

Chapter 22

PIM Configuration Guidelines

To configure Protocol Independent Multicast (PIM), include the `pim` statement:

```
pim {
  dense-groups {
    addresses;
  }
  disable;
  graceful-restart {
    disable;
    restart-duration seconds;
  }
  import [ policy-names ];
  interface interface-name {
    disable;
    hello-interval seconds;
    mode (dense | sparse | sparse-dense);
    priority number;
    version version;
  }
  rib-group group-name;
  rp {
    auto-rp (announce | discovery | mapping);
    bootstrap-export [ policy-names ];
    bootstrap-import [ policy-names ];
    bootstrap-priority number;
    embedded-rp {
      maximum-rps limit;
      group-ranges {
        destination-mask;
      }
    }
  }
  local {
    family (inet | inet6) {
      address address;
      anycast-pim {
        rp-set {
          address address [forward-msdp-sa];
        }
        local-address address;
      }
    }
    disable;
    group-ranges {
      destination-mask;
    }
  }
}
```

```

    }
    hold-time seconds;
    priority number;
  }
}
static {
  address address {
    version version;
    group-ranges {
      destination-mask;
    }
  }
  traceoptions {
    file name <replace> <size size> <files number> <no-stamp>
      <(world-readable | no-world-readable)>;
    flag flag <flag-modifier> <disable>;
  }
}
}

```

You can include this statement at the following hierarchy levels:

- [edit protocols]
- [edit routing-instance *instance-name* protocols]
- [edit logical-routers *logical-router-name* protocols]
- [edit routing-instance *instance-name* logical-routers *logical-router-name* protocols]

By default, PIM is disabled.

This chapter includes the following PIM tasks:

- Configuring PIM Mode-Independent Interface Properties on page 196
- Configuring Other PIM Mode-Independent Properties on page 198
- Configuring PIM Dense Mode Properties on page 201
- Configuring PIM Sparse Mode Properties on page 202
- Configuring Sparse-Dense Mode Properties on page 214
- Configuring Multicast for Layer 3 VPNs on page 215
- Configuring Multicast for Virtual Routers on page 218
- Configuration Examples on page 219

Configuring PIM Mode-Independent Interface Properties

You can configure the following properties regardless of whether the PIM interface is configured in sparse, dense, or sparse-dense mode:

- Changing the PIM Version on page 197

- Configuring the Designated Router Priority on page 197
- Modifying the Hello Interval on page 197
- Disabling the PIM Interface on page 198

For information about configuring the PIM mode on an interface, see “Configuring PIM Dense Mode Properties” on page 201, “Configuring PIM Sparse Mode Properties” on page 202, and “Configuring Sparse-Dense Mode Properties” on page 214.

Changing the PIM Version

All systems on a subnet must run the same version of PIM.

By default, the JUNOS software uses PIM version 2. To configure PIM version 1, include the `version` statement:

```
version 1;
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.



NOTE: The default PIM version can be version 1 or version 2, depending on the mode you are configuring. PIM version 1 is the default for rendezvous point (RP) mode (at the [`pim rp static address address`] hierarchy level). However, PIM version 2 is the default for interface mode (at the [`pim interface interface-name`] hierarchy level). Explicitly configured versions override the defaults.

Configuring the Designated Router Priority

By default, a PIM interface has the lowest of being selected as the designated router (DR). To change this, include the `priority` statement:

```
priority number;
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

The default priority is 1. Use a larger number to increase the of the interface’s being elected as the DR.

Modifying the Hello Interval

Routers send hello messages at a fixed interval on all PIM-enabled interfaces. Using hello messages, routers advertise their existence as a PIM router on the subnet. With all PIM-enabled routers advertised, a single DR for the subnet is established.

When a router is configured for PIM, it sends out a hello message at a 30-second default interval. The interval range is from 0 through 255. When the interval counts down to 0, it sends out another hello message, and the timer is reset. A router that gets no response from a neighbor in 3.5 times the interval value drops the neighbor. In the case of a 30-second interval, the amount of time a router would wait for a response is 105 seconds.

To modify how often the router sends hello messages out of an interface, include the `hello-interval` statement:

```
hello-interval seconds;
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Disabling the PIM Interface

To disable PIM on an interface, include the `disable` statement:

```
disable;
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Configuring Other PIM Mode-Independent Properties

You can configure the following properties regardless of whether PIM is configured in sparse, dense, or sparse-dense mode:

- Configuring a PIM RPF Routing Table on page 198
- Filtering PIM Join Messages on page 199
- Configuring PIM Trace Options on page 200

Configuring a PIM RPF Routing Table

By default, PIM uses `inet.0` as its reverse-path-forwarding (RPF) routing table group. PIM uses an RPF routing table group to resolve its RPF neighbor for a particular multicast source address and to resolve the RPF neighbor for the RP address. PIM can optionally use `inet.2` as its RPF routing table group. To do this, add the `rib-groups` statement to the `routing-options` hierarchy level, and name the routing table group in the `pim` statement:

```
routing-options {
  rib-groups {
    pim-rg {
      import-rib inet.2;
    }
  }
}
protocols {
  pim {
    rib-group inet pim-rg;
  }
}
```

For a list of the hierarchy levels at which you can include these statements, see the statement summary section for this statement.

Specifying additional import routing table groups or an export routing table group in the routing table group has no effect on PIM operation. PIM uses the first routing table group specified as an import routing table group.

PIM uses a single routing table group as its RPF routing table group. This ensures that the route with the longest matching prefix is chosen as the RPF route.

You can also configure IS-IS to populate `inet.2` with IS-IS routes that have regular IP next hops. This allows RPF to work properly even when MPLS is configured for traffic engineering, or IS-IS is configured to use “shortcuts” for local traffic.

For more information on RPF tables and routing protocols, see the *JUNOS Routing Protocols Configuration Guide*.

Filtering PIM Join Messages

While multicast scopes prevent the actual multicast data packets from flowing in or out of an interface, PIM join filters prevent a state from being created in a router. A state—the (*,G) or (S,G) entries—is the information used for forwarding unicast or multicast packets. Using PIM join filters prevents the transport of multicast traffic across a network and the dropping of packets at a scope at the edge of the network. Also, PIM join filters reduce the potential for denial-of-service (DoS) attacks and PIM state explosion—large numbers of PIM join messages forwarded to each router on the rendezvous-point tree (RPT), resulting in memory consumption.

To use PIM join filters to efficiently restrict multicast traffic from certain source addresses, create and apply the routing policy across all routers in the network. See Table 8 for a list of match conditions.

Table 8: PIM Join Filter Match Conditions

Match Condition	Matches On
interface	Router interface or interfaces specified by name or IP address
neighbor	Neighbor address (the source address in the IP header of the join and prune message)
route-filter	Multicast group address embedded in the join and prune message
source-address-filter	Multicast source address embedded in the join and prune message

To create a routing policy to reject a PIM join request for a source, include a policy name at the [edit policy-options policy-statement] or [edit logical-routers *logical-router-name* policy-options policy-statement] hierarchy level.

To apply one or more policies to routes being imported into the routing table from PIM, include the import statement:

```
import [ policy-names ];
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

For a PIM join filter example, see “Example: Configuring PIM Join Filters” on page 225.



NOTE: Configuring multicast scoping on all routers filters the actual data and might be preferable to a PIM join filter solution. For more information about multicast scoping, see “Multicast Scoping Overview” on page 123.

Configuring PIM Trace Options

To trace PIM protocol traffic, you can specify options in the global `traceoptions` statement at the [edit `routing-options`] or [edit `logical-routers logical-router-name routing-options`] hierarchy level. Options applied at the routing options level trace all packets, and options applied at the protocol level trace only IGMP traffic.

```
traceoptions {
  file name <replace> <size size> <files number> <no-stamp>
    <(world-readable | no-world-readable)>;
  flag flag <flag-modifier> <disable>;
}
```

You can configure PIM-specific options by including the `traceoptions` statement at the PIM hierarchy level. For a list of the hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

You can specify the following PIM-specific options in the `traceoptions` statement:

- **assert**—Trace assert messages, which are used to resolve which of the parallel routers connected to a multiaccess LAN is responsible for forwarding packets to the LAN.
- **bootstrap**—Trace bootstrap messages, which are sent periodically by the PIM domain’s bootstrap router and are forwarded, hop by hop, to all routers in that domain.
- **cache**—Trace the packets in the PIM routing cache.
- **graft**—Trace graft and graft acknowledgment messages.
- **hello**—Trace hello packets, which are sent so that neighboring routers can discover each other.
- **join**—Trace join messages, which are sent to join a branch onto the multicast distribution tree.
- **packets**—Trace all PIM packets.
- **prune**—Trace prune messages, which are sent to prune a branch off the multicast distribution tree.
- **register**—Trace register and register-stop messages. Register messages are sent to the RP when a multicast source first starts sending to a group.
- **rp**—Trace candidate RP advertisements.

For general information about tracing, see the *JUNOS System Basics Configuration Guide*. For a PIM tracing example, see “Example: Tracing PIM Protocol Traffic” on page 227.

Configuring PIM Dense Mode Properties

To configure the router properties for PIM dense mode, enable the minimum PIM dense mode configuration. For information about operating interfaces in PIM dense mode, see “PIM Modes” on page 171.

By default, PIM is disabled. When you enable PIM, it operates in sparse mode by default. To enable PIM dense mode on the router, include the `pim` statement:

```
pim {
  rib-group group-name;
  interface interface-name;
  mode dense;
}
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

To specify that PIM dense mode use `inet.2` as its RPF routing table instead of `inet.0`, include the `rib-group` statement. For more information about configuring RPF routing tables, see “Configuring a PIM RPF Routing Table” on page 198.

You can specify the interfaces on which to enable PIM. Specify the full name, including the physical and logical address components. For details about specifying interfaces, see the *JUNOS Network Interfaces Configuration Guide*. If you do not specify any interfaces, PIM is enabled on all router interfaces. Generally, you specify interface names only if you are disabling PIM on certain interfaces.



NOTE: You cannot configure both PIM and Distance Vector Multicast Routing Protocol (DVMRP) in forwarding mode on the same interface. You can configure PIM on the same interface only if you configured DVMRP in unicast-routing mode.

Configuring PIM Sparse Mode Properties

To configure PIM sparse mode properties, see the following sections:

- Minimum PIM Sparse Mode Configuration on page 202
- Logical Routers and PIM Sparse Mode on page 203
- Enabling PIM Sparse Mode on page 203
- Configuring PIM Sparse Mode Graceful Restart on page 203
- Configuring the Router's Local RP Properties on page 205
- Configuring Static RPs on page 208
- Configuring Bootstrap Properties on page 208
- Configuring Auto-RP Announcement, Mapping, and Discovery on page 209
- Configuring Embedded RP for IPv6 on page 213
- Configuring the Assert Timeout on page 214

For information about operating interfaces in PIM sparse mode, see “PIM Modes” on page 171.

Minimum PIM Sparse Mode Configuration

Each any-source multicast (ASM) group has a shared tree through which receivers learn about new multicast sources and new receivers learn about all multicast sources. The RP is the root of this shared tree.

To configure this router's properties as the candidate RP, include the `rp` statement:

```
rp {
  local {
    family (inet | inet6) {
      disable;
      address address;
      group-ranges {
        destination-mask;
      }
      hold-time seconds;
      priority number;
    }
  }
  auto-rp (announce | discovery | mapping);
  bootstrap-export [ policy-names ];
  bootstrap-import [ policy-names ];
  bootstrap-priority number;
}
```

```

embedded-rp {
  maximum-rps limit;
  group-ranges {
    destination-mask;
  }
}
static {
  address address {
    version version;
    group-ranges {
      destination-mask;
    }
  }
}
}

```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Logical Routers and PIM Sparse Mode

Logical routers partition a single physical router into multiple logical devices that perform independent routing tasks. Because logical routers perform a subset of the tasks once handled by the physical router, logical routers offer an effective way to maximize the use of a single router platform.

For an overview of logical routers and a detailed example of logical router configuration, see the logical routers chapter of the *JUNOS Feature Guide*.

Enabling PIM Sparse Mode

You can configure PIM interfaces to operate in sparse, dense, or sparse-dense mode. Sparse mode is the default. There is no need to explicitly configure sparse mode on a PIM interface, but this is often done for clarity or when you configure a change from dense to sparse mode.

To explicitly configure PIM to operate in sparse mode on an interface, include the `mode sparse` statement:

```
mode sparse;
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Configuring PIM Sparse Mode Graceful Restart

You can configure PIM sparse mode to continue to forward existing multicast packet streams during a routing process failure and restart. Only PIM sparse mode can be configured this way. The routing platform does not forward multicast packets for protocols other than PIM during graceful restart, because all other multicast protocols must restart after a routing process failure.



NOTE: If you configure PIM sparse-dense mode, only sparse multicast groups benefit from graceful restart.

The routing platform does not forward new streams until after the restart is complete. After restart, the routing platform refreshes the forwarding state with any updates that were received from neighbors during the restart period. For example, the routing platform relearns the join and prune states of neighbors during the restart, but it does not apply the changes to the forwarding table until after the restart.

When PIM sparse mode is enabled, the routing platform generates a unique 32-bit random number called a generation identifier. Generation identifiers are included by default in PIM hello messages, as specified in the Internet draft Internet draft-draft-ietf-pim-sm-v2-new-10.txt. When a routing platform receives PIM hello messages containing generation identifiers on a point-to-point interface, the JUNOS software activates an algorithm that optimizes graceful restart.

Before PIM sparse mode graceful restart occurs, each routing platform creates a generation identifier and sends it to its multicast neighbors. If a routing platform with PIM sparse mode restarts, it creates a new generation identifier and sends it to neighbors. When a neighbor receives the new identifier, it resends multicast updates to the restarting router to allow it to exit graceful restart efficiently. The restart phase is complete when the restart interval timer expires. On platforms that support PIM sparse mode graceful restart, the restart can be completed within 30 to 300 seconds. The default restart duration is 60 seconds.



NOTE: Multicast forwarding can be interrupted in two ways. First, if the underlying routing protocol is unstable, multicast RPF checks can fail and cause an interruption. Second, because the forwarding table is not updated during the graceful restart period, new multicast streams are not forwarded until graceful restart is complete.

To configure graceful restart for PIM sparse mode, include the **graceful-restart** statement at the **routing-options** hierarchy level, and include the **pim** statement at the appropriate hierarchy levels:

```
[edit routing-options]
  graceful-restart;

  pim {...};
```

For a list of the hierarchy levels at which you can configure the **pim** statement, see the statement summary section for this statement.

To disable graceful restart for PIM, include the **disable** statement:

```
disable;
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

By default, the router allows 60 seconds for the restart duration. The range is from 30 to 300 seconds. After this restart time, the Routing Engine resumes normal multicast operation. To configure the restart duration, include the **restart-duration** statement:

```
restart-duration seconds;
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

For more information about graceful restart for PIM, see “Multicast Redundancy” on page 30. For more information about graceful restart and other routing protocols, see the *JUNOS Routing Protocols Configuration Guide* and *JUNOS Feature Guide*.

Configuring the Router’s Local RP Properties

Local RP configuration makes the router a statically defined RP. To configure the router’s RP properties, include the `local` statement:

```
local {
  family (inet | inet6) {
    disable;
    address address;
    group-ranges {
      destination-mask;
    }
    hold-time seconds;
    priority number;
  }
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols pim rp]
- [edit routing-instances *routing-instance-name* protocols pim rp]

For an overview of logical routers and a detailed example of logical router configuration, see the logical routers chapter of the *JUNOS Feature Guide*.

For information about the RP configuration statements, see the following sections:

- Configuring the IP Protocol Family on page 205
- Configuring the Local RP Address on page 206
- Configuring the Router’s RP Priority on page 207
- Configuring the Groups for Which the Router Is the RP on page 207
- Modifying the Local RP Hold Time on page 207

Configuring the IP Protocol Family

PIM supports both IP version 4 (IPv4) and IP version 6 (IPv6) addressing.

IPv6 PIM hello messages are sent to every interface on which you configure `family inet6`, whether at the PIM level of the hierarchy or not. As a result, if you configure an interface with both `family inet` at the [edit interface *interface-name*] hierarchy level and `family inet6` at the [edit protocols pim interface *interface-name*] hierarchy level, PIM sends both IPv4 and IPv6 hellos to that interface.

By default, PIM operates in sparse mode on an interface. If you explicitly configure sparse mode, PIM uses this setting for all IPv6 multicast groups. However, if you configure sparse-dense mode, PIM does not accept IPv6 multicast groups as dense groups and operates in sparse mode over them.

For correct operation of PIM sparse mode, the RP address should be known to a router. The JUNOS IPv6 PIM implementation supports only static RP configuration. Automatic RP announcement and bootstrap routers are not available with IPv6. You configure the static IPv6 RP address in the same way as IPv4 addresses, by including the `address` statement. However, on a router that is itself the RP, include the `address` statement at the `[edit protocols pim rp local family inet6]` or `[edit routing-instances routing-instance-name protocols pim rp local family inet6]` hierarchy level.

For an overview of logical routers and a detailed example of logical router configuration, see the logical routers chapter of the *JUNOS Feature Guide*.

The Multicast Listener Discovery (MLD) protocol is automatically enabled on any broadcast interfaces on which you configure PIM and `family inet6`. For an overview of MLD, see “MLD Overview” on page 69.

To specify whether IPv4 or IPv6 local RP properties apply to the configuration values, include the `family` statement:

```
family (inet | inet6);
```

You can include this statement at the following hierarchy levels:

- `[edit protocols pim rp local]`
- `[edit routing-instances routing-instance-name protocols pim rp local]`

Configuring the Local RP Address

To specify the local RP address, include the `address` statement:

```
address address;
```

You can include this statement at the following hierarchy levels:

- `[edit protocols pim rp local family]`
- `[edit routing-instances routing-instance-name protocols pim rp local family]`

Configuring the Router's RP Priority

The router's priority value for becoming the RP is included in the bootstrap messages that the router sends. Use a smaller number to increase the likelihood that the router becomes the RP for local multicast groups. Each PIM router uses the priority value and other factors to determine the candidate RPs for a particular group range. After the set of candidate RPs is distributed, each router determines algorithmically the RP from the candidate RP set using a well-known hash function.

By default, the priority value is set to 1. If this value is set to 0, the bootstrap router can override the group range being advertised by the candidate RP. To modify the router's priority, include the `priority` statement:

```
priority number;
```

You can include this statement at the following hierarchy levels:

- [edit protocols pim rp local family]
- [edit routing-instances *routing-instance-name* protocols pim rp local family]

The priority can be a number from 0 through 255.

Configuring the Groups for Which the Router Is the RP

By default, a router running PIM is eligible to be the RP for all groups (224.0.0.0/4). To limit the groups for which this router can be the RP, include the `group-ranges` statement:

```
group-ranges number {
    destination-mask;
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols pim rp local family]
- [edit routing-instances *routing-instance-name* protocols pim rp local family]

Modifying the Local RP Hold Time

For candidate RPs, the hold time is used by the bootstrap router to time out RPs, and applies to the bootstrap RP-set mechanism. The RP hold time is part of the candidate RP advertisement message sent to the bootstrap router. If the bootstrap router does not receive a candidate RP advertisement from an RP within the hold time, it removes that router from its list of candidate RPs. The default hold time is 150 seconds.

To modify the hold-time value for the local RP, include the `hold-time` statement:

```
hold-time seconds;
```

You can include this statement at the following hierarchy levels:

- [edit protocols pim rp local family]
- [edit routing-instances *routing-instance-name* protocols pim rp local family]

Configuring Static RPs

Static RP configuration directs the router to another statically defined RP. To configure static RPs, include the **static** statement:

```
static {
  address address {
    version version;
    group-ranges {
      destination-mask;
    }
  }
}
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.



NOTE: You can configure a static RP in a logical router only if the logical router is not directly connected to a source.

To configure other static RPs, include one or more **address** statements. The default multicast address group range is 224.0.0.0/4.

For each static RP address, you can optionally specify the PIM version and the groups for which this address can be the RP. The default PIM version is version 1.

The RP that you select for a particular group must be consistent across all routers in a multicast domain.



NOTE: The default PIM version can be version 1 or version 2, depending on the mode you are configuring. PIM version 1 is the default for RP mode (at the [pim rp static address address] hierarchy level). However, PIM version 2 is the default for interface mode (at the [pim interface interface-name] hierarchy level). Explicitly configured versions override the defaults.

Configuring Bootstrap Properties

To configure bootstrap properties, see the following sections:

- Configuring the Router's Bootstrap Router Priority on page 208
- Filtering PIM Bootstrap Messages on page 209

Configuring the Router's Bootstrap Router Priority

To determine which router is the RP, all routers within a PIM domain collect bootstrap messages. A PIM domain is a contiguous set of routers that implement PIM; all are configured to operate within a common boundary. The domain's bootstrap router originates bootstrap messages, which are sent hop by hop within the domain. The routers use bootstrap messages to distribute RP information dynamically and to elect a bootstrap router when necessary.

By default, the router has a bootstrap priority of 0, which means the router can never be the bootstrap router. To modify this priority, include the `bootstrap-priority` statement. The router with the highest priority value is elected to be the bootstrap router. In the case of a tie, the router with the highest IP address is elected to be the bootstrap router.

```
bootstrap-priority number;
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Filtering PIM Bootstrap Messages

You can create import and export policies to control the flow of bootstrap messages to and from the RP, and apply them to PIM. To apply one or more import policies to bootstrap messages imported into the RP, include the `bootstrap-import` statement:

```
bootstrap-import [ policy-names ];
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

To apply one or more export policies to bootstrap messages exported from the RP, include the `bootstrap-export` statement:

```
bootstrap-export [ policy-names ];
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

For an example, see “Example: Rejecting PIM Bootstrap Messages at the Boundary of a PIM Domain” on page 227.

Configuring Auto-RP Announcement, Mapping, and Discovery

You can configure a mode-dynamic way of assigning RPs in a multicast network by means of auto-RP. When you configure auto-RP for a router, the router learns the address of the RP in the network automatically. Auto-RP operates in PIM version 1 and version 2.

Although auto-RP is a nonstandard (non-RFC-based) function requiring dense mode PIM to advertise control traffic, it provides an important failover advantage that static RP assignment does not: you can configure multiple routers as RP candidates. If the elected RP stops operating, one of the other preconfigured routers takes over the RP functions. This capability is controlled by the auto-RP mapping agent.

If PIM is operating in sparse or sparse-dense mode, configure how the router operates in auto-RP by specifying the following auto-RP options:

- Use the **discovery** option to let the router receive and process discovery messages from the mapping agent. This is the most basic auto-RP option.
- Add the **announce** option on the router to allow the router to send announce messages in the network, advertising itself as a candidate RP. Routers configured with this option must also be configured as RPs, or announce messages are not sent.
- Add the **mapping** option on the router to allow the router to perform the mapping function. If the router is also an RP, the **mapping** option also allows the router to send auto-RP announcements (mapping on an RP allows the router to perform both the announcement and mapping functions).

The router joins the auto-RP groups on the configured interfaces and on the loopback interface `lo0.0`. For auto-RP to work correctly, configure a routable IP address on the loopback interface. The router ID is used as the address for auto-RP updates. You cannot use the loopback address `127.0.0.1`. Also, you must enable PIM sparse-dense mode on the `lo0.0` interface if you do not specify `interface all`.

In most cases, how the router handles auto-RP discovery, announce, or mapping messages depends on whether the router is an RP (configured as local RP) or not. Table 9 shows how the router behaves depending on the local RP configuration.

Table 9: Local RP and Auto-RP Message Types

Auto-RP Message Type	Local RP?	Router Behavior
discovery	No	Listen for auto-RP mapping messages.
discovery	Yes	Listen for auto-RP mapping messages.
announce	No	Listen for auto-RP mapping messages.
announce	Yes	Listen for auto-RP mapping messages. Send auto-RP announce messages.
mapping	No	Listen for auto-RP mapping messages. Listen for auto-RP announce messages. If elected mapping agent, send auto-RP mapping messages.
mapping	Yes	Listen for auto-RP mapping messages. Send auto-RP announce messages. Listen for auto-RP announce messages If elected mapping agent, send auto-RP mapping messages.

To configure auto-RP at the main hierarchy level, follow these steps:

1. Include the `mode` statement, and specify the option `sparse-dense` on all interfaces at the `[edit protocols pim]` hierarchy level:

```
[edit protocols pim]
interface all {
    mode sparse-dense;
}
```

This configuration allows the router to operate in sparse mode for most groups and dense mode for others. The default is to operate in sparse mode unless the router is specifically informed of a dense mode group.

2. Configure two multicast dense mode groups (224.0.1.39 and 224.0.1.40) using the `dense-groups` statement at the `[edit protocols pim]` hierarchy level:

```
[edit protocols pim]
dense-groups {
    224.0.1.39/32;
    224.0.1.40/32;
}
```

Auto-RP requires multicast flooding to announce potential RP candidates and to discover the elected RPs in the network. Multicast flooding occurs through a PIM dense mode model where group 224.0.1.39 is used for **announce** messages and group 224.0.1.40 is used for **discovery** messages.

3. Finally, include the **auto-RP** statement at the `[edit protocols pim rp]` hierarchy level to configure auto-RP on each router in the group. There are four possible categories for each router.
 - a. If the router is not a local RP and listens only for auto-RP mapping messages, include the **auto-rp discovery** statement to the router RP configuration at the `[edit protocols pim rp]` hierarchy level:

```
[edit protocols pim rp]
auto-rp discovery;
```

- b. If the router is a local RP, sends auto-RP announcements, and listens for auto-RP mapping messages, configure the router as a local RP and include the **auto-rp announce** statement to the router RP configuration at the `[edit protocols pim rp]` hierarchy level:

```
[edit protocols pim rp]
local {
    address 10.0.1.1;
}
auto-rp announce;
```

- c. If the router performs only the mapping function to listen for auto-RP announcements, performs the auto-RP-to-group mapping, and sends auto-RP mapping messages, include the **auto-rp mapping** statement at the `[edit protocols pim rp]` hierarchy level. When multiple candidate RP routers announce their capabilities to support multicast groups, there must be a single router in the network to act as mapping agent. The mapping agent sends out discovery messages to the network, informing all routers in a multicast group of the RP to use:

```
[edit protocols pim rp]
auto-rp mapping;
```

- d. If the router combines the local RP function to send announcements and also perform the mapping function, configure the router as a local RP and include the `auto-rp mapping` statement to the router RP configuration at the `[edit protocols pim rp]` hierarchy level:

```
[edit protocols pim rp]
local {
  address 10.0.1.1;
}
auto-rp mapping;
```

All routers must also have a routable IP address on the `lo0` interface:

```
interface lo0 {
  unit 0 {
    family inet {
      address 127.0.0.1; /* this address cannot be used by auto-rp */
      address 192.168.27.1 { /* this example uses a private IP address */
        preferred;
      }
    }
  }
}
```

You can include these statements at the following hierarchy levels (auto-RP announce is not supported in logical routers):

- `[edit protocols]`
- `[edit routing-instances routing-instance-name protocols]`
- `[edit logical-routers logical-router-name protocols]` (all statements except `auto-rp announce`)
- `[edit logical-routers logical-router-name routing-instances routing-instance-name protocols]` (all statements except `auto-rp announce`)

Use the `show pim rps` command to verify the auto-RP information:

```
user@host> show pim rps
RP address      Type      Holdtime   Timeout   Active groups   Group prefixes
192.168.5.1     auto-rp   150        123       1                224.0.0.0/4
```

Use the `show pim rps extensive` command to see information about how an RP is learned, what groups it handles, and the number of groups actively using the RP:

```
user@host> show pim rps extensive
RP: 192.168.5.1
Learned from 192.168.5.1 via: auto-rp
Time Active: 00:34:29
Holdtime: 150 with 108 remaining
Device Index: 6
Subunit: 32769
Interface: pd-0/0/0.32769
Group Ranges:
  224.0.0.0/4
Active groups using RP:
  224.2.2.100
```

```
total 1 groups active
```

Register State for RP:

Group Holdtime	Source Timeout	FirstHop	RP Address	State	RP address	Type
-------------------	-------------------	----------	------------	-------	------------	------

In the example, the RP at 192.168.5.1 was learned through auto-RP. The RP is able to support all groups in the 224.0.0.0/4 range (all possible groups). The local router has sent PIM control traffic for the 224.2.2.100 group to the RP.

Additionally, the presence of a Tunnel Physical Interface Card (PIC) in an RP router creates a de-encapsulation interface, which allows the RP to receive multicast traffic from the source. This interface is indicated by `pd-0/0/0.32769`.

Configuring Embedded RP for IPv6

You configure embedded RP to allow multidomain IPv6 multicast networks to find RPs in other routing domains. Embedded RP embeds an RP address inside PIM join messages and other types of messages sent between routing domains.

Embedded RP is disabled by default. To configure embedded RP for IPv6 PIM sparse mode, include the `embedded-rp` statement:

```
embedded-rp {
  maximum-rps limit;
  group-ranges {
    destination-mask;
  }
}
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

The `maximum-rps` statement limits the number of embedded RPs created in a specific routing instance. The range is from 1 through 500. The default is 100.

The `group-range` statement determines which multicast addresses or prefixes can embed RP address information. If messages within a group range contain embedded RP information and the group range is not configured, the embedded RP in that group range is ignored. Any valid unicast-prefix-based ASM address can be used as a group range. The default group range is `FF70::/12` to `FFF0::/12`. Messages with embedded RP information that do not match any configured group ranges are treated as normal multicast addresses.

If the derived RP address is not a valid IPv6 unicast address, it is treated as any other multicast group address and not used for RP information. Verification fails if the extracted RP address is a local interface, unless the routing platform is configured as an RP and the extracted RP address matches the configured RP address. Then the local RP decides whether it is configured to act as an RP for the embedded RP multicast address.

When you configure embedded RP for IPv6, embedded RPs are preferred to RPs discovered by IPv6 any other way. You configure embedded RP independent of any other IPv6 multicast properties. This feature is applied only when IPv6 multicast is properly configured.

For more information about the use of embedded RP, see “Embedded RP for IPv6 Multicast” on page 193.

Configuring the Assert Timeout

You configure the assert timeout to determine how often multicast routers running PIM sparse mode enter a PIM assert message cycle. Multicast routers running PIM sparse mode often forward the same stream of multicast packets onto the same LAN through the rendezvous-point tree (RPT) and shortest-path tree (SPT). PIM assert messages help routers determine which router forwards the traffic and prunes the RPT for this group.

Assert messages are useful for LANs that connect multiple routers and no hosts. For more information about network configurations using assert timeouts, see “PIM Sparse-Mode SPT Cutover” on page 181.

To configure the assert timeout for PIM sparse mode, include the `assert-timeout` statement:

```
assert-timeout seconds;
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

The range is from 5 through 210 seconds. The default is 210 seconds.

Configuring Sparse-Dense Mode Properties

To configure PIM to operate in sparse-dense mode on an interface, include the `mode sparse-dense` statement. Include the `dense-groups` statement to specify which groups are operating in dense mode:

```
dense-groups {
    addresses;
}
interface interface-name {
    mode sparse-dense;
}
```

For a list of the hierarchy levels at which you can include these statements, see the statement summary section for this statement.

You can configure graceful restart with PIM sparse-dense mode, but only sparse multicast groups benefit from graceful restart. For more information about graceful restart for PIM sparse mode, see “Configuring PIM Sparse Mode Graceful Restart” on page 203.

For an example of a sparse-dense mode configuration, see “Example: Configuring Sparse-Dense Mode” on page 221.

Configuring Multicast for Layer 3 VPNs

If the service provider supports PIM, you can configure multicast for a Layer 3 virtual private network (VPN) using PIM version 2 as the multicast protocol. The JUNOS software complies with RFC 2547 and Internet draft draft-rosen-vpn-mcast-07.txt, *BGP/MPLS VPNs* and *Multicast in MPLS/BGP VPNs*, Section 2 (Multicast Domains).

For multicast to work on Layer 3 VPNs, each of the following routers must have a Tunnel Services PIC, which is hardware used to encapsulate and de-encapsulate data packets into tunnels:

- Each provider edge (PE) router
- Any provider (P) router acting as the RP
- Any customer edge (CE) router that is acting as a source's DR or as an RP. A receiver's designated router does not need a Tunnel PIC.

When you complete the configuration, two multicast tunnel interfaces are configured automatically. You do not need to configure the tunnel interfaces. The interface `mt-[xxxx]`, used for encapsulation, is in the range from 32,768 through 49,151. The interface `mt-[yyyy]`, used for de-encapsulation, is in the range from 49,152 through 65,535. For each VPN, the PE routers build a multicast distribution tree within the service provider core network. After the tree is created, each PE router encapsulates all multicast traffic (data and control messages) from the attached VPN and sends the encapsulated traffic to the VPN group address. Because all the PE routers are members of the outgoing interface list in the multicast distribution tree for the VPN group address, they all receive the encapsulated traffic. When the PE routers receive the encapsulated traffic, they de-encapsulate the messages and send the data and control messages to the CE routers.



NOTE: It is possible for the PE router to be configured as the VPN customer RP (C-RP) router. The PE router can also act as the DR. This type of PE configuration can simplify configuration of customer DRs and VPN C-RPs for multicast VPNs. However, the BSR and auto-RP features are not supported. This section does not discuss the use of the PE as the VPN C-RP.


This section describes how to configure multicast for Layer 3 VPNs:

- Configuring the VPN on page 216
- Configuring PIM Connectivity Between the Provider and PE Routers on page 216
- Configuring Multicast Connectivity on the CE Routers on page 217
- Configuring Multicast Connectivity for the VPN on the PE Router on page 217
- Configuring the Routing Group on page 218

Configuring the VPN

You must first configure the VPN. Figure 30 shows a configuration for VPN-A, used as an example later in this section. For more information about configuring VPNs, see the *JUNOS VPNs Configuration Guide*.

Figure 30: Configuring the VPN



1487


Configuring PIM Connectivity Between the Provider and PE Routers

To configure PIM on the main routing instance for all provider and PE routers, include statements at the [edit protocols pim] hierarchy level:

1. Configure the interfaces between each provider router and the PE routers by including the `interface` statement at the [edit protocols pim] hierarchy level. On all PE routers, enable PIM version 2 and sparse mode on interface `lo0` of the PE routers, either by configuring that specific interface or by including the statement `set version 2 mode sparse` for `interface all` at the [edit protocols pim] hierarchy level on a PE router.
2. Configure PIM version 2 by including the `version` statement at the [edit protocols pim interface *interface-name*] hierarchy level.
3. Configure sparse mode (the mode in which the PIM interfaces operate) by including the `mode` statement at the [edit protocols pim interface *interface-name*] hierarchy level.
4. Configure the RP address by including the `static` statement at the [edit protocols pim rp] hierarchy level. In Figure 31, the provider router is the RP.

Figure 31 shows a multicast configuration on the provider network.

Figure 31: Multicast Configuration on the Provider Network



1488


Configuring Multicast Connectivity on the CE Routers

To configure PIM for the master routing instance on all CE routers, include statements at the [edit protocols pim] hierarchy level:

1. Configure the interfaces going toward the provider router acting as the RP by including the `interface` statement at the [edit protocols pim] hierarchy level. In Figure 32, the interfaces are labeled A and B.
2. Configure PIM version 2 by including the `version` statement at the [edit protocols pim interface *interface-name*] hierarchy level.
3. Configure sparse mode or sparse-dense mode (the mode in which the PIM interfaces operate) by including the `mode` statement at the [edit protocols pim interface *interface-name*] hierarchy level.
4. Configure the RP address by including the `static` statement at the [edit protocols pim rp] hierarchy level. In Figure 32 on page 217, CE2 is the RP router; however, the RP router can be anywhere in the customer network.

Figure 32 shows multicast connectivity on the customer edge.

Figure 32: Multicast Connectivity on the CE Routers



1489

Configuring Multicast Connectivity for the VPN on the PE Router

To configure multicast connectivity for the VPN on the PE router, you must configure a VPN group address and configure the interfaces toward the router acting as RP. To configure the VPN group address, include the `vpn-group-address` statement at the [edit routing-instances *instance-name* protocols pim] hierarchy level:


```
[edit routing-instances instance-name protocols pim]
  vpn-group-address address;
```

The PIM configuration in the VPN routing and forwarding (VRF) instance on the PE routers should match the master PIM instance on the CE router. Therefore, the PE router contains both a master PIM instance (to communicate with the provider core) and the VRF instance (to communicate with the CE routers). See the *JUNOS VPNs Configuration Guide* for information about configuring VPNs on PE routers.



NOTE: VRF instances that are part of the same VPN share the same VPN group address. For example, all PE routers containing multicast-enabled routing instance VPN-A share the same VPN group address configuration. In Figure 33, the shared VPN group address configuration is 239.1.1.1.

Figure 33: Multicast Connectivity for the VPN



Configuring the Routing Group

Routing groups are usually configured at the [edit routing-instances routing-options] hierarchy level. However, with multicast in VRF instances, you must configure routing groups differently. Configure the multicast routing group by adding the `rib-groups` statement at the [edit routing-options] hierarchy level.

After you configure the multicast routing group in the main routing instance, add the routing group to the VPN’s VRF instance. To do this, include the `rib-group` statement at the [edit routing-instances instance-name protocols pim] hierarchy level.

For a multicast for Layer 3 VPN example, see “Example: Configuring PIM Sparse Mode over Layer 3 VPNs” on page 228.

Configuring Multicast for Virtual Routers

You can configure PIM for the `virtual-router` routing instance type as well as for the `vrf` instance type. The `virtual-router` instance type is similar to the `vrf` instance type used with Layer 3 VPNs, but is used for non-VPN-related applications.

The `virtual-router` instance type has no VRF import, VRF export, VRF target, or route distinguisher requirements. The `virtual-router` instance type is used for non-Layer 3 VPN situations, for example, to allow the use of IP Security (IPSec) tunnels within VPNs.

When PIM is configured under the `virtual-router` instance type, the VPN configuration is not based on RFC 2547, *BGP/MPLS VPNs*, so PIM operation does not comply with the Internet draft draft-rosen-vpn-mcast-07.txt *Multicast in MPLS/BGP VPNs*. For more information about multicast draft support, see “IP Multicast Standards” on page 28. In the `virtual-router` instance type, PIM operates in a routing instance by itself, forming adjacencies with PIM neighbors over the routing instance interfaces as the other routing protocols do with neighbors in the routing instance.

To configure PIM for a `virtual-router` instance type, include the `pim` statement and specify the `virtual-router` instance type:

```
instance-type virtual-router;
protocols {
  pim {
    ...pim-configuration...
```

```
    }
}
```

You can include these statements at the following hierarchy levels:

- [edit routing-instances *routing-instance-name*]
- [edit logical-routers *logical-router-name* routing-instances *routing-instance-name*]

Do not include the `vpn-group-address` statement for the `virtual-router` instance type.

Configuration Examples

This section contains the following PIM configuration examples:

- Example: Configuring PIM Dense Mode on page 219
- Example: Configuring PIM Sparse Mode on page 220
- Example: Configuring Sparse-Dense Mode on page 221
- Example: Configuring Anycast RP on page 221
- Example: Configuring PIM BSR Filters on page 225
- Example: Configuring PIM Join Filters on page 225
- Example: Configuring Externally-Facing Border Routers on page 226
- Example: Tracing PIM Protocol Traffic on page 227
- Example: Rejecting PIM Bootstrap Messages at the Boundary of a PIM Domain on page 227
- Example: Configuring PIM Sparse Mode over Layer 3 VPNs on page 228
- Example: Configuring PIM Dense Mode over Layer 3 VPNs on page 236
- Example: Configuring PIM Sparse-Dense Mode over Layer 3 VPNs on page 239

Example: Configuring PIM Dense Mode

The following example shows a configuration for PIM dense mode:

```
[edit protocols]
pim {
  interface so-5/0/1 {
    mode dense;
  }
  interface so-5/0/2 {
    mode dense;
  }
  traceoptions {
    file log-pim;
    flag normal;
  }
}
```

```

        flag state;
    }
}

```

Example: Configuring PIM Sparse Mode

The following example shows a configuration for the RP router and for non-RP routers.

Configuring the RP Router

This example shows a static RP configuration. Add the `address` statement at the `[edit protocols pim rp local]` hierarchy level.

For all interfaces, use the `mode` statement to set the mode to sparse, and use the `version` statement to set the PIM version to 2 at the `[edit protocols PIM rp interface all]` hierarchy level. When configuring all interfaces, exclude the `fxp0.0` management interface by adding the `disable` statement for that interface.



NOTE: You do not need to configure Internet Group Management Protocol (IGMP) version 2 for a sparse mode configuration. When PIM is enabled, by default, IGMP version 2 is also enabled.

```

[edit]
protocols {
  pim {
    rp {
      local {
        address 198.58.3.253;
      }
      interface all {
        mode sparse;
        version 2;
      }
      interface fxp0.0 {
        disable;
      }
    }
  }
}

```

Configuring All Non-RP Routers

In this example, configure a non-RP router for PIM sparse mode. To specify a static RP address, add the `address` statement at the `[edit protocols pim rp static]` hierarchy level. Use the `version` statement at the `[edit protocols pim rp static address]` hierarchy level to specify PIM version 2.

Add the `mode` statement at the `[edit protocols pim interface all]` hierarchy level to configure the interfaces for sparse mode operation. Then add the `version` statement at the `[edit protocols pim interface all mode]` to specify PIM version 2 for all interfaces. When configuring all interfaces, exclude the `fxp0.0` management interface by adding the `disable` statement for that interface.

```
[edit]
protocols {
  pim {
    rp {
      static {
        address 198.58.3.253 {
          version 2;
        }
      }
    }
    interface all {
      mode sparse;
      version 2;
    }
    interface fxp0.0 {
      disable;
    }
  }
}
```

Example: Configuring Sparse-Dense Mode

Configure PIM sparse-dense mode on all interfaces, specifying that the groups 224.0.1.39 and 224.0.1.40 are using dense mode:

```
[edit protocols]
pim {
  dense-groups {
    224.0.1.39;
    224.0.1.40;
  }
  interface all {
    version 1;
    mode sparse-dense;
  }
}
```

Example: Configuring Anycast RP

When you configure anycast RP, you bypass the restriction of having one active RP per multicast group, and instead deploy multiple RPs for the same group range. The RP routers share one unicast IP address. Sources from one RP are known to other RPs that use Multicast Source Discovery Protocol (MSDP). Sources and receivers use the closest RP, as determined by the interior gateway protocol (IGP).

You can use anycast RP within a domain to provide redundancy and RP load sharing. When an RP goes down, sources and receivers are taken to a new RP by means of unicast routing.

You can configure anycast RP to use PIM and MSDP for IPv4, or PIM alone for both IPv4 and IPv6 scenarios. Both are covered in this section.

For information about standards supported for anycast RP, see “IP Multicast Standards” on page 28.

We recommend a static RP mapping with anycast RP over a bootstrap router and auto-RP configuration because it provides all the benefits of a bootstrap router and auto-RP without the complexity of the BSR and auto-RP mechanisms.

The following example shows an anycast RP configuration for the RP routers, first with MSDP and then using PIM alone, and for non-RP routers.

Configuring the RP Router with MSDP

In this example, configure an RP using the lo0 or loopback interface, which is always up. Use the `address` statement to specify the unique and routable router ID and the RP address at the `[edit interfaces lo0 unit 0 family inet]` hierarchy level. In this case, the router ID is `198.58.3.254/32` and the shared RP address is `198.58.3.253/32`. Add the flag statement `primary` to the first address. Using this flag selects the router's primary address from all the preferred addresses on all interfaces.

```
[edit]
interfaces {
  lo0 {
    description "PIM RP";
    unit 0 {
      family inet {
        address 198.58.3.254/32;
        primary;
        address 198.58.3.253/32;
      }
    }
  }
}
```

Add the `address` statement at the `[edit protocols pim rp local]` hierarchy level to specify the RP address (the same address as the secondary lo0).

For all interfaces, use the `mode` statement to set the mode to `sparse` and the `version` statement to specify PIM version 2 at the `[edit protocols pim rp local interface all]` hierarchy level. When configuring all interfaces, exclude the `fxp0.0` management interface by adding the `disable` statement for that interface.

```
[edit]
protocols {
  pim {
    rp {
      local {
        family inet;
        address 198.58.3.253;
      }
      interface all {
        mode sparse;
        version 2;
      }
      interface fxp0.0 {
        disable;
      }
    }
  }
}
```

To configure MSDP peering, add the `peer` statement to configure the address of the MSDP peer at the `[edit protocols msdp]` hierarchy level. For MSDP peering, use the unique, primary addresses instead of the anycast address. To specify the local address for MSDP peering, add the `local-address` statement at the `[edit protocols msdp peer]` hierarchy level.

```
[edit]
protocols {
  msdp {
    peer 198.58.3.250 {
      local-address address 198.58.3.254;
    }
  }
}
```

Configuring the RP Router Using Only PIM

In this example, configure an RP using the `lo0` or loopback interface, which is always up. Use the `address` statement to specify the unique and routable router address and the RP address at the `[edit interfaces lo0 unit 0 family inet]` hierarchy level. In this case, the router ID is `198.58.3.254/32` and the shared RP address is `198.58.3.253/32`. Add the flag statement `primary` to the first address. Using this flag selects the router's primary address from all the preferred addresses on all interfaces.

```
[edit]
interfaces {
  lo0 {
    description "PIM RP";
    unit 0 {
      family inet {
        address 198.58.3.254/32;
        primary;
        address 198.58.3.253/32;
      }
    }
  }
}
```

Add the `address` statement at the `[edit protocols pim rp local]` hierarchy level to specify the RP address (the same address as the secondary `lo0` interface).

For all interfaces, use the `mode` statement to set the mode to `sparse`, and the `version` statement to specify PIM version 2 at the `[edit protocols pim rp local interface all]` hierarchy level. When configuring all interfaces, exclude the `fxp0.0` management interface by adding the `disable` statement for that interface.

Use the `anycast-pim` statement to configure anycast RP without MSDP (for example, if IPv6 is used for multicasting). The other RP routers that share the same IP address are configured using the `rp-set` statement. There is one entry for each RP, and the maximum that can be configured is 15. For each RP, specify the routable IP address of the router and whether MSDP source active (SA) messages should be forwarded to the RP.

```
[edit]
protocols {
  pim {
    rp {
      local {
        family inet {
          address 198.58.3.253;
          anycast-pim {
            rp-set {
              address 198.58.3.240;
              address 198.58.3.241 forward-msdp-sa;
            }
            local-address 198.58.3.254; #If not configured, lo0 primary is
                                         used
          }
        }
      }
    }
  }
  interface all {
    mode sparse;
    version 2;
  }
  interface fxp0.0 {
    disable;
  }
}
}
```

MSDP configuration is not necessary for this type of IPv4 anycast RP configuration.

Configuring All Non-RP Routers

Whether MSDP is used or not, the anycast RP configuration for a non-RP router is the same as a static RP configuration for a non-RP router. Specify a static RP by adding the address at the [edit protocols pim rp static] hierarchy level. Use the **version** statement at the [edit protocols pim rp static address] hierarchy level to set PIM version 2.

```
[edit]
protocols {
  pim {
    rp {
      static {
        address 198.58.3.253 {
          version 2;
        }
      }
    }
  }
}
}
```

Use the **mode** statement at the [edit protocols pim rp interface all] hierarchy level to specify sparse mode on all interfaces. Then add the **version** statement at the [edit protocols pim rp interface all mode] to configure all interfaces for PIM version 2. When configuring all interfaces, exclude the **fxp0.0** management interface by adding the **disable** statement for that interface.

```
[edit]
```

```

protocols {
  pim {
    interface all {
      mode sparse;
      version 2;
    }
    interface fxp0.0 {
      disable;
    }
  }
}

```

Example: Configuring PIM BSR Filters

Configure a filter to prevent BSR messages from entering or leaving your network. Add this configuration to all routers.

```

[edit]
protocols {
  pim {
    rp {
      bootstrap-import no-bsr;
      bootstrap-export no-bsr;
    }
  }
}
policy-options {
  policy-statement no-bsr {
    then reject;
  }
}

```

Example: Configuring PIM Join Filters

In this example, you create the PIM join filter by including the `import pim-join-filter` statement at the `[edit protocols pim]` hierarchy level. Define `pim-join-filter` by adding the `policy-statement pim-join` filter statement at the `[edit policy-options]` hierarchy level. The filter is composed of a route filter and a source address filter—`bad-groups` and `bad-sources`, respectively. Policy `bad-groups` prevents (*,G) or (S,G) join messages from being received for all groups listed. Policy `bad-sources` prevents (S,G) join messages from being received for all sources listed. The `bad-groups` filter and `bad-sources` filter are in two different terms. If route filters and source address filters are in the same term, they are logically ANDed.

```

[edit]
protocols {
  pim {
    import pim-join-filter;
  }
}

policy-statement pim-join-filter {
  term bad-groups {
    from {
      route-filter 224.0.1.2/32 exact;
      route-filter 224.0.1.3/32 exact;
    }
  }
}

```

```

route-filter 224.0.1.8/32 exact;
route-filter 224.0.1.22/32 exact;
route-filter 224.0.1.24/32 exact;
route-filter 224.0.1.25/32 exact;
route-filter 224.0.1.35/32 exact;
route-filter 224.0.1.39/32 exact;
route-filter 224.0.1.40/32 exact;
route-filter 224.0.1.60/32 exact;
route-filter 224.0.2.1/32 exact;
route-filter 224.0.2.2/32 exact;
route-filter 225.1.2.3/32 exact;
route-filter 229.55.150.208/32 exact;
route-filter 234.42.42.42/30 orlonger;
route-filter 239.0.0.0/8 orlonger;
}
then reject;
}
term bad-sources {
  from {
    source-address-filter 10.0.0.0/8 orlonger;
    source-address-filter 127.0.0.0/8 orlonger;
    source-address-filter 172.16.0.0/12 orlonger;
    source-address-filter 192.168.0.0/16 orlonger;
  }
  then reject;
}
term last {
  then accept;
}
}

```

Example: Configuring Externally-Facing Border Routers

In this example, you add the `scope` statement at the `[edit routing-options multicast]` hierarchy level to prevent auto-RP traffic from “leaking” into or out of your PIM domain. Two scopes defined below, `auto-rp-39` and `auto-rp-40`, are for specific addresses. The `scoped-range` statement defines a group range, thus preventing group traffic from leaking.

```

[edit]
routing-options {
  multicast {
    scope auto-rp-39 {
      prefix 224.0.1.39/32;
      interface t1-0/0/0.0;
    }
    scope auto-rp-40 {
      prefix 224.0.1.40/32;
      interface t1-0/0/0.0;
    }
    scope scoped-range {
      prefix 239.0.0.0/8;
      interface t1-0/0/0.0;
    }
  }
}

```

Example: Tracing PIM Protocol Traffic

Trace only unusual or abnormal operations to a routing log file, and trace detailed information about all PIM messages to a PIM log file:

```
[edit]
routing-options {
  traceoptions {
    file routing-log;
    flag errors;
  }
}
protocols {
  pim {
    interface so-0/0/0;
    traceoptions {
      file pim-log;
      flag packets;
    }
  }
}
```

Example: Rejecting PIM Bootstrap Messages at the Boundary of a PIM Domain

In this example, the policy statement from interface so-0-1/0 then reject rejects bootstrap messages from the specified interface:

```
[edit]
protocols {
  pim {
    rp {
      bootstrap-import pim-import;
      bootstrap-export pim-export;
    }
  }
}
policy-options {
  policy-statement pim-import {
    from interface so-0/1/0;
    then reject;
  }
  policy-statement pim-export {
    to interface so-0/1/0;
    then reject;
  }
}
```


Example: Configuring PIM Sparse Mode over Layer 3 VPNs

This section illustrates how multicast is configured in PIM sparse mode for a multicast range for VPN-A (see Figure 34), and shows how to configure the following:

- Configuring PIM on the P Router on page 228
- Configuring PIM on the PE1 Router on page 229
- Configuring PIM on the PE2 Router on page 229
- Configuring PIM on the CE1 Router on page 230
- Configuring PIM on the CE2 Router on page 230
- Configuring the Routing Instance on the PE1 Router on page 231
- Configuring the Routing Instance on the PE2 Router on page 232
- Configuring the PE Router for Interoperability on page 234
- Configuring the Routing Table Group on page 234

For information about configuring VPNs, see the *JUNOS VPNs Configuration Guide*.

Figure 34: Customer Edge and Service Provider Networks



1492

Configuring PIM on the P Router

Configure PIM on the P router. The P router acts as the P (RP) router in this example. Specify the P router’s address (10.255.71.47) at the [edit protocols pim rp local] hierarchy level.

```
[edit]
protocols {
  pim {
    dense-groups {
      224.0.1.39/32;
      224.0.1.40/32;
    }
    rp {
      local {
        address 10.255.71.47;
      }
    }
  }
}
```

```

        interface all {
            mode sparse;
            version 2;
        }
        interface fxp0.0 {
            disable;
        }
    }
}

```

Configuring PIM on the PE1 Router

Configure PIM on the provider edge 1 (PE1) router. Specify a static route to the service provider RP router—the P router (10.255.71.47).

```

[edit]
protocols {
  pim {
    rp {
      static {
        address 10.255.71.47;
      }
    }
    interface all {
      mode sparse;
      version 2;
    }
    interface fxp0.0 {
      disable;
    }
  }
}

```

Configuring PIM on the PE2 Router

Configure PIM on the provider edge 2 (PE2) router. Specify a static route to the service provider RP—the P router (10.255.71.47).

```

[edit]
protocols {
  pim {
    rp {
      static {
        address 10.255.71.47;
      }
    }
    interface all {
      mode sparse;
      version 2;
    }
    interface fxp0.0 {
      disable;
    }
  }
}

```

Configuring PIM on the CE1 Router

Configure PIM on the customer edge (CE1) router. Specify the RP address for the VPN RP—router CE2 (10.255.245.91).

```
[edit]
protocols {
  pim {
    rp {
      static {
        address 10.255.245.91;
      }
    }
    interface all {
      mode sparse;
      version 2;
    }
    interface fxp0.0 {
      disable;
    }
  }
}
```

Configuring PIM on the CE2 Router

Configure PIM on the customer edge 2 (CE2) router, which acts as the VPN RP. Specify router CE2's address (10.255.245.91) at the [edit protocols pim rp local] hierarchy level:

```
[edit]
protocols {
  pim {
    rp {
      local {
        address 10.255.245.91;
      }
    }
    interface all {
      mode sparse;
      version 2;
    }
    interface fxp0.0 {
      disable;
    }
  }
}
```

Configuring the Routing Instance on the PE1 Router

Configure the routing instance (VPN-A) for the Layer 3 VPN on router PE1. As part of the configuration, you need to establish the PIM instance for the VPN. Use the `vpn-group-address` statement at the `[edit routing-instances routing-instance-name protocols pim]` hierarchy level to specify the VPN group address, which is needed for multicast over a Layer 3 VPN configuration.

Set the RP configuration for the VRF instance at the `[edit routing-instances routing-instance-name protocols pim]` hierarchy level. The RP configuration within the VRF instance provides explicit knowledge of the RP address, so that the (*,G) state can be forwarded.

For Release 5.5 or later, configure an additional unit on the loopback interface of the PE router at the `[edit interfaces]` hierarchy level, and assign an address from the VPN address space. Then add the newly created loopback interface in two places:

- Routing instance (VPN-A) at the `[edit routing-instances routing-instance-name]` hierarchy level.
- Routing instance (VPN-A) at the `[edit routing-instances routing-instance-name protocols pim]` hierarchy level.

Also, add the loopback interface to the IGP and Border Gateway Protocol (BGP) policies to advertise the interface in the VPN address space. For more information about how to configure a logical unit on a loopback interface, see the *JUNOS VPNs Configuration Guide*.

In multicast Layer 3 VPNs, the multicast PE routers must use the primary loopback address (or router ID) for sessions with their internal BGP peers. If the PE routers use a route reflector with next-hop self configured, Layer 3 multicast over VPN will not work because PIM cannot transmit upstream interface information for multicast sources behind remote PEs into the network core. Multicast Layer 3 VPNs require the BGP next-hop address of the VPN route to match the BGP next-hop address of the loopback VRF instance address.

```
[edit]
routing-instances {
  VPN-A {
    instance-type vrf;
    interface t1-1/0/0:0.0;
    interface lo0.1;
    route-distinguisher 10.255.71.46:100;
    vrf-import VPNA-import;
    vrf-export VPNA-export;
    protocols {
      ospf {
        export bgp-to-ospf;
        area 0.0.0.0 {
          interface t1-1/0/0:0.0;
          interface lo0.1;
        }
      }
    }
    pim {
      vpn-group-address 239.1.1.1;
      rp {
```



```

protocols {
  ospf {
    export bgp-to-ospf;
    area 0.0.0.0 {
      interface t1-2/0/0:0.0;
      interface lo0.1;
    }
  }
  pim {
    vpn-group-address 239.1.1.1;
    rp {
      static {
        address 10.255.245.91;
      }
    }
    interface t1-2/0/0:0.0 {
      mode sparse;
      version 2;
    }
    interface lo0.1 {
      mode sparse;
      version 2;
    }
  }
}
interfaces {
  lo0 {
    description "unit 1 has the important PIM address"
  }
  unit 0 {
    family inet {
      address 192.168.27.14/32;
      primary;
      address 127.0.0.1/32;
    }
  }
  unit 1 {
    family inet {
      address 10.10.47.102/32;
    }
  }
}
}

```



NOTE: Multicast Layer 3 VPNs require the BGP next-hop address of the VPN route to match the BGP next-hop address of the loopback VRF instance address.

Configuring the PE Router for Interoperability

When one of the PE routers is running Cisco Systems IOS software, you must configure the Juniper Networks PE router to support this multicast interoperability requirement. The Juniper Networks PE router must have the `lo0.0` interface in the master routing instance and the `lo0.1` interface assigned to the VPN routing instance. You must configure the `lo0.1` interface with the same IP address that the `lo0.0` interface uses for BGP peering in the provider core in the master routing instance.

Configure the same IP address on the `lo0.0` and `lo0.1` loopback interfaces of the Juniper Networks PE router at the `[edit interfaces]` hierarchy level, and assign the address used for BGP peering in the provider core in the master routing instance.

```
[edit]
interfaces {
  lo0 {
    description "unit 0 and unit 1 configured for Cisco IOS interoperability";
    unit 0 {
      family inet {
        address 192.168.27.14/32;
        primary;
        address 127.0.0.1/32;
      }
    }
    unit 1 {
      family inet {
        address 192.168.27.14/32;
      }
    }
  }
}
```

Configuring the Routing Table Group

Configure the multicast routing table group by adding the `VPNA-mcast-rib` statement at the `[edit routing-options]` hierarchy level. This group accesses `inet.2` when doing RPF checks. However, if you are using `inet.0` for multicast RPF checks, this step will prevent your multicast configuration from working.

You must also include the interface routes in `inet.2`. For more information about creating routing table groups, see the *JUNOS Routing Protocols Configuration Guide*.

```
[edit]
routing-options {
  interface-routes {
    rib-group VPNA-mcast-rib;
  }
  rib-groups {
    VPNA-mcast-rib {
      export-rib VPN-A.inet.2;
      import-rib VPN-A.inet.2;
    }
  }
}
```

After you configure the multicast routing table group, activate it by including the statement `rib-group inet VPN-A-mcast-rib` at the [edit routing-instances *instance-name* protocols pim] hierarchy level of the VPN's VRF instance.

```
[edit]
routing-instances {
  VPN-A {
    protocols {
      pim {
        rib-group inet VPN-A-mcast-rib;
      }
    }
  }
}
```

Use the following commands to verify the configuration:

- To display all PE tunnel interfaces, issue the command `show pim join` from the provider router acting as the RP.
- To display multicast tunnel information and the number of neighbors, issue the command `show pim interfaces instance instance-name` from the PE1 or PE2 router. When issued from the PE1 router, the output display is:

```
user@host> show pim interfaces instance VPN-A
Instance: PIM.VPN-A

Name                Stat Mode      IP V State Count DR address
-----
lo0.1                Up   Sparse      4 2 DR         0 10.10.47.101
mt-1/1/0.32769       Up   Sparse      4 2 DR         1
mt-1/1/0.49154       Up   Sparse      4 2 DR         0
pe-1/1/0.32769       Up   Sparse      4 1 P2P        0
t1-2/1/0:0.0         Up   Sparse      4 2 P2P        1
```

- To display multicast tunnel interface information, DR information, and the PIM neighbor status between VRF instances on PE1 and PE2, issue the command `show pim neighbors instance instance-name` from either PE router. When issued from the PE1 router, the output display is:

```
user@host> show pim neighbors instance VPN-A
Instance: PIM.VPN-A

Interface           IP V Mode      Option      Uptime Neighbor addr
-----
mt-1/1/0.32769      4 2            HPL         01:40:46 10.10.47.102
t1-1/0/0:0.0        4 2            HPL         01:41:41 192.168.196.178
```

Example: Configuring PIM Dense Mode over Layer 3 VPNs

Multicast over Layer 3 VPNs for dense mode works much the same way as in sparse mode. In the following example, the VPN network uses dense mode for the entire multicast group range. Compare this with the configuration used in “Example: Configuring PIM Sparse Mode over Layer 3 VPNs” on page 228. In that configuration, sparse mode is used for the entire multicast group range.

To support PIM dense mode over Layer 3 VPNs, follow the same steps used in “Example: Configuring PIM Sparse Mode over Layer 3 VPNs” on page 228, with the following differences:

- Configure dense mode for the CE router using the `mode` statement at the `[edit protocols pim interface]` hierarchy level. In the example below, the CE-facing interface is `t1-1/0/0:0`.
- Configure dense mode in the routing instance of the PE router facing the CE router (configured for dense mode) using the `mode` statement at the `[edit routing-instances instance-name protocols pim]` hierarchy level.
- Remove the RP configurations from the CE router and from the routing instance on the PE router.

This section shows how to configure the following:

- Configuring PIM on the P Router on page 236
- Configuring PIM on the PE Router on page 237
- Configuring PIM on the CE Router on page 237
- Configuring the Routing Instance on the PE Router on page 238

For information about configuring VPNs, see the *JUNOS VPNs Configuration Guide*.

Configuring PIM on the P Router

Configure PIM on the P router as in the PIM sparse mode example:

```
[edit]
protocols {
  pim {
    dense-groups {
      224.0.1.39/32;
      224.0.1.40/32;
    }
    rp {
      local {
        address 10.255.71.47;
      }
    }
    interface all {
      mode sparse;
      version 2;
    }
  }
}
```

```

        interface fxp0.0 {
            disable;
        }
    }
}

```

Configuring PIM on the PE Router

Configure PIM on the PE router. Use the `mode` statement at the [edit protocols pim interface] hierarchy level to specify `sparse` mode.

```

[edit]
protocols {
  pim {
    rp {
      static {
        address 10.255.71.47;
      }
    }
    interface all {
      mode sparse;
      version 2;
    }
    interface fxp0.0 {
      disable;
    }
  }
}

```

Configuring PIM on the CE Router

Configure PIM on the CE router. Use the `mode` statement at the [edit protocols pim interface] hierarchy level to specify `dense` mode. An RP is not used with `dense` mode, so no RP statements are required on the CE router.

```

[edit]
protocols {
  pim {
    interface all {
      mode dense;
      version 2;
    }
    interface fxp0.0 {
      disable;
    }
  }
}

```

Configuring the Routing Instance on the PE Router

Use the `mode` statement at the `[edit routing-instances instance pim interface]` hierarchy level to specify `dense` mode for interface `t1-1/0/0:0.0`. An RP is not used with `dense` mode, so no RP statements are required for the routing instance on the PE router.

```
[edit]
routing-instances {
  VPN-A {
    instance-type vrf;
    interface t1-1/0/0:0.0;
    interface lo0.1;
    route-distinguisher 10.255.71.46:100;
    vrf-import VPN-A-import;
    vrf-export VPN-A-export;
    protocols {
      ospf {
        export bgp-to-ospf;
        area 0.0.0.0 {
          interface t1-1/0/0:0.0;
          interface lo0.1;
        }
      }
      pim {
        vpn-group-address 239.1.1.1;
        interface t1-1/0/0:0.0 {
          mode dense;
          version 2;
        }
        interface lo0.1 {
          mode dense;
          version 2;
        }
      }
    }
  }
}
interfaces {
  lo0 {
    description "unit 1 has the important PIM address";
    unit 0 {
      family inet {
        address 192.168.27.13/32;
        primary;
        address 127.0.0.1/32;
      }
    }
    unit 1 {
      family inet {
        address 10.10.47.101/32;
      }
    }
  }
}
```

Example: Configuring PIM Sparse-Dense Mode over Layer 3 VPNs

Multicast over Layer 3 VPNs for sparse-dense mode works much the same way as in sparse mode. In the following example, the VPN network uses dense mode for group range 229.0.0.0/8 and sparse mode for the remaining multicast group range outside 229.0.0.0/8. Compare this with the configuration used in “Example: Configuring PIM Sparse Mode over Layer 3 VPNs” on page 228. In that configuration, sparse mode is used for the entire multicast group range.

To support PIM dense mode over Layer 3 VPNs, follow the same steps used in “Example: Configuring PIM Sparse Mode over Layer 3 VPNs” on page 228, with the following differences:

- Configure sparse-dense mode for the CE router and PE router interfaces using the `mode` statement at the `[edit protocols pim interface]` hierarchy level. In the example below, the CE-facing interface is `t1-1/0/0:0`.
- Configure the `dense-groups` statement to define the desired group range on the CE router at the `[edit protocols pim]` hierarchy level and in the routing instance at the `[edit routing-instances instance-name protocols pim]` hierarchy level on the PE router.

This section shows how to do the following tasks:

- Configuring PIM on the P Router on page 240
- Configuring PIM on the PE Router on page 240
- Configuring PIM on the CE Router on page 241
- Configuring the Routing Instance on the PE Router on page 241

For information about configuring VPNs, see the *JUNOS VPNs Configuration Guide*.

Configuring PIM on the P Router

Configure PIM on the P router as in the PIM sparse mode example:

```
[edit]
protocols {
  pim {
    dense-groups {
      224.0.1.39/32;
      224.0.1.40/32;
    }
    rp {
      local {
        address 10.255.71.47;
      }
    }
    interface all {
      mode sparse;
      version 2;
    }
    interface fxp0.0 {
      disable;
    }
  }
}
```

Configuring PIM on the PE Router

Configure PIM on the PE router. Use the mode statement at the [edit protocols pim interface] hierarchy level to specify sparse-dense mode.

```
[edit]
protocols {
  pim {
    rp {
      static {
        address 10.255.71.47;
      }
    }
    interface all {
      mode sparse-dense;
      version 2;
    }
    interface fxp0.0 {
      disable;
    }
  }
}
```

Configuring PIM on the CE Router

Configure PIM on the CE router. Use the `dense-groups` statement at the `[edit protocols pim]` hierarchy level to define the desired group range on the CE router. Use the `mode` statement at the `[edit protocols pim interface]` hierarchy level to specify `sparse-dense` mode.

```
[edit]
protocols {
  pim {
    dense-groups {
      229.0.0.0/8;
    }
    rp {
      static {
        address 10.255.245.91;
      }
    }
    interface all {
      mode sparse-dense;
      version 2;
    }
    interface fxp0.0 {
      disable;
    }
  }
}
```

Configuring the Routing Instance on the PE Router

Use the `dense-groups` statement at the `[edit routing-instances instance-name protocols pim]` hierarchy level to define the desired group range for the routing instance on the PE router. Use the `mode` statement at the `[edit routing-instances instance pim interface]` hierarchy level to specify `sparse-dense` mode for interface `t1-1/0/0:0.0`.

```
[edit]
routing-instances {
  VPN-A {
    instance-type vrf;
    interface t1-1/0/0:0.0;
    interface lo0.1;
    route-distinguisher 10.255.71.46:100;
    vrf-import VPN-A-import;
    vrf-export VPN-A-export;
    protocols {
      ospf {
        export bgp-to-ospf;
        area 0.0.0.0 {
          interface t1-1/0/0:0.0;
          interface lo0.1;
        }
      }
    }
    pim {
      dense-groups {
        229.0.0.0/8;
      }
      vpn-group-address 239.1.1.1;
    }
  }
}
```

