <i>area-id</i> interface <i>type-fpc/pic/port</i> set interface fxp0.0 disable set area 0.0.0.0 interface lo0 set area 0.0.0.0 interface lo0 passive set traffic engineering [edit routing-options] set router-id <i>router-id</i> show commit</pre> |
| b. Verify That OSPF Neighbors Are Established on page 19 | show ospf neighbor |
| 4. Set Up BGP on Routers in Your Network on page 21 | |
| a. Define the Local Autonomous System on page 22 | <pre>[edit] edit routing-options set autonomous-system <i>as-number</i> show commit</pre> |
| b. Configure BGP Neighbor Connections on page 23 | <pre>[edit] edit protocols bgp set group <i>group-name</i> type <i>type</i> neighbor <i>neighbor-address</i> set group <i>group-name</i> local-address <i>local-address</i> show commit</pre> |
| c. Configure a Simple Routing Policy on page 24 | <pre>[edit] edit routing-options set static route <i>destination/24</i> reject [edit policy-options] set policy-statement <i>policy-name</i> term <i>term-name</i> from route-filter <i>address</i> exact set policy-statement <i>policy-name</i> term <i>term-name</i> then accept [edit protocols bgp] set export <i>policy-name</i> show commit</pre> |
| d. Verify That BGP Sessions Are Up on page 26 | show bgp summary |
| 5. Enable MPLS and RSVP on page 28 | |
| a. Enable MPLS and RSVP on Routers on page 28 | <pre>[edit] edit protocols set mpls interface all set rsvp interface all [edit protocols mpls] set interface fxp0.0 disable [edit protocols rsvp] set interface fxp0.0 disable show commit</pre> |

| Verifying a Network Configured with MPLS Tasks | Command or Action |
|---|---|
| b. Enable MPLS on Transit Interfaces on page 29 | [edit] edit interfaces set <i>type-fpc/pic/port</i> unit <i>number</i> family mpls show commit |
| 6. Establish an LSP in Your Network on page 30 | |
| a. Configure the LSP on page 30 | [edit] edit protocols mpls set label-switched-path <i>lsp-path-name</i> to <i>address</i> show commit |
| b. Verify the LSP on page 32 | show mpls lsp extensive |
| Example Configurations for an MPLS Topology on page 34 | show configuration no-more |

Configuring MPLS on Your Network

Purpose For MPLS to run on the routers in your network, you must enable MPLS and the Resource Reservation Protocol (RSVP), configure an interior gateway protocol (IGP) and Border Gateway Protocol (BGP) to run over the relevant interfaces, and configure each interface with the following:

- Basic IP information
- MPLS support

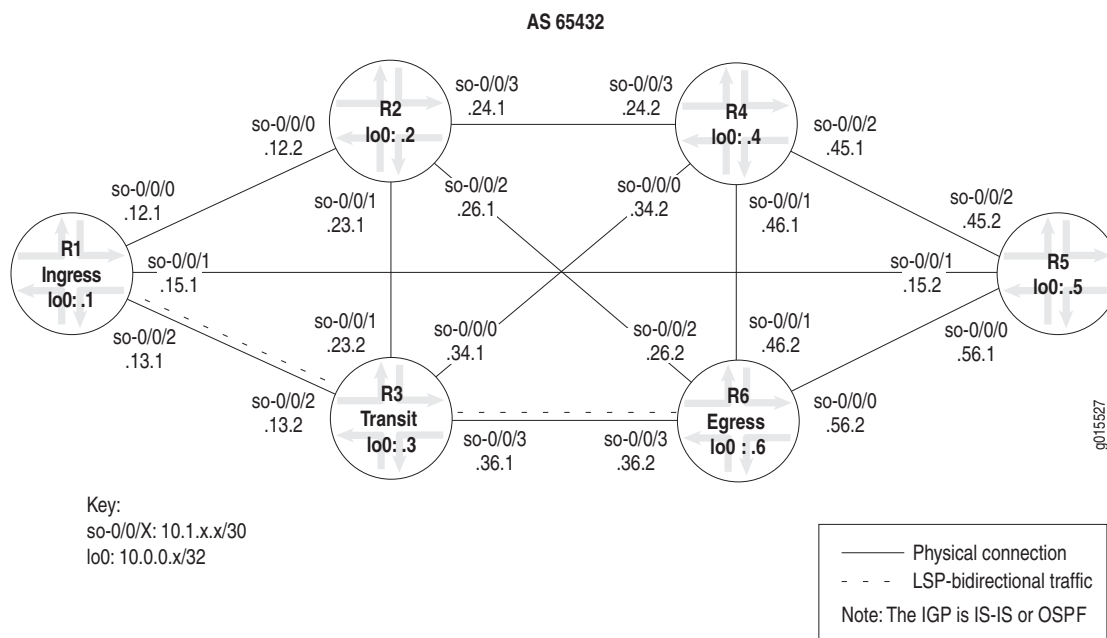
In addition, you must configure a label-switched path (LSP) from the ingress router to the egress router. For more information on ingress and egress routers, see the *JUNOS MPLS Applications Configuration Guide*.

You can configure your MPLS network with either Intermediate System-to-Intermediate System (IS-IS) or Open Shortest Path First (OSPF) as the IGP. The example network in Figure 1 is configured with IS-IS. To configure interfaces with OSPF, see the *JUNOS Routing Protocols Configuration Guide*.

An IGP is required for the Constrained Shortest Path First (CSPF) LSP, which is the default with the JUNOS software. The example network in Figure 1 focuses on CSPF LSPs.

Figure 1 illustrates the example MPLS network topology used in this section and throughout this book. The example network uses IS-IS Level 2 and a policy to create traffic. However, IS-IS Level 1 or an OSPF area can be used and the policy omitted if the network has existing BGP traffic.

Figure 1: MPLS Network Topology



The MPLS network in Figure 1 on page 6 illustrates a router-only network with SONET interfaces that consist of the following components:

- A full-mesh interior BGP (IBGP) topology, using AS 65432
- MPLS and RSVP enabled on all routers
- A `send-statics` policy on routers R1 and R6 that allow a new route to be advertised into the network
- Two unidirectional LSPs between routers R1 and R6, which allow for bidirectional traffic

The network shown in Figure 1 is a 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.

See “Example Configurations for an MPLS Topology” on page 34 for complete configurations for all routers in this example MPLS network. The following sections outline the steps for configuring MPLS on a network based on the topology shown in Figure 1.

You can enable MPLS throughout the rest of the network by repeating Step 1, “Configure IP Addresses on Router Interfaces” on page 8 through Step 5, “Enable MPLS and RSVP” on page 28 as appropriate on other routers until all routers and interfaces are enabled for MPLS.

Steps To Take To configure the MPLS network, follow these steps:

1. Configure IP Addresses on Router Interfaces on page 8
2. Configure IS-IS as the IGP on page 9
3. Configure OSPF as the IGP on page 16
4. Set Up BGP on Routers in Your Network on page 21
5. Enable MPLS and RSVP on page 28
6. Establish an LSP in Your Network on page 30

Step 1: Configure IP Addresses on Router Interfaces

Purpose Before you can run MPLS on your network, you must have an IP address configured on all interfaces. Repeat this procedure as appropriate on other router interfaces in your network until all interfaces have an IP address.

Action To configure an IP address, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```
[edit]
user@host# edit interfaces type-fpc/pic/port unit logical-unit-number
```

2. Configure the IP address:

```
[edit interfaces type-fpc/pic/port unit number]
user@host# set family inet address address
```

3. Verify the configuration:

```
user@host# show
user@host# commit
```

Sample Output

```
user@R1> edit
Entering configuration mode

[edit]
user@R1# edit interfaces so-0/0/2 unit 0

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

[edit interfaces so-0/0/2 unit 0]
user@R1# show
family inet {
    address 10.1.13.1/30;
}

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

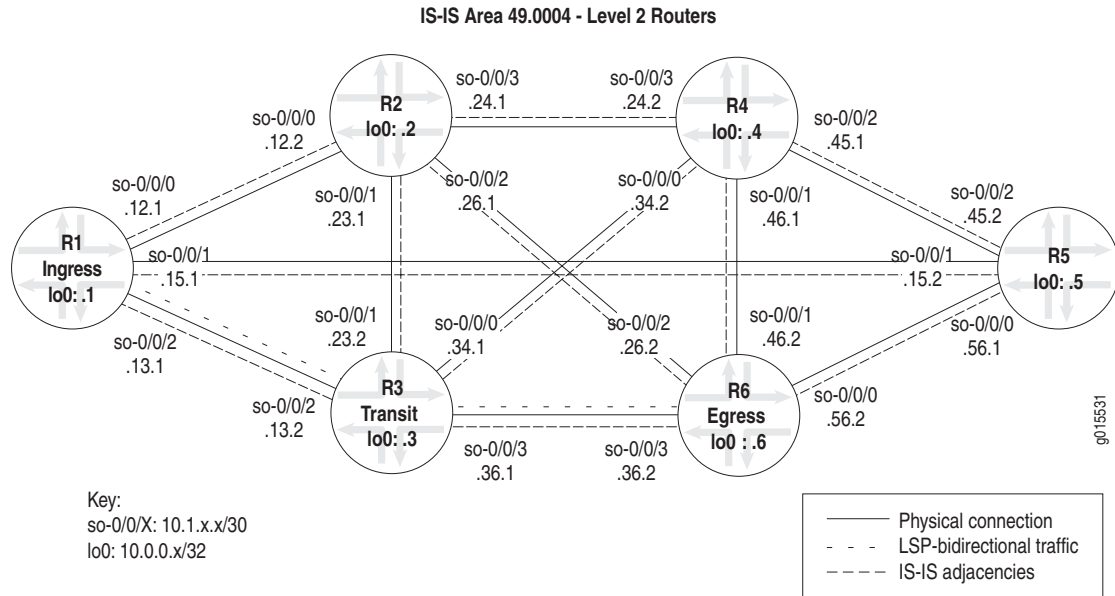
What It Means The sample output shows an interface configured with an IP address. The IP address is assigned when you configure the protocol family. In this instance, the IP address is included in the **inet** family. The **family** statement identifies which protocol packets are accepted into the interfaces. For example, valid IP packets are dropped if the interface is not configured with the **family inet** statement.

For more information on interface addressing, see the *JUNOS Network Interfaces Configuration Guide*.

Step 2: Configure IS-IS as the IGP

Purpose Before you can run MPLS on your network, you should have an IGP running on all specified routers and interfaces. The IGP can be either IS-IS or OSPF. For the steps to configure OSPF, see “Configure OSPF as the IGP” on page 16.

Figure 2: IS-IS Network Topology



The IS-IS IGP in the MPLS network in Figure 2 consists of the following:

- All routers are configured for Level 2, therefore default CSPF LSPs can occur.
- All routers are in IS-IS area 49.0004. However, the routers in this network could be in any area because Level 2 adjacencies occur between all directly connected Level 2 routers regardless of which area they are in.
- Level 2 adjacencies between all directly connected Level 2 routers as follows:
 - R1 is adjacent to R2, R3, and R5
 - R2 is adjacent to R1, R3, R4, and R6
 - R3 is adjacent to R1, R2, R4, and R6
 - R4 is adjacent to R2, R3, R5, and R6
 - R5 is adjacent to R1, R4, and R6
 - R6 is adjacent to R2, R3, R4, and R5

When you configure IS-IS as the IGP, you must enable IS-IS on the router, configure International Organization for Standardization (ISO) addressing, and enable IS-IS on all router interfaces.

You can enable IS-IS throughout the rest of the network by repeating Step 1, “Enable IS-IS on Routers in Your Network” on page 10 through Step 3, “Enable IS-IS on Router Interfaces” on page 14 as appropriate on other routers until all routers and interfaces establish IS-IS adjacencies.

Steps To Take To configure IS-IS and establish IS-IS adjacencies, follow these steps:

1. Enable IS-IS on Routers in Your Network on page 10
2. Configure ISO Addressing on page 12
3. Enable IS-IS on Router Interfaces on page 14
4. Verify That IS-IS Adjacencies Are Established on page 15

1. Enable IS-IS on Routers in Your Network

Action To enable IS-IS on routers in your network, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```
[edit]
user@host# edit protocols isis
```

2. Disable Level 1 if appropriate for your network:

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

3. Configure the interface:

```
[edit protocols isis]
user@host# edit interface type-fpc/pic/port level level-number metric metric
```

4. Disable the management interface if you have included the **interface all** statement, as shown in “Sample Output 2” on page 11:

```
[edit protocols isis]
user@host# set interface fxp0.0 disable
```

5. Include the loopback interface (lo0) if you have listed all interfaces separately, as shown in “Sample Output 1” on page 11:

```
[edit protocols isis]
user@host# set interface lo0.0
```

6. Set the loopback interface (lo0) to passive:

```
[edit protocols isis]
user@R1# set interface lo0 passive
```


7. Verify and commit the configuration:

```
user@host# show
user@host# commit
```

Sample Output 1

```
user@R1> edit
Entering configuration mode

[edit]
user@R1# edit protocols isis

[edit protocols isis]
user@R1# set level 1 disable

[edit protocols isis]
user@host# edit interface all level 2 metric 10

[edit protocols isis]
user@host# set interface lo0.0

[edit protocols isis]
user@host# set interface lo0 passive

[edit protocols isis]
user@R1# show
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;
}

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

Sample Output 2

```
[edit protocols isis]
user@R6# show
level 1 disable;
interface all {
    level 2 metric 15;
}
interface fxp0.0 {
    disable;
}
interface lo0.0 {
    passive;
}
```

What It Means Sample Output 1 shows that IS-IS Level 1 is disabled, making this a Level 2 router. All routers in the network shown in Figure 1 on page 6 are running at one IS-IS level (Level 2), therefore default CSPF LSPs can occur.

Because R1 in Sample Output 1 has all IS-IS enabled interfaces listed, including the loopback interface (lo0), you do not need to include the **disable** statement for the management interface (fxp0). All interfaces have unit number 0, the default if a unit number is not specified. When you configure an interface at the [edit protocols isis] hierarchy level, and you do not include the logical unit, the default 0 is appended to the interface name, for example, so-0/0/1.0.

Sample Output 2 does not list the interfaces configured with IS-IS; instead, all interfaces are configured, including the loopback interface (lo0) and the management interface (fxp0). Therefore, you do not need to include a separate statement for the loopback (lo0) interface. However, in this instance, it is best practice to disable the management interface (fxp0) so that IS-IS packets are not sent over it. If you do not disable the management interface (fxp0) when you include the **interface-all** statement, the IS-IS protocol can form adjacencies over the management backbone, but traffic does not flow because transit traffic does not go out of the management interface.

Sample Output 2 also shows that all interfaces on R6 are configured with a metric of 15. A metric is not required to configure IS-IS on your interfaces. The default metric value is 10 (with the exception of the loopback [lo0] interface, which has a default metric of 0). A metric is included to demonstrate that you can configure a metric for IS-IS if the default (10) is not appropriate for your network.

Both sample outputs show the **passive** statement included in the configuration of the loopback (lo0) interface. Including the **passive** statement is considered best practice and ensures the following:

- Protocols are not run over the loopback (lo0) interface
- When the router ID (RID) is configured manually, ensures that the loopback (lo0) interface is advertised to other networks.



NOTE: It is considered best practice to configure the RID manually to avoid duplicate RID problems.

2. Configure ISO Addressing

Purpose For a router to support IS-IS, you must configure an ISO network entity title (NET) address on one of the router's interfaces, preferably the loopback interface (lo0).

Action To configure ISO addressing, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```
[edit]
user@host# edit interfaces
```

2. Include a NET address for the loopback interface:

```
[edit interfaces]
user@host# set lo0 unit number family iso address address
```

3. Verify and commit the configuration:

```
user@host# show
user@host# commit
```

Sample Output

```

user@R1> edit
Entering configuration mode

edit]
user@R1# edit interfaces

[edit interfaces]
user@R1# set lo0 unit 0 family iso address 49.0004.1000.0000.0001.00

[edit interfaces]
user@R1# show
[...Output truncated...]
lo0 {
    unit 0 {
        family inet {
            address 10.0.0.1/32;
        }
        family iso {
            address 49.0004.1000.0000.0001.00;
        }
    }
}

[edit interfaces]
user@R1# commit
commit complete

```

What It Means The sample output shows that the loopback (lo0) interface is configured with the NET address 49.0004.1000.0000.0001.00. The loopback interface (lo0) becomes a point of connection from the router to the IS-IS network. Every router in an IS-IS network must have at least one ISO NET address that identifies a point of connection to the IS-IS network. The NET address is generally configured on the loopback (lo0) interface. Routers that participate in multiple areas can have multiple NET addresses.

All the routers in the network shown in Figure 1 on page 6 share a Level 2 database containing identical information. A common Level 2 database occurs in this case because all adjacencies are Level 2, and all routers are within the same IS-IS area (49.0004). Level 2 LSP flooding reaches all routers in the network due to the presence of a single level. For more information on determining the NET address, see the *JUNOS Routing Protocols Configuration Guide*.

3. Enable IS-IS on Router Interfaces

Purpose Enable reception and transmission of ISO protocol data units (PDUs) on each router interface in the network with the **family** statement, which identifies which protocol packets are accepted into the interfaces. For example, valid IS-IS packets are dropped if the interface is not configured with the **family iso** statement.

Action To configure support for IS-IS on router interfaces in your network, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```
[edit]
user@host# edit interfaces
```

2. Configure IS-IS:

```
[edit interfaces]
user@host# set type-fpc/pic/port unit number family iso
```

3. Verify and commit the configuration:

```
user@host# show
user@host# commit
```

Sample Output

```
user@R1> edit
Entering configuration mode

edit]
user@R1# edit interfaces

[edit interfaces]
user@R1# set so-0/0/2 unit 0 family iso

[edit interfaces]
userR1# show
[...Output truncated...]
so-0/0/2 {
    unit 0 {
        family inet {
            address 10.1.13.1/30;
        }
        family iso;
    }
}

[edit interfaces]
user@R1# commit
commit complete
```

What It Means The sample output shows that the interface **so-0/0/2** is configured with IS-IS.

4. Verify That IS-IS Adjacencies Are Established

Purpose After configuring IS-IS, you must verify that neighboring routers have formed adjacencies with each other.

Action To verify IS-IS adjacencies, enter the following JUNOS command-line interface (CLI) operational mode command:

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

Sample Output

```
user@R1> show isis adjacency
Interface          System      L State      Hold (secs) SNPA
so-0/0/0.0         R2          2 Up         25
so-0/0/1.0         R5          2 Up         23
so-0/0/2.0         R3          2 Up         20

user@R3> show isis adjacency
Interface          System      L State      Hold (secs) SNPA
so-0/0/0.0         R4          2 Up         25
so-0/0/1.0         R2          2 Up         25
so-0/0/2.0         R1          2 Up         26
so-0/0/3.0         R6          2 Up         25

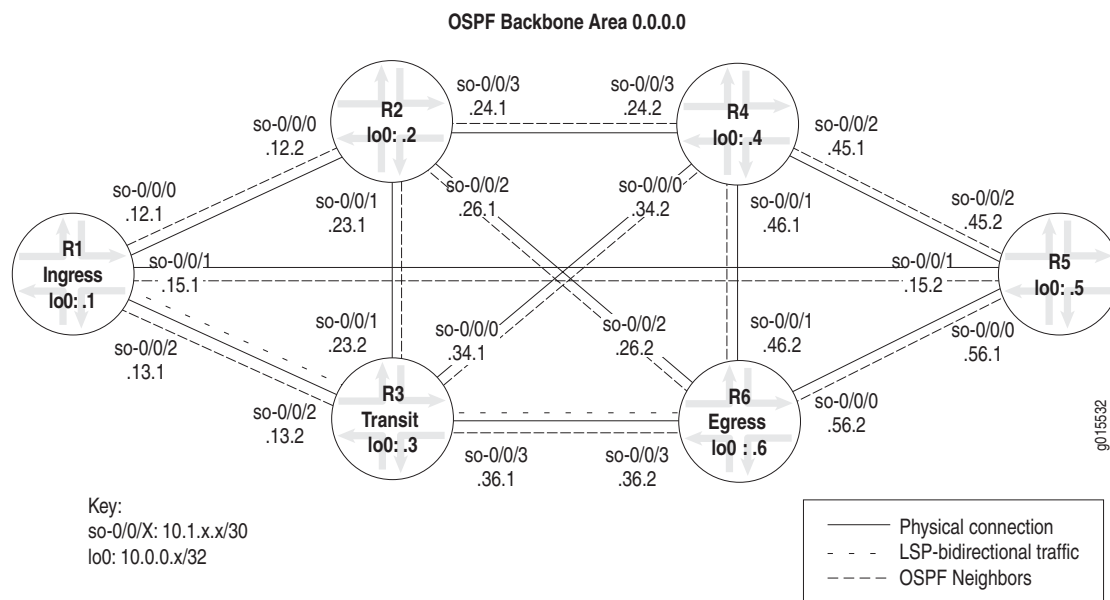
user@R6> show isis adjacency
Interface          System      L State      Hold (secs) SNPA
so-0/0/0.0         R5          2 Up         19
so-0/0/1.0         R4          2 Up         22
so-0/0/2.0         R2          2 Up         22
so-0/0/3.0         R3          2 Up         19
```

What It Means The sample output from the ingress, transit, and egress routers shows that all routers in the network shown in Figure 1 on page 6 have formed IS-IS adjacencies.

Step 3: Configure OSPF as the IGP

Purpose Before you can run MPLS on your network, you must have an IGP running on all specified routers and interfaces. The IGP can be either OSPF or IS-IS. For the steps to configure IS-IS, see “Configure IS-IS as the IGP” on page 9.

Figure 3: OSPF Network Topology



The OSPF IGP in the MPLS network in Figure 2 consists of the following:

- All routers are configured for the backbone OSPF area 0.0.0.0.
- All routers have the RID manually configured to avoid possible problems when the OSPF RID changes; for example, when multiple loopback addresses are configured.
- All routers have traffic engineering enabled. When traffic engineering is enabled for OSPF, the SPF algorithm takes into account the various LSPs configured under MPLS and configures OSPF to generate link-state advertisements (LSAs) that carry traffic engineering parameters. These routes are installed into the primary routing table `inet.0`, but the LSPs are installed by default into the `inet.3` routing table.
- Adjacencies between all OSPF neighbors are as follows:
 - R1 is adjacent to R2, R3, and R5
 - R2 is adjacent to R1, R3, R4, and R6
 - R3 is adjacent to R1, R2, R4, and R6
 - R4 is adjacent to R2, R3, R5, and R6

- R5 is adjacent to R1, R4, and R6
- R6 is adjacent to R2, R3, R4, and R5

When you configure OSPF as the IGP, you must enable OSPF and traffic engineering on the router. We also recommend that you manually configure the RID and include the loopback interface (lo0) at the `[edit protocols ospf]` hierarchy level.

You can enable OSPF throughout the rest of the network by repeating this step as appropriate on other routers until all routers and interfaces establish OSPF neighbors.

Steps To Take To configure OSPF and establish OSPF neighbors, follow these steps:

1. Enable OSPF on Routers in Your Network on page 17
2. Verify That OSPF Neighbors Are Established on page 19

1. Enable OSPF on Routers in Your Network

Action To enable OSPF on routers in your MPLS network, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```
[edit]
user@host# edit protocols ospf
```

2. Configure the area and the interface:

```
[edit protocols ospf]
user@host# set area area-id interface type-fpc/pic/port
```

3. Disable the management interface if you have included the `interface all` statement in the previous step:

```
[edit protocols ospf]
user@host# set interface fxp0.0 disable
```

4. Include the loopback (lo0) interface if you intend to manually configure the RID:

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

5. Set the loopback interface (lo0) to passive:

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

6. Configure traffic engineering:

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

- Manually configure the RID at the [routing-options] hierarchy level:

```
[edit]
user@host# edit routing-options

[edit routing-options]
user@host# set router-id router-id
```

- Verify and commit the entire configuration:

```
user@host# show
user@host# commit
```

Sample Output

```
user@R6> edit
Entering configuration mode

[edit]
user@R6# edit protocols ospf

[edit protocols ospf]
user@R6# set area 0.0.0.0 interface so-0/0/0.0

[edit protocols ospf]
user@R6# set area 0.0.0.0 interface lo0

[edit protocols ospf]
user@R6# set area 0.0.0.0 interface lo0 passive

[edit protocols ospf]
user@R6# set traffic-engineering

[edit protocols ospf]
user@R6# 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 so-0/0/3.0;
    interface lo0.0 {
        passive;
    }
}

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

[edit]
user@R6# edit routing-options

[edit routing-options]
user@R6# set router-id 10.0.0.6
```



```
[edit routing-options]
user@R6# show
[...Output truncated...]
router-id 10.0.0.6;
autonomous-system 65432;

[edit routing-options]
user@R6# commit
commit complete
```

What It Means The sample output shows that OSPF, with traffic engineering, is enabled on the interfaces on egress router R6. In addition, the RID is configured manually to avoid possible problems when the OSPF RID changes; for example, when multiple loopback addresses are configured. The RID uniquely identifies the router within the OSPF network. It is transmitted within the LSAs used to populate the link-state database and calculate the shortest-path tree. In a link-state network, it is important that two routers do not share the same RID value, otherwise IP routing problems may occur.

The sample outputs also shows the **passive** statement included in the configuration of the loopback (lo0) interface. Including the **passive** statement is considered best practice and ensures the following:

- Protocols are not run over the loopback (lo0) interface
- When the router ID (RID) is configured manually, ensures that the loopback (lo0) interface is advertised to other networks.

2. Verify That OSPF Neighbors Are Established

Purpose After configuring OSPF, you must verify that neighboring routers have formed adjacencies with each other.

Action To verify OSPF neighbors, enter the following JUNOS CLI operational mode command:

```
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 | 37 |
| 10.1.15.2 | so-0/0/1.0 | Full | 10.0.0.5 | 128 | 35 |
| 10.1.13.2 | so-0/0/2.0 | Full | 10.0.0.3 | 128 | 38 |

```
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 | 38 |
| 10.1.23.1 | so-0/0/1.0 | Full | 10.0.0.2 | 128 | 35 |
| 10.1.13.1 | so-0/0/2.0 | Full | 10.0.0.1 | 128 | 37 |
| 10.1.36.2 | so-0/0/3.0 | Full | 10.0.0.6 | 128 | 36 |

```
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 | 39 |
| 10.1.46.1 | so-0/0/1.0 | Full | 10.0.0.4 | 128 | 37 |
| 10.1.26.1 | so-0/0/2.0 | Full | 10.0.0.2 | 128 | 36 |
| 10.1.36.1 | so-0/0/3.0 | Full | 10.0.0.3 | 128 | 37 |


What It Means The sample output from the ingress, transit, and egress routers shows that all routers in the network shown in Figure 1 on page 6 have formed OSPF neighbor adjacencies.

Step 4: Set Up BGP on Routers in Your Network

Purpose Before BGP can function in your MPLS network, you must define the autonomous system (AS) number on the routers in your network, and configure at least one group that includes at least one peer.

Optionally, you can configure a routing policy. The routing policy allows you to control the information shared with BGP neighbors and provides the opportunity to filter and modify the information you receive.

Figure 4: BGP Network Topology



The BGP configuration in the MPLS network in Figure 4 consists of the following:

- A full-mesh IBGP topology, using AS 65432.
- All IBGP sessions peer between loopback addresses because significant stability advantages are gained.
- All routers are configured with one group, **group internal**.
- A **send-statics** policy on routers R1 and R6 allows a new route to be advertised into the network.

The example network uses IS-IS Level 2 and a policy to create routes that are reachable through the LSP. However, IS-IS Level 1 or an OSPF area can be used and the policy omitted if the network has existing BGP traffic.

You can set up BGP throughout the rest of the network by repeating Step 1, “Define the Local Autonomous System” on page 22 through Step 3, “Configure a Simple Routing Policy” on page 24 as appropriate on other routers until all routers are set up with BGP.

Steps to Take To set up BGP on routers in your network, follow these steps:

1. Define the Local Autonomous System on page 22
2. Configure BGP Neighbor Connections on page 23
3. Configure a Simple Routing Policy on page 24
4. Verify That BGP Sessions Are Up on page 26

1. Define the Local Autonomous System

Purpose Before BGP can function, you need to define a local AS number on the routers in your network. In the example network in Figure 4 on page 21, all routers are in AS 65432.

Action To define an AS number on routers in your network, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```
[edit]
user@host# edit routing-options
```

2. Configure all interfaces to a specific AS:

```
[edit routing-options]
user@host# set autonomous-system as-number
```

3. Verify the configuration:

```
user@host# show
user@host# commit
```

Sample Output

```
user@R1> edit
Entering configuration mode

[edit]
user@R1# edit routing-options

[edit routing-options]
user@R1# set autonomous-system 65432
```

```
[edit routing-options]
user@R1# show
[...Output truncated...]
autonomous-system 65432;

[edit routing-options]
user@R6# commit
commit complete
```

What It Means The output shows that router R1 resides in AS 65432. All other routers in the example network shown in Figure 4 on page 21 also reside in AS 65432.

2. Configure BGP Neighbor Connections

Purpose You must configure at least one group that includes at least one peer for BGP to run in your network. First determine which neighbors are internal or external to your local AS boundary. Internal neighbors are inside your local AS boundary. In the example network shown in Figure 4 on page 21, all the routers are in one AS and are therefore internal. In this example, all IBGP sessions peer between loopback addresses because significant stability advantages are gained. For more information about configuring BGP neighbor connections, see the *JUNOS Routing Protocols Configuration Guide*.

Action To configure BGP neighbor connections, follow these steps:

1. In configuration mode, go to the following hierarchy level:

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

2. Configure the group and peer's IP address:

```
[edit protocols bgp]
user@host# set group group-name type type neighbor neighbor-address
```



NOTE: For external neighbors, use the following form of the command that includes the peer's AS number:

```
user@host# set group group-name neighbor neighbor-address peer-as
peer-as-number
```

3. Configure the local address:

```
[edit protocols bgp]
user@host# set group group-name local-address local-address
```

4. Verify and commit the configuration:

```
user@host# show
user@host# commit
```

Sample Output

```
user@R1> edit
Entering configuration mode

[edit]
user@R1# edit protocols bgp

[edit protocols bgp]
user@R1# set group internal type internal neighbor 10.0.0.2

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

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

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

What It Means The sample output shows that router R1 is in an internal group with five BGP neighbors. The `local-address` statement is included in this example configuration because IBGP is used. It is considered best practice to configure a local address when you use an IBGP. BGP messages are sourced from the loopback address because the `local-address` statement is included in the configuration. Generally, you would not configure a local address when external BGP is configured.

3. Configure a Simple Routing Policy

Purpose Routing policy allows you to control the information shared with BGP neighbors and provides the opportunity to filter and modify the information you receive. Typically, a network is injected into BGP using a policy. This may also be done through a static route. In the network in Figure 4 on page 21, a static route export policy is used to inject routes into BGP.

Action To configure a simple routing policy, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```
[edit]
user@host# edit routing-options
```

2. Configure a static route for redistribution to other autonomous systems:

```
[edit routing-options]
user@host# set static route destination/24 reject
```

3. Configure a routing policy that matches and accepts the configured static routes into BGP updates:

```
[edit]
user@host# edit policy-options

[edit policy-options]
user@host# set policy-statement policy-name term term-name from
route-filter address exact
user@host# set policy-statement policy-name term term-name then accept
```

4. Apply the policy created in Step 3 to all BGP neighbors:

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

[edit protocols bgp]
user@host# set export policy-name
```

5. Verify and commit the configuration:

```
user@host# show
user@host# commit
```

Sample Output

```
user@R1> edit
Entering configuration mode

[edit]
user@R1# edit routing-options

[edit routing-options]
user@R1# set static route 100.100.1.0/24 reject

[edit routing-options]
user@R1# show
[...Output truncated...]
    route 100.100.1.0/24 reject;
}
router-id 10.0.0.1;
autonomous-system 65432;

[edit routing-options]
user@R1# top

[edit]
user@R1# edit policy-options

[edit policy-options]
user@R1# set policy-statement send-statics term statics from route-filter
100.100.1.0/24 exact
```

```

[edit policy-options]
user@R1# set policy-statement send-statics term statics then accept

[edit policy-options]
user@R1# top

[edit]
user@R1# edit protocols bgp

[edit protocols bgp]
user@R1# set export send-statics

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

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

```

What It Means The sample output shows that routing policy **send-statics** is configured on the router. The routing policy matches and accepts the configured static routes into the routing table and injects the routes into BGP updates. Typically, a routing policy is applied at the group level, although it can be applied at the global level, as shown in this example.

4. Verify That BGP Sessions Are Up

Purpose After configuring BGP, you must verify that BGP peers are established and the sessions are up.

Action To verify BGP peers and sessions, enter the following JUNOS CLI operational mode command:

```
user@host> show bgp summary
```

Sample Output

```

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 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn State|#Active/Received/Damped...
10.0.0.2 65432 1369 1373 0 0 11:25:11 0/0/0 0/0/0
10.0.0.3 65432 1369 1372 0 0 11:24:55 0/0/0 0/0/0
10.0.0.4 65432 1369 1372 0 0 11:25:03 0/0/0 0/0/0
10.0.0.5 65432 1369 1372 0 0 11:25:07 0/0/0 0/0/0
10.0.0.6 65432 1343 1344 0 1 11:10:55 1/1/0 0/0/0

```



```
user@R3> show bgp summary
```

```
Groups: 1 Peers: 4 Down peers: 0
```

| Table | Tot Paths | Act Paths | Suppressed | History | Damp | State | Pending |
|----------|-----------|-----------|------------|---------|-------|-------------|----------------------------------|
| inet.0 | 2 | 2 | 0 | 0 | | 0 | 0 |
| Peer | AS | InPkt | OutPkt | OutQ | Flaps | Last Up/Dwn | State #Active/Received/Damped... |
| 10.0.0.1 | 65432 | 1375 | 1375 | 0 | 6 | 11:26:57 | 1/1/0 0/0/0 |
| 10.0.0.2 | 65432 | 43016 | 43016 | 0 | 0 | 2w0d22h | 0/0/0 0/0/0 |
| 10.0.0.4 | 65432 | 74460 | 74461 | 0 | 0 | 3w4d20h | 0/0/0 0/0/0 |
| 10.0.0.6 | 65432 | 1347 | 1347 | 0 | 6 | 11:13:10 | 1/1/0 0/0/0 |

```
user@R6> 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... |
| 10.0.0.1 | 65432 | 1348 | 1350 | 0 | 0 | 11:13:46 | 1/1/0 0/0/0 |
| 10.0.0.2 | 65432 | 1347 | 1351 | 0 | 0 | 11:14:02 | 0/0/0 0/0/0 |
| 10.0.0.3 | 65432 | 1347 | 1350 | 0 | 0 | 11:13:58 | 0/0/0 0/0/0 |
| 10.0.0.4 | 65432 | 1347 | 1350 | 0 | 0 | 11:13:54 | 0/0/0 0/0/0 |
| 10.0.0.5 | 65432 | 1347 | 1350 | 0 | 0 | 11:13:50 | 0/0/0 0/0/0 |

What It Means The sample output from the ingress, transit, and egress routers shows that all routers in the network shown in Figure 4 on page 21 have BGP peers established and sessions up.

Step 5: Enable MPLS and RSVP

Purpose You can enable MPLS and RSVP throughout the rest of the network by repeating Step 1, “Enable MPLS and RSVP on Routers” on page 28 and Step 2, “Enable MPLS on Transit Interfaces” on page 29 as appropriate on other routers until all routers are enabled with MPLS and RSVP.



NOTE: Even though the MPLS and RSVP protocols are enabled, you must complete all five steps in this chapter to have the MPLS protocol running on your network.

- Steps To Take**
1. Enable MPLS and RSVP on Routers on page 28
 2. Enable MPLS on Transit Interfaces on page 29

1. Enable MPLS and RSVP on Routers

Action To enable MPLS and RSVP on routers in your network, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```
[edit]
user@host# edit protocols
```

2. Configure MPLS and RSVP:

```
[edit protocols]
user@host# set mpls interface all
user@host# set rsvp interface all
```

3. Disable the management interface for MPLS and RSVP:

```
[edit protocols mpls]
user@host# set interface fxp0.0 disable
```

```
[edit protocols rsvp]
user@host# set interface fxp0.0 disable
```

4. Verify and commit the configuration:

```
user@host# show
user@host# commit
```

Sample Output

```
user@R1> edit
Entering configuration mode

[edit]
user@R1# edit protocols

[edit protocols]
user@R1# set mpls interface all

[edit protocols]
user@R1# set rsvp interface all
```

```
[edit protocols]
user@R1# show
rsvp {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
mpls {
  interface all;
  interface fxp0.0 {
    disable;
  }
}

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

What It Means The sample output shows that router R1 has MPLS and RSVP enabled on all interfaces, except for the management interface (fxp0.0), which is disabled. It is considered best practice to disable the management interlace (fxp0.0) for MPLS and RSVP to preempt any problems. The sample network shown in Figure 1 on page 6 has all interfaces (with the management interface [fxp0.0]) disabled on all routers configured with the MPLS and RSVP protocols.

Typically every interface that you want to use is listed. For an example of a router configured with specific interfaces, see “Enable IS-IS on Routers in Your Network” on page 10.

2. Enable MPLS on Transit Interfaces

Purpose Even though transit interfaces are enabled with MPLS when you include the family mpls statement in the configuration, MPLS as a whole is not configured on your router or in your network. You must complete all five steps in this chapter to have the MPLS protocol running on your network.



NOTE: The management interface (fxp0) and the loopback interface (lo0) are not transit interfaces.

Action To configure transit interfaces to support MPLS, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```
[edit]
user@host# edit interfaces
```

2. Configure MPLS:

```
[edit interfaces]
user@host# set type-fpc/pic/port unit number family mpls
```

3. Verify and commit the configuration:

```
user@host# show
user@host# commit
```

Sample Output

```

user@R1> edit
Entering configuration mode

[edit]
user@R1# edit interfaces

[edit interfaces]
user@R1# set so-0/0/2 unit 0 family mpls

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

[edit interfaces]
user@R1# commit
commit complete

```

What It Means The sample output shows that the interface `so-0/0/2` is configured to support MPLS. The family statement identifies which protocol packets are accepted into the interfaces. For example, valid MPLS packets are dropped if the interface is not configured with the MPLS protocol.

Step 6: Establish an LSP in Your Network

Purpose Create a label-switched path on specified routers in your network using the loopback address of the ingress and egress routers.

Steps To Take To establish an LSP in your network, follow these steps:

1. Configure the LSP on page 30.
2. Verify the LSP on page 32.

1. Configure the LSP

Action To configure an LSP in your network, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```

[edit]
user@host# edit protocols mpls

```

2. Configure the LSP on the ingress and egress routers:

```

[edit protocols mpls]
user@host# set label-switched-path lsp-path-name to address

```

3. Verify and commit the configuration:

```
user@host# show
user@host# commit
```

Sample Output 1

```
user@R1> edit
Entering configuration mode

[edit]
user@R1# edit protocols mpls

[edit protocols mpls]
user@R1# set label-switched-path R1-to-R6 to 10.0.0.6

[edit protocols mpls]
user@R1# show
label-switched-path R1-to-R6 {
    to 10.0.0.6;
}
interface all;
interface fxp0.0 {
    disable;
}

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

Sample Output 2

```
[edit protocols mpls]
user@R6# show
label-switched-path R6-to-R1 {
    to 10.0.0.1;
}
interface all;
interface fxp0.0 {
    disable;
}
```

What It Means The sample output shows that two CSPF LSPs (R1-to-R6 and R6-to-R1) are configured between routers R1 and R6. CSPF is enabled by default with the JUNOS software. The example network shown in Figure 1 on page 6 focuses on CSPF LSPs.

The CSPF algorithm is an advanced form of the SPF algorithm used in OSPF and IS-IS route computations. CSPF is used in computing paths for LSPs that are subject to multiple constraints. When computing paths for LSPs, CSPF considers not only the topology of the network, but also the attributes of the LSP and the links, and attempts to minimize congestion by intelligently balancing the network load.

Typically in a network, LSPs are configured to every other egress router, resulting in a full mesh of LSPs that correspond to the BGP full mesh. In the example network shown in Figure 1 on page 6, two LSPs are configured between R1 and R6 to allow for bidirectional traffic. The first LSP is from R1 to R6 (R1-to-R6) and the second is from R6 to R1 (R6-to-R1). If only one LSP was configured, for example, from R1 to R6, only unidirectional traffic would be allowed.

2. Verify the LSP

Purpose After configuring the LSP, you must verify that the LSP is up. LSPs can be ingress, transit, or egress. Use the `show mpls lsp` command for quick verification of the LSP state, with the `extensive` option (`show mpls lsp extensive`) as a follow-up if the LSP is down. 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 verify that the LSP is up, enter the following command from the ingress router:

```
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
    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
        6 Dec 13 11:50:15 Selected as active path
        5 Dec 13 11:50:15 Record Route: 10.1.13.2 10.1.36.2
        4 Dec 13 11:50:15 Up
        3 Dec 13 11:50:15 Originate Call
        2 Dec 13 11:50:15 CSPF: computation result accepted
        1 Dec 13 11:49:45 CSPF failed: no route toward 10.0.0.6[6 times]
      Created: Mon Dec 13 11:47:19 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: 127, Since: Mon Dec 13 11:50:10 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 39136 protocol 0
  PATH rcvfrom: 10.1.13.2 (so-0/0/2.0) 28709 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

```

What It Means The sample output from ingress router **R1** show two LSPs in which this router participates: ingress LSP **R1-to-R6** and egress LSP **R6-to-R1** (the reverse LSP which allows bidirectional traffic). Both LSPs have active routes to the destination: **10.0.0.6** for the ingress LSP and **10.0.0.1** for the egress LSP. The state for both LSPs is up.

For more information on verifying the LSP, see “Determining the LSP State” on page 59.

Example Configurations for an MPLS Topology

Purpose The configurations in this section are for the six routers in the example network illustrated in Figure 1 on page 6.

Action To display the configuration of a router, use the following JUNOS CLI operational mode command:

```
user@host> show configuration | no-more
```

Sample Output 1 user@R1> show configuration | no-more

```
system {
  host-name R1;
  [...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;
      }
    }
  }
}
```

#family mpls is not
#configured because the
#loopback (lo0) interface is
#not a transit interface


```

routing-options {
  static {
    [...Output truncated...]
    route 100.100.1.0/24 reject;
  }
  router-id 10.0.0.1;
  autonomous-system 65432;
}
protocols {
  rsvp {
    inactive: traceoptions {
      file rsvp.log;
      flag packets;
    }
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  mpls {
    label-switched-path R1-to-R6 {
      to 10.0.0.6;
    }
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  bgp {
    export send-statics;
    group internal {
      type internal;
      local-address 10.0.0.1;
      neighbor 10.0.0.2;
      neighbor 10.0.0.3;
      neighbor 10.0.0.5;
      neighbor 10.0.0.4;
      neighbor 10.0.0.6;
    }
  }
  isis {
    level 1 disable;
    interface all {
      level 2 metric 10;
    }
    interface fxp0.0 {
      disable;
    }
    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@R2> show configuration | no-more
system {
    host-name R2;
    [...Output truncated...]
}
interfaces {
    so-0/0/0 {
        unit 0 {
            family inet {
                address 10.1.12.2/30;
            }
            family iso;
            family mpls;
        }
    }
    so-0/0/1 {
        unit 0 {
            family inet {
                address 10.1.23.1/30;
            }
            family iso;
            family mpls;
        }
    }
    so-0/0/2 {
        unit 0 {
            family inet {
                address 10.1.26.1/30;
            }
            family iso;
            family mpls;
        }
    }
    so-0/0/3 {
        unit 0 {
            family inet {
                address 10.1.24.1/30;
            }
            family iso;
            family mpls;
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.70.144/21;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 10.0.0.2/32;
            }
            family iso {
                address 49.0004.1000.0000.0002.00;
            }
        }
    }
}

```

#family mpls is not
#configured because the
#loopback (lo0) interface is
#not a transit interface

```

routing-options {
  [...Output truncated...]
  router-id 10.0.0.2;
  autonomous-system 65432;
}
protocols {
  rsvp {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  mpls {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  bgp {
    group internal {
      type internal;
      local-address 10.0.0.2;
      neighbor 10.0.0.1;
      neighbor 10.0.0.3;
      neighbor 10.0.0.4;
      neighbor 10.0.0.6;
    }
  }
  isis {
    level 1 disable;
    interface all {
      level 2 metric 10;
    }
    interface fxp0.0 {
      disable;
    }
    interface lo0.0;
    passive
  }
}

```

Sample Output 3

```

user@R3> show configuration | no-more
system {
  host-name R3;
  [...Output truncated...]
}
interfaces {
  so-0/0/0 {
    unit 0 {
      family inet {
        address 10.1.34.1/30;
      }
      family iso;
      family mpls;
    }
  }
  so-0/0/1 {
    unit 0 {
      family inet {
        address 10.1.23.2/30;
      }
      family iso;
      family mpls;
    }
  }
}

```

```

}
so-0/0/2 {
  unit 0 {
    family inet {
      address 10.1.13.2/30;
    }
    family iso;
    family mpls;
  }
}
so-0/0/3 {
  unit 0 {
    family inet {
      address 10.1.36.1/30;
    }
    family iso;
    family mpls;
  }
}
fxp0 {
  unit 0 {
    family inet {
      address 192.168.70.145/21;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.0.0.3/32;
    }
    family iso {
      address 49.0004.1000.0000.0003.00;
    }
  }
}
}
routing-options {
  static {
    [...Output truncated...]
    router-id 10.0.0.3;
    autonomous-system 65432;
  }
}
protocols {
  rsvp {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  mpls {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
}

```

#family mpls is not
#configured because the
#loopback (lo0) interface is
#not a transit interface

```

bgp {
  group internal {
    type internal;
    local-address 10.0.0.3;
    neighbor 10.0.0.1;
    neighbor 10.0.0.2;
    neighbor 10.0.0.4;
    neighbor 10.0.0.6;
  }
}
isis {
  level 1 disable;
  interface all {
    level 2 metric 10;
  }
  interface fxp0.0 {
    disable;
  }
  interface lo0.0;
  passive
}
}

```

Sample Output 4

```

user@R4> show configuration | no-more
system {
  host-name R4;
  [...Output truncated...]
}
interfaces {
  so-0/0/0 {
    unit 0 {
      family inet {
        address 10.1.34.2/30;
      }
      family iso;
      family mpls;
    }
  }
  so-0/0/1 {
    unit 0 {
      family inet {
        address 10.1.46.1/30;
      }
      family iso;
      family mpls;
    }
  }
  so-0/0/2 {
    unit 0 {
      family inet {
        address 10.1.45.1/30;
      }
      family iso;
      family mpls;
    }
  }
  so-0/0/3 {
    unit 0 {
      family inet {
        address 10.1.24.2/30;
      }
      family iso;
      family mpls;
    }
  }
}

```

```

    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.70.146/21;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 10.0.0.4/32;
            }
            family iso {
                address 49.0004.1000.0000.0004.00;
            }
        }
    }
}
routing-options {
    static {
        [...Output truncated...]
        router-id 10.0.0.4;
        autonomous-system 65432;
    }
}
protocols {
    rsvp {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
    mpls {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
    bgp {
        group internal {
            type internal;
            local-address 10.0.0.4;
            neighbor 10.0.0.2;
            neighbor 10.0.0.3;
            neighbor 10.0.0.5;
            neighbor 10.0.0.6;
        }
    }
    isis {
        level 1 disable;
        interface all {
            level 2 metric 10;
        }
        interface fxp0.0 {
            disable;
        }
        interface lo0.0;
        passive
    }
}

```

#family mpls is not
#configured because the
#loopback (lo0) interface is
#not a transit interface

Sample Output 5

```

user@R5> show configuration | no-more
system {
    host-name R5;
    [...Output truncated...]
}
interfaces {
    so-0/0/0 {
        unit 0 {
            family inet {
                address 10.1.56.1/30;
            }
            family iso;
            family mpls;
        }
    }
    so-0/0/1 {
        unit 0 {
            family inet {
                address 10.1.15.2/30;
            }
            family iso;
            family mpls;
        }
    }
    so-0/0/2 {
        unit 0 {
            family inet {
                address 10.1.45.2/30;
            }
            family iso;
            family mpls;
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.70.147/21;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 10.0.0.5/32;
            }
            family iso {
                address 49.0004.1000.0000.0005.00;
            }
        }
    }
}
routing-options {
    static {
        [...Output truncated...]
        router-id 10.0.0.5;
        autonomous-system 65432;
    }
}

```

#family mpls is not
#configured because the
#loopback (lo0) interface is
#not a transit interface

```

protocols {
  rsvp {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  mpls {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  bgp {
    group internal {
      type internal;
      local-address 10.0.0.5;
      neighbor 10.0.0.1;
      neighbor 10.0.0.4;
      neighbor 10.0.0.6;
    }
  }
  isis {
    level 1 disable;
    interface all {
      level 2 metric 10;
    }
    interface fxp0.0 {
      disable;
    }
    interface lo0.0;
    passive
  }
}

```

Sample Output 6

```

user@R6> show configuration | no-more
system {
  host-name R6;
  [...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;
    }
    family iso {
      address 49.0004.1000.0000.0006.00;
    }
  }
}
}
routing-options {
  static {
    [...Output truncated...]
    route 100.100.6.0/24 reject;
  }
  router-id 10.0.0.6;
  autonomous-system 65432;
}
protocols {
  rsvp {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  mpls {
    label-switched-path R6-to-R1 {
      to 10.0.0.1;
    }
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
}

```

#family mpls is not
#configured because the
#loopback (lo0) interface is
#not a transit interface

```

bgp {
  group internal {
    type internal;
    local-address 10.0.0.6;
    export send-statics;
    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;
  }
}
isis {
  level 1 disable;
  interface all {
    level 2 metric 10;
  }
  interface fxp0.0 {
    disable;
  }
  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 Sample Outputs 1 through 6 show the configurations of all six routers in the example network illustrated in Figure 1 on page 6. LSPs **R1-to-R6** and **R6-to-R1** are configured on **R1** and **R6**, respectively.

Two static routes, **100.100.1/24** on **R1** and **100.100.6/24** on **R6**, are configured at the `[edit routing-options static route]` hierarchy level. Both prefixes are included in the `send-statics` policy at the `[edit policy-options send statics]` hierarchy level so the routes can become BGP routes.

In addition, the RID is configured manually at the `[edit routing-options]` hierarchy level to avoid duplicate RID problems, and the `passive` statement is included at the `[edit protocols isis interface lo0]` hierarchy level to ensure that protocols are not run over the loopback (**lo0**) interface and the loopback (**lo0**) interface is advertised correctly throughout the network.